

Public goods games on networks and in tumors

Citation for published version (APA):

Bayer, P. (2019). *Public goods games on networks and in tumors*. [Doctoral Thesis, Maastricht University]. Maastricht University. <https://doi.org/10.26481/dis.20190213pb>

Document status and date:

Published: 01/01/2019

DOI:

[10.26481/dis.20190213pb](https://doi.org/10.26481/dis.20190213pb)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Public Goods Games on Networks and in Tumors

DISSERTATION

To obtain the degree of Doctor at
Maastricht University on the authority of the Rector Magnificus
Prof. dr. Rianne M. Letschert
in accordance with the decision of the Board of Deans,
to be defended in public,
on February 13, 2019 at 12:00 hours

by

Péter Bayer

Promotor

Prof. dr. P. Jean-Jacques Herings

Supervisors

Prof. dr. Ronald Peeters

Prof. dr. Frank Thuijsman

Assessment committee

Prof. dr. Dries Vermeulen (chair)

Dr. Yann Bramoullé (Aix-Marseille School of Economics)

Dr. Mark Voorneveld (Stockholm School of Economics)

Prof. dr. Gerhard Weiss

To the memory of István Bayer

Acknowledgement

“I don’t know half of you half as well as I should like; and I like less than half of you half as well as you deserve.”

J.R.R. Tolkien

My time as a PhD student, the last three and a half years, has been some of the best of my life. I had to/chose to divide this time between Maastricht, Budapest, Lausanne, and, to a lesser degree, Marseille and Chicago, and every single instance of the many-many times I traveled I had hated to leave but I was made to feel right at home when I arrived. I had the immense luck and privilege to be surrounded by wonderful people in every one of these places. There is no way to thank every one of them here in a manner that would do them justice, so I mostly limit myself to the Maastricht bunch.

The first and greatest set of thanks is to my supervisors, Ronald Peeters and Frank Thuijsman, for the faith they put in me and the guidance they gave me. I can only hope that I met their expectations at least as well as they met mine.¹ I would also take this opportunity to say that they are most welcome; when I started neither of them were full professors and when I finish they both are, so supervising me must be an excellent boost to one’s career. I extend the same gratitude to my promotor, Jean-Jacques Herings, who went far beyond what his role demanded, not only aiding in the crafting of this dissertation but in filling the gaps in some of its hardest parts. Working with the trio of them was the absolute professional highlight of this period and I can only hope that my future cooperations will be half as fruitful and enjoyable as this one was.

Second, I thank my earliest mentors, Miklós Pintér and Florian Englmaier for being wise enough to recommend Maastricht to me, kind enough to recommend me

¹Which, thanks to a person who will not be named but whose initials consist of the first and the last letters of the alphabet, were quite high at the start already.

to Maastricht, and for a great deal more. The former is singularly responsible for me wanting to pursue an academic career, somehow mustering the patience to teach and guide me at a time when I thought I already knew everything, while the latter helped me find my way in a completely new environment.

Third, I thank my collaborators, Kateřina Staňková, Joel Brown, György Kozics, and Mark Broom. With nothing to take away from my team of supervisors, I've come to think of Joel and Kateřina as foster-supervisors. I thank the two of them for introducing me to their field, for working with me, and for keeping my mind open. I also thank Kateřina for being amazing as the leader of Work Package 4 (the "cancer group") in the 4C Modelling Workshops and arranging my visit to Chicago. I thank Joel for hosting me at Chicago, for our conversations about history, and for being the wisest, kindest person in my vicinity (which should insult no one). I thank Mark for his deft leadership of the 4C Modelling Workshops and for the collaboration we started. Finally, Gyuri, who remained one of my closest friends for over a decade, I thank his conversion from econometrics to game theory with the hope and the promise that he will not regret it.

I thank the two external members of my reading committee, Yann Bramoullé and Mark Voorneveld, for their comments on my papers as well as this dissertation. Additionally, I am grateful to Yann for hosting me in Marseille and to Mark for our conversations in Durham, Lund, and via e-mail. I thank the internal members, Dries Vermeulen and Gerhard Weiss, for their swift evaluation.

I thank the whole AE1 department for providing a friendly, comfortable, and constructive environment for my work. I thank the course coordinators with whom I had the pleasure of performing my teaching duties, Christian Kerckhoffs, Elias Tsakas, Peter Werner, and Matthias Wibrál. Working with them gave me deep insights on how a course should be run and I intend to make good use of my experience with them for my own style of teaching. I thank Elke Lucas for keeping said department together and helping me find my way through the paperwork, and Sylvia Beenen for respecting my preferences in assigning my teaching duties.

I thank two of my old Hungarian friends with whom I did my PhD in the Netherlands in parallel, Mánuel Mágó and Anna Zseleva. Manu has been my best friend and worst in-game ally for as long as I can remember, while Anna is one of the most positive people I know, a great friend to me, as well as a huge help in figuring out what was what in Maastricht.

Emre Ergin was the first friend I made here, a person with whom I share many hobbies, eating well, gaming, history, and a passion for debate; I thank him for all the

meals, games, and conversations we shared (unfortunately he is still wrong in many things, despite my best efforts). Giang Tran is the main reason I went to office instead of working from the quiet comfort of home; I thank her for her kindness and curiosity which inspired me to come out of my shell and have a social life. Li You is one of the most uninhibited and fun people I know; I thank her for exploring Chicago with me and for her kind words of comfort in my time of mourning for my grandfather. Jasmine Maes was the perfect neighbor for me, quiet, helpful, but always ready for playing board games; I thank her for her awkward compliments and for her keeping me company in the long (in Maastricht terms) commutes between home and work. Anastas Tenev I found to be the perfect conversationalist, eloquent, knowledgeable, but opinionated enough to be interesting; I thank him for all our talks, be it about research, philosophy, politics, or, my clear favorite, cheap gossip. I thank you all for making the PhD experience so enjoyable. Additionally, I would like to thank Eveline, Evi, Frauke, Hande, Henrik, Indra, Katharina, Marcello, Monica, Nick, Raginee, Riccardo, and Sena for all the memories we shared in the office, or the various board game events, parties, billiard nights, watching awful movies, and visiting conferences.

Finally, I thank my family for the financial and, more importantly, the emotional support they have given me. I thank my sister Ada for designing the slides of my Science Slam talk as well as the cover of this dissertation, my aunt Évi for being the role model for my career, and my late grandfather István, whose memory this dissertation is dedicated to, for being my role model in many-many things. Finally-finally, I thank my wife Nóri for her strength enduring our times apart, her support in keeping me going, and her love.

Veldwezelt, January 13, 2019

Contents

1	Introduction	1
1.1	Public goods games	1
1.2	Networks	2
1.3	Learning	4
1.4	Public goods and tumors	6
1.5	Structure	9
2	Adaptive learning in weighted network games	13
2.1	Introduction	13
2.2	Weighted network games	16
2.3	Nash equilibria	20
2.4	Cycling of learning processes	26
2.5	Convergence results	32
2.6	Conclusion	37
2.A	Appendix: Proofs	38
3	Myopic and farsighted players in the local public goods game	43
3.1	Introduction	44
3.2	The model	48
3.3	Existence of absorbing effort profiles	53
3.4	Rising inequality between myopic players	59
3.5	Reaching a single effort profile	64
3.6	Lower and upper bounds on long-run payoffs	73
3.7	Conclusion	82
3.A	Toolkit: Buffers	84
3.B	Proofs of the lemmas of Section 3.5	90

3.C	Proofs of Propositions 3.6.2 and 3.6.8	98
4	Best-response dynamics in directed network games	103
4.1	Introduction	103
4.2	The model	106
4.3	Classes of games with acyclical best-response dynamics	110
4.4	Conclusion	115
5	A two-phenotype model of immune evasion by cancer cells	117
5.1	Introduction	118
5.2	The model	120
5.3	Steady-state analysis	124
5.3.1	Nonexistence of steady states in mixed tumors	126
5.3.2	Steady states in selfish tumors	127
5.3.3	Cooperative steady states	130
5.4	Treatment effects	134
5.4.1	Effective treatment strategies	136
5.4.2	Ambiguous effects of treating against cancer crypticity	140
5.5	Conclusion	143
5.A	Appendix	146
	Bibliography	151
	Valorization	159
	Resumé	163

Chapter 1

Introduction

“Science is the great antidote to the poison of enthusiasm and superstition.”

Adam Smith

This introductory chapter provides an informal introduction to the concepts used, presents the motivation of the topics under study, and gives a roadmap for the rest of the dissertation.

1.1 Public goods games

Game theory is the mathematical discipline of decision-making in strategic situations. The minimal requirements needed to describe such situations – games – are the interacting agents – players –, the possible decisions the players can make – strategies –, and the players’ preferences over the outcomes of their decisions. Under some permissive assumptions, these preferences can be represented by numerical (von Neumann-Morgenstern) utility functions – called payoffs. Crucially, the models of game theory allow a player’s payoffs to (partially or fully) depend on the other players’ decisions. Games with only these three components are called normal-form games. Depending on the complexity of the model, games may have several other components describing the timing of the decision-making process, the information available to the players, and the heuristics by which they evaluate information and make their choices. As it is apparent by the number and difficulty of these concepts, game theory is a particularly rich, deep field, with close ties to real-life situations.

Zero-sum games have the property that in every outcome the sum of payoffs of

all players is always zero (or, equivalently, always the same constant). These games capture situations where the players are in direct conflict with each other: one player's gain is another's loss and vice versa. It was the study of two-player zero-sum games by John von Neumann in 1928, titled *Zur Theorie der Gesellschaftsspiele* that is counted as the first contribution to modern game theory.

Conversely, non-zero-sum games allow for mutual gains via coordination and cooperation. Some animals associate into social groups in order to realize the gains from playing non-zero-sum games. Arguably, the defining characteristic of the homo sapiens is the innovative drive to seek out and play non-zero-sum games. Civilization and high economic development is associated with a society's ability to design, create, and maintain institutions that ensure coordination and cooperation in those games.

The **public goods game** is one of the most well-known and most relevant of non-zero sum games. The original version of the game reads as follows: There are $n > 2$ agents with an endowment of 100 euros each. Every agent decides individually and simultaneously on how much of the endowment they wish to contribute to a common pot. Once the decisions are made, the amount of money in the pot is multiplied by a number $1 < m < n$ and redistributed among the agents equally. In this game every euro of a player's contribution increases the wealth of the whole society by $m > 1$ euros, meaning that contributing the entire endowment of 100 euros is the socially optimal choice. However, as every euro of contribution increases the wealth of every player, including the donor's, only by $m/n < 1$ euros, every donor loses $1 - m/n$ euros for each euro of contribution, making the contribution of 0 euros the individually optimal choice. Hence, payoff-maximizing agents will not contribute. This game, together with its generalizations, is one of the most widely studied problems of game theory with applications in every field of economics, as well as in sociology, biology, medicine, computer science, political science, and psychology. Such profound interest in what is a simple tradeoff between self-regard and prosocial behavior is due to the synergy of individual efforts towards the collective good, universally found in every human and animal society.

1.2 Networks

While the depth of the public goods game, especially given its simplicity, makes it an attractive choice to model a variety of situations, its original form cannot capture the intricacies of real societal interactions. There are two points of critique of the original

game. The first is the process by which the societal gain is created. In the original version, the extra money is provided by an external source, i.e. the arbitrator, representing society's joint benefit from the public good created by a centralized institution, i.e. the government. However, value may also be created by private individuals, i.e. the players themselves, in which case the gains are possibly realized by only a subset of players. Second, in real societies, agents are not symmetric with respect to their benefits gained from centrally created public goods either. In fact, many public goods are spatially excludable. For example, an investment into a city's commuter train network benefits the city's residents unequally, those living in the city center will realize lower gains than those living in the metropolitan area. An elegant way to resolve these critical points and construct a more closely applicable model in the spirit of the public goods game is by describing the nature of interactions within the population by the use of networks. In this setting, the players reside in the nodes of the network, and players contribute to create localized goods in the nodes they inhabit. Players gain benefits from the goods they themselves created, as well as the goods to which they are connected in the network.

Networks are a representation of pairwise relations between abstract objects. Depending on what type of objects are considered, we may distinguish social networks, biological networks, computer networks, and many others. By only considering pairs, the number of parameters that define a network is kept low, a polynomial function of the number of objects. Because of this, the framework strikes a balance between richness and efficiency, making networks accurate, elegant, and comprehensive descriptions of complex spatial systems. In this dissertation three types of networks are considered. In order of increasing generality these are simple graphs, weighted networks, and directed networks.

In **simple graphs** relations are binary, two objects may be either linked or not linked. This is the most common framework to capture social networks. For example, scientific collaboration networks are typically imagined as simple graphs with researchers being the nodes and co-authorships being the links. In these models, every link is assigned the same (or similar) meaning, such as friendship, communication, or some other type of connection. While not always an accurate representation of reality, such networks are widely used since they allow for concise graphical representations of even a large number of nodes and links. Chapter 3's model is a game played on a simple graph.

Weighted networks are generalizations of simple graphs, allowing each link to have a unique meaning. This is accomplished by assigning numbers – weights – to

links, usually indicating a stronger relationship in case of a large weight or a weaker one in case of a small weight. Weighted networks constitute a richer model, more closely capturing the details of a web of real relationships, while still allowing for good representations. Depending upon the application the weights may be positive integers, indicating a multiplicity of links, such as multiple joint papers in a scientific collaboration network, positive real numbers, indicating intensity, strength, or capacity such as the value of a collaboration, or both positive and negative real numbers. A possible interpretation is positive numbers indicating good/beneficial/healthy relationships and negative numbers indicating bad/detrimental/toxic relationships, with the absolute value of the weight measuring its strength. A game played on weighted networks is discussed in Chapter 2.

Both simple graphs and weighted networks make the assumption that links go both ways between pairs, and, in case of weighted networks, the weight of a link is equal for both players. **Directed networks**, a generalization of both, is a framework that exists to capture situations which do not adhere to this property. These situations include unequal relationships, where one participant is more important or more powerful than the other, parasitic relationships, from which one party benefits but the other is hurt, and hierarchical relationships, where one party is affected by the other but not the other way around. A game played on directed networks is discussed in Chapter 4.

1.3 Learning

Any analysis of a normal-form game begins with the application of a solution concept. A solution is an ex-ante prediction of the outcome of the game satisfying some intuitive properties. Players may be expected to agree on the solution as a result of some bargaining process, or the solution may be the benchmark at the initiation of the process. Unless the option to cooperate is specified by the rules of the game we typically assume that players have no regard for each other's payoffs and therefore such solutions, especially when viewed as a possible outcome of the game, must be multilaterally self-enforcing. By far the most commonly used and well-known concept is the **Nash equilibrium**. Developed by John Nash in 1950, the Nash equilibrium is a combination of the players' strategies such that no player can be strictly better off in payoffs terms by unilaterally deviating from it. The Nash equilibrium is an intuitive, succinct, and powerful concept, as its existence is proven for every normal-form game (possibly requiring a probabilistic extension of the game). Its predictive potential,

however, is limited.

The main issue of Nash equilibria which hinders its use in real life situations is multiplicity. Even in the simplest case of two players and two strategies the existence of a unique Nash equilibrium is not guaranteed, and even for a handful of players and strategies the number of Nash equilibria may be in the dozens or hundreds. The average number of Nash equilibria in games with 4 players and 6 strategies for each player is above 2000.

This vast multiplicity causes inaccuracy of the predictions on two levels: coordination and selection. Many realistic scenarios feature a lack of coordination mechanisms (e.g. binding agreements), meaning that even a few players are unlikely to choose strategies that correspond to the same Nash equilibrium. In a 2×2 game with two pure (not requiring probabilistic mixing of strategies) Nash equilibria, such as coordination games, both players may justifiably play both of their strategies. This means that all four outcomes may be reached, the likelihood of a successful coordination is 50%, and the Nash equilibrium's predictive power is nil. In case of larger games coordination becomes much harder.

Even if the game rules include ways to achieve coordination, the problem of equilibrium selection persists. The original equilibrium concept gives no insight into which of the dozens, hundreds, or thousands of Nash equilibria should be expected to be played. As these equilibria may lead to very different combinations of payoffs, both coordination and selection remain key obstacles to making plausible and useful predictions. Other solution concepts exist to address these issues, many of them are refinements of the Nash equilibrium, but, just like the original concept itself, they have proven extremely hard to validate in general classes of games.

Learning is the way by which agents accumulate, evaluate, and react to new information. For a wide range of games learning has proven to be highly successful in providing motivation for various equilibrium concepts, and specifically for the Nash equilibrium. In a game featuring learning, the players are playing slightly different versions of the same normal-form game many times. As players discover more about the game and their opponents, they make improvements on their strategies the result of which should, in time, resemble equilibrium play. This approach is attractive from an applied point of view, as, rather than one-shot game with near-simultaneous decision-making, most strategic interactions take place in a repeated setting. The improvements happen as the players receive revision opportunities, the timing of which, and the heuristic by which the revisions themselves are made, depend on the model.

The concept of learning is central to all chapters appearing in this dissertation.

Chapter 2 studies how well the use of Nash equilibria is motivated in public goods games played on weighted networks. Chapter 3 presents a probabilistic learning model in a public goods game played on a simple graph with two types of decision makers, farsighted and myopic. Chapter 4 examines a simple learning process, the best-response dynamic, in public goods games played on directed networks. In the learning models of Chapters 2, 3, and 4, the processes governing the asymptotic behavior of strategies are one-sided, meaning that the revision opportunities are assigned to only one player in each time period, while the strategy of every other player in the player sets remains fixed. In subsequent periods, the identity of the player holding that period's revision opportunity may change, but the set of players remains fixed throughout, irrespective of how the players perform. In Chapter 5's model of the immunotherapy of cancer, the tumor's reaction to various forms of therapy is analyzed. This model takes an evolutionary approach, meaning that players – cells – using strategies yielding high payoffs in the given circumstances will thrive, while those with strategies yielding low payoffs will perish. The modeling choices in the four chapters were made to reflect the plausible ways by which populations in their respective applications learn to play the game. As a general rule of thumb, we can say that short-term economic and biological interactions (compared to the lifespan of the model's subjects), as well as sociological models have a fixed player set with players accumulating information as they learn through trial and error. In long-term economic and biological interactions, in genetic algorithms of computer science, and in intergenerational sociological studies the player set is typically fluctuating with new players arriving to the game, possibly by having more successful players proliferate, as well as living players exiting it, possibly by less successful players dying off. By this process, the population goes through the Darwinian dynamics of evolution via natural selection.

1.4 Public goods and tumors

A **tumor** is a group of body cells which have undergone mutations to an extent that they satisfy five properties not found in normal cells: (1) Self-sufficiency in growth signals; the ability of a cell to reproduce in the absence of hormones or other signals from the body. (2) Insensitivity to anti-growth signals; the resistance to environmental conditions that would halt proliferation in normal cells. (3) Evading apoptosis; apoptosis is a type of programmed cell death occurring in response to, among other factors, genetic instability or infections, which tumor cells have evolved to disobey. (4)

Limitless replicative potential; normal mammalian cells are unable to reproduce after about 60-70 doublings, called the Hayflick limit, which is overcome via mutation in the order of one in every ten million cells, producing an immortalized cell with unlimited doubling capability. (5) Sustained angiogenesis; angiogenesis is the process by which new blood vessels are formed from pre-existing ones, a process which tumor cells are able to hijack in order to ensure a continuous supply of oxygen and nutrients into the tumor. These five properties are present in malignant and benign tumors alike. Malignant tumors satisfy an additional property: (6) Metastasis; the ability to break away from their original site, invade neighboring tissues and organs and spread through the body. These six properties are collectively known as the Original Hallmarks of Cancer. In recent years further hallmarks were added as typical characteristics of malignant tumors, only one of which is touched upon by this dissertation: (7) Immune evasion; the ability to withstand, evade, or suppress attempts by the body's immune system to destroy the tumor.

Cancer is an umbrella term for over 100 diseases of multi-cellular organisms involving malignant tumors. It is one of the leading causes of death in the world, affecting up to 14 million people each year, killing up to 9 million (2015), and costing over 1 trillion USD annually (2010).

In 1971, U.S. President Richard Nixon signed the National Cancer Act, starting an increased research effort into finding a cure for cancer, a project nicknamed the "War on Cancer". In spite of significant research achievements into the causes, prognosis, and new forms of therapy of cancer, the reduction in cancer death rates has remained modest and is mostly attributed to lifestyle improvements and a greater awareness rather than advances in therapy. The war rhetoric, with its positive effects on public awareness notwithstanding, allows the persistence of potentially harmful paradigms. The most widely known paradigm in therapy is that of the Maximum Tolerated Dose (MTD).

Unlike bacteria or viruses, cancer cells are highly similar in structure to the normal, unmutated cells of the human body, meaning that any non-targeted therapy that is able to destroy tumors, such as chemotherapy, is also highly toxic for the patient. MTD is the highest dose of a drug or treatment the patient can receive without causing unacceptable side effects. One of the standards of cancer treatment is to identify the MTD of the given drug, typically chemotherapy, introduce it to the patient, ceasing only when the toxicity caused by the drug becomes as bad as the cancer prognosis. By doing so the treating physician is waging his or her own war with the patient's cancer, and the therapy becomes a race between cancer's resilience to the drug, and the

patient's resilience to the cancer and the drug. The only two outcomes of such therapies are a cancer-free patient, with all detectable tumors eradicated, or a deceased patient. In both outcomes the patient pays a high price both in financial terms and in enduring the painful side effects. While viewing the ongoing efforts to cure cancer as a war both on a macro level (by rushing to develop new, revolutionary forms of treatment) and on a micro level (by trying to completely eradicate all tumors of a patient) is helpful in taking the necessary immediate, decisive, and radical actions, it may obscure situations when treating to kill as many cancer cells as quickly as possible is a suboptimal strategy. This is because, although MTD treatments provide good initial results in shrinking the tumor, the high dose of treatment promotes the evolution of resistance by cancer cells. Due to their defective DNA, cancer cells accumulate mutations with an increased rate, hence adaptation happens rapidly. Since cancer cells without resistance are wiped out by the high dose, the time during which the entire tumor adapts, and therefore the disease remains sensitive to the therapy is minimized by applying the MTD.

The game theory of cancer emerged as a direct response to the war-on-cancer paradigm, treating cancer cells as payoff-maximizing players rather than the enemy. The intention was to model cancer's response to existing therapies, but the field has quickly grown beyond this level. Today, game theoretic papers exist or are in preparation modeling many properties of tumors from the initial evolution of the hallmarks of cancer, to tumor growth and metastasis, to more effective timing of treatments. As any successful society, cancer cells play their own non-zero-sum games, notably the public goods game, between themselves. Examples include the joint efforts by cancer cells to promote angiogenesis and attract nutrients, to systematically destroy stromal structures and thus metastasize into different tissues, and to suppress the immune system. Game theory's insights into these areas are potentially extremely valuable. Thus far the most consequential insight comes from the direct application of evolutionary game theory and it states that the best control for cancer cells resistant to therapy are non-resistant cancer cells. Hence, by lowering the dose of treatment, the time until resistance can be greatly increased while at the same time keeping the cancer burden low enough that the patient does not succumb, rendering cancer into a chronic disease. The lower dose means that this approach is also less costly and less painful for the patient. By applying treatment to control the tumor rather than destroy it, called "adaptive therapy", Moffitt Cancer Center has achieved two extremely promising results: animal trials have shown that prostate cancer may be kept under control indefinitely, while presently ongoing clinical trials with human patients living with advanced prostate cancer have shown that the time until cancer progression and death may be at least

tripled compared to the standard of care. As the clinical trials are still continuing, with new trials currently under development at the time this dissertation is written, the true potentials of adaptive therapy, and other applications of game theory in cancer are still yet to be determined.

Chapter 5 of this dissertation presents a model of the immunotherapy of cancer. Some immune evasion strategies of cancer cells, specifically the tumor’s immunosuppression, are modeled as public goods games played by the cancer cells. Immunotherapy is a form of treatment that induces or enhances immune response. In recent years forms of immunotherapy have proven successful in what appeared to be hopeless cases of cancer, which resulted in widespread mainstream attention. Its general efficacy, however, is still uncertain, as a successful application requires much more tailoring than existing methods of treatment.

1.5 Structure

This dissertation is structured as follows. The introduction is followed by four self-contained chapters. Three of them, Chapters 2, 3, and 4 are related to public goods games played on networks, and one, Chapter 5, is related to the immunotherapy of cancer. The dissertation is concluded by a valorization addendum, exploring the value-creating potential of the research presented in the four chapters.

Chapter 2 examines the convergence of basic one-sided learning processes in public goods games played on weighted networks. It is shown that adaptive learning converges to the set of Nash equilibria of a game as long as two conditions are satisfied. First, all players must receive revision opportunities with some frequency. It does not matter how often or how rarely a player gets a revision opportunity as long as we can specify a length of time such that during any interval of that length every player revises at least once. Second, every revision must take the updating player closer to his contemporary best response. Crucially, we do not impose that players always take the optimal action at the time of their updates, or that they must have a higher payoff after their revisions than before. The condition merely states that the distance from their optimal choice, given the other players’ contributions, must decrease as a result of the update. If a third condition, the finiteness of the set of Nash equilibria, also holds then every such process converges to a single Nash equilibrium point. Finally, the chapter proves that this third condition is generically satisfied.

Chapter 3 builds a model of exploitation with two player types with different abil-

ities. One sophisticated agent plays against a group of myopic agents and is trying to maximize his payoffs by manipulating his neighbors into producing the public good. The players are connected by a simple graph, hence the game involves asymmetries in both ability and position. Sophistication is modeled as a form of foresight, the ability to take future payoffs into account, while the rest of the population's myopia represents a general societal naiveté. As in Chapter 2, the game's dynamics take place in discrete time with one player receiving a revision opportunity in each period, while every other player's actions are kept the same. In each period the player who receives the revision opportunity is assigned randomly, independently from past periods, and with a uniform distribution across players. The chapter shows that an optimal exploitative strategy exists for the farsighted player and it is Markovian, meaning that every choice depends only on the current state of the world rather than the game's entire history until that point. Furthermore, playing optimally guarantees that the game eventually reaches a rest point, at which the myopic players are playing their best-response contributions and the farsighted player has no ways to advance the game such that he receives a higher expected payoff. This result is driven by the fact that in an unweighted setting, the best-responding behavior of myopic players propagates a rising inequality between players, reflecting increased production of players with lower access to the public good, and increased free-riding behavior by players with more access to it. Since the asymptotic behavior of this game in equilibrium is once again converging to the rest points of the system, Chapter 2's own convergence results are shown to be robust with respect to the inclusion of a farsighted player for games played on unweighted graphs. Finally, the consequences of changes in the network are explored. Adding a link between the farsighted player and a previously unlinked myopic player makes the farsighted player better off, while adding a link between two of the farsighted player's previously unlinked myopic neighbors will make the farsighted player worse off.

Chapter 4 considers games played on directed networks. In both Chapters 2 and 3 the reciprocal nature of the interactions, guaranteed in simple graphs and weighted networks, is strongly exploited. Both convergence results rely heavily on the fact that a player moving towards a best response always advances the game towards one of its Nash equilibria, a property derived through the theory of potential games. The results of this strand of literature can only be used in undirected networks. This chapter provides insight for the thus far unexplored, but more general case of directed networks. Examples and situations that lead to either cycling, and therefore non-convergence are shown. Two classes of networks are identified with acyclical best-response dynamics: hierarchical relationships with a clear, unambiguous ranking of players, and networks

where the internal effects of a player's own production to his own payoffs are larger than their external effects on the rest of the population.

Chapter 5 presents an ecological model of immunotherapy based on game theoretical concepts. Two populations are considered, cancer cells and killer T-cells engaged in direct conflict. Cancer cells have two types with respect to their immune evasion strategies, selfish and cooperative. Selfish cells employ strategies that benefit no cancer cell other than itself. These strategies involve the evolution of crypticity by down-regulating tumor-specific antigens on the cells' surface, resulting in a failure by the killer T-cells of the immune system to recognize them as cancerous cells. Cooperative cells on the other hand use strategies that benefit every tumor cell within the neighborhood. Such strategies involve immunosuppression by sending signals to regulatory T-cells that will in turn deactivate killer T-cells. Selfish and cooperative cells may be viewed as defectors and contributors in a public goods game, as being selfish is always the individually optimal strategy – maximizing the proliferative potential of the given cell – while cooperation is the socially optimal one – maximizing the proliferative potential of the tumor. The chapter contains analytical results as well as simulations, considering multiple types of immunotherapy with two main results: (1) Introducing treatment focusing on the selfish phenotype, e.g. by unmasking the cells evolving crypticity and activating the killer T-cells, pushes the tumor towards cooperation. As the cells' private incentive to be selfish diminishes, cooperation inside the tumor will evolve and cancer's immunosuppression will intensify. As a result, the patient's prognosis may worsen by the treatment. (2) Introducing treatment interfering with cancer's immunosuppressive ability has less effect on the cancer's composition, as cancer cells have an incentive to be selfish with and without the therapy. As a result, the tumor's immunosuppression will diminish, and patient prognosis improves by the treatment.

Chapter 2

Adaptive learning in weighted network games

“Learning never exhausts the mind.”

Leonardo da Vinci

This chapter¹ studies adaptive learning in the class of weighted network games. This class of games includes applications like research and development within interlinked firms, crime within social networks, the economics of pollution, and defense expenditures within allied nations. We show that for every weighted network game, the set of pure Nash equilibria is non-empty and, generically, finite. Pairs of players are shown to have jointly profitable deviations from interior Nash equilibria. If all interaction weights are either non-negative or non-positive, then Nash equilibria are Pareto inefficient. We show that quite general learning processes converge to a Nash equilibrium of a weighted network game if every player updates with some regularity.

2.1 Introduction

The theory of learning is of fundamental importance in game theory. With most of the focus in the non-cooperative game theory literature being devoted to the study of equilibria – various concepts, characterizations of the equilibrium set, properties,

¹This chapter is based upon Bayer, Herings, Peeters, and Thuijsman (2017). I would like to thank Péter Bíró, Yann Bramoullé, and Mark Voorneveld for their comments, as well as Sebastian Bervoets and Mathieu Faure for feedback and suggestions.

refinements – it is critical to understand how equilibrium is reached. However, the main concepts of equilibrium theory, and in particular, the concept of Nash equilibrium, have proven difficult to validate, especially in one-shot games. To quote Fudenberg and Levine (1998): “One traditional explanation of equilibrium is that it results from analysis and introspection by the players in a situation where the rules of the game, the rationality of the players, and the players’ payoff functions are common knowledge. Both conceptually and empirically, these theories have many problems.” One of the main goals of learning in game theory is to provide such a motivation. For one-shot games this is typically achieved by interpreting the equilibrium points as results of a series of updates by the players acting in a recurrent setting of that game. These updates are made in response to observed moves by their opponents, with various assumptions on rationality. Ideally, as the players discover more about the game and about their opponents, their collective decisions should, in time, resemble equilibrium play. As such, the learning literature focuses mainly on the stability and convergence properties of various learning processes.

The class of games in which we frame our analysis is the class of weighted network games. This class of games corresponds to the private provision of local public goods games introduced by Bramoullé and Kranton (2007) for undirected graphs and generalized by Bramoullé et al. (2014) for weighted networks. For a comprehensive overview of related models, see Sections 3 and 4 of Jackson and Zenou (2014). The main practical reason this class of games is worth studying, is its wide range of applications in various subfields of economic theory, including R&D within interlinked firms (König et al., 2014), crime within a social network (Ballester et al., 2006), and peer effects with spatial interactions (Blume et al., 2010). Further applications include pollution models as in Leontief (1970) as well as defense expenditures within an international community as studied by Sandler and Hartley (1995) and Sandler and Hartley (2007).

Networks offer a simple way to model complex interactions between many decision makers. The simplest network models are undirected graphs, in which a link between a pair of players indicates a direct interaction. Since players may be indirectly affected by the neighbors of their neighbors, and so on, each interaction may be relevant for each player, resulting in a profoundly rich model. In weighted networks, interaction weights with arbitrary values, either positive or negative, are used to characterize the way that pairs of interacting players influence each other.

The parameters of a weighted network game are the weighted network itself, describing the interactions between the players, a vector of targets that describes the players’ needs, and a vector of upper bounds representing the players’ highest possible

activity levels. Each player has a concave benefit function of the weighted aggregate activity and a linear cost function of his own activity. We show that the set of Nash equilibria of weighted network games is non-empty and generically finite. Additionally, we show that under quite general conditions pairs of players can jointly improve their payoffs, so Nash equilibria are not strong. We also give conditions such that they are not Pareto efficient.

We study a class of learning processes with the following features. The players update their decisions at discrete points in time, maximizing their payoffs for a single period. The updates determine the status quo of the next period. At any given period, only one player is allowed to update, the actions of every other player remain the same. This class of learning processes includes e.g. the improvement paths of Monderer and Shapley (1996).

Weighted network games are generalized aggregative games (Dubey et al., 2006), as well as best-response potential games (Voorneveld, 2000), but may not belong to the class of ordinal potential games (Monderer and Shapley, 1996). Since weighted network games generally do not have an ordinal potential, better-response dynamics may not converge, and we show the possibility of non-convergence by an example.

Our main results concern the properties of adaptive learning processes centered around the best responses. We find that convergence to the set of Nash equilibria requires two conditions: (1) each update has to take the player closer to his contemporary best response, and (2) with some regularity, every player must have the possibility to update. Furthermore, we show that such processes converge to a Nash equilibrium point, if (3) the set of Nash equilibria is finite. The first and second condition concerns the players and may be interpreted as assumptions of cautiousness and activity, respectively. The third condition concerns the parameters of the weighted network game and is generically satisfied. The main significance of our results is in the fact that the convergence of learning processes to a Nash equilibrium can be achieved with relatively weak assumptions on the behavior of the players. Our convergence conditions are less demanding than those studied before, and weakening either of (1) and (2) would lead to non-convergence.

As revision opportunities are generally scarce in real-life applications, the analysis of general, discrete-time learning processes are necessary to provide motivation for the use of Nash equilibria. To our knowledge, the paper on which chapter is based upon is the first to do so in the setting of weighted network games. Bramoullé et al. (2014) and Bervoets and Faure (2019) study best-response dynamics in continuous time. Bervoets et al. (2016) considers a two-stage stochastic learning process with experimenting play-

ers that converges with probability one. Eksin et al. (2012) considers a similar game of incomplete information played on a graph. Unlike the papers examining continuous learning processes, our results make full use of the global properties of the best-response potential and hence we are able to make predictions on processes with erratic movements. These processes may correspond to decision makers who make mistakes either by result of a “trembling hand” or due to a lack of sophistication. Therefore, both our setting and our results are behaviorally motivated.

This chapter is structured as follows. Section 2.2 introduces weighted network games. Section 2.3 contains the characterization of the set of Nash equilibria and its welfare properties. In Section 2.4 we define learning processes and discuss their cycling properties. Section 2.5 contains our main results, the convergence conditions of learning processes. Section 2.6 concludes.

2.2 Weighted network games

Let $I = \{1, \dots, n\}$ denote the set of players with $n \geq 2$. The action set of player $i \in I$ is $X_i = [0, \bar{x}_i]$ with $\bar{x}_i > 0$. Let $x_i \in X_i$ denote player i 's action. The action profile of all players is denoted by $x = (x_j)_{j \in I}$ and the action profile of all players except i by $x_{-i} = (x_j)_{j \neq i}$. Similarly, $X = \prod_{i \in I} X_i$ denotes the set of action profiles and $X_{-i} = \prod_{j \in I \setminus \{i\}} X_j$ the set of action profiles for all players other than i .

Definition 2.2.1. The tuple $G = (I, X, (\pi_i)_{i \in I})$ is called a *weighted network game* if for every $i \in I$ the payoff function $\pi_i: X \rightarrow \mathbb{R}$ is given by:

$$\pi_i(x) = f_i\left(\sum_{j \in I} w_{ij}x_j\right) - c_i x_i,$$

with cost parameters $c_i > 0$, interaction weights $w_{ij} \in \mathbb{R}$, and benefit functions $f_i: \mathbb{R} \rightarrow \mathbb{R}$.

Assumption 2.2.2. For every $i \in I$, $w_{ii} = 1$, and for every $i, j \in I$, $w_{ij} = w_{ji}$. Furthermore, for every $i \in I$, the benefit function f_i is twice continuously differentiable and satisfies the following properties: (1) $f'_i > 0$, (2) $f''_i < 0$, and (3) there exists $t_i \in \mathbb{R}$ such that $f'_i(t_i) - c_i = 0$.

The interpretation is the following. Each player $i \in I$ produces a specialized good using a linear production technology. The costs of producing one unit of the good are equal to c_i . The production of player i is denoted by x_i . Each player consumes his own good as well as a weighted sum of his neighbors' goods. The total amount of consumption

of player i is $\sum_{j \in I} w_{ij} x_j$ and the benefit of consumption is $f_i(\sum_{j \in I} w_{ij} x_j)$. The consumption level t_i where the marginal benefits of consumption are equal to c_i is called the target value. Note that since t_i can be above or below values achievable by using action profiles in X , property (3) of the benefit function is without loss of generality.

For player i , the interaction weight w_{ij} captures the substitutability of one unit of player j 's good to his own. If $w_{ij} > 0$, then player j 's production generates positive externalities for player i . Specifically, if $w_{ij} = 1$, then player i 's enjoyment of player j 's good equals that of his own good. If $w_{ij} \in (0, 1)$, then player i derives less benefits from player j 's good than from his own. If $w_{ij} \in (1, \infty)$, then player i enjoys the good of player j more than that of his own. Negative values of w_{ij} indicate that player j 's production has negative external effects on player i 's benefits, with $w_{ij} \in (0, -1)$, $w_{ij} = -1$, and $w_{ij} \in (-1, -\infty)$ indicating that the negative effects are smaller, equal, or greater in magnitude than the positive effects of equal amounts of the own good. The assumption $w_{ii} = 1$ is a normalization. The symmetry assumption $w_{ij} = w_{ji}$ for $i, j \in I$ is also made in previous studies like Dubey et al. (2006), Bramoullé and Kranton (2007), and Bramoullé et al. (2014). The asymmetric case $w_{ij} \neq w_{ji}$ is relatively unexplored in the local public good setting and hence an interesting direction for future research, see Bourlès et al. (2017) for a model of transfers with asymmetric interactions. The interaction weights are collected in a matrix W .

Example 2.2.3. Let $I = \{1, 2\}$ be a set of two countries that have to decide on the level of their defense expenditures. We take $X_i = [0, Z_i]$, where Z_i denotes the GDP of Country $i \in I$, and

$$W = \begin{pmatrix} 1 & w_{12} \\ w_{12} & 1 \end{pmatrix}.$$

The increasing, concave functions f_1 and f_2 indicate the countries' benefits from defense. Let $t_i = 0.01Z_i$, indicating that both countries have a target value for defense expenditure of 1% of their GDP. This is the amount they would spend on defense if the other nation spends nothing.

If $w_{12} = 0$, then neither country benefits from the other's defense expenditure, nor are they threatened by it. This may indicate neutrality or a significant geographical distance. If $w_{12} > 0$, the two nations are allies and the game becomes a game of strategic substitutes. In this case both nations benefit from the other's defense spending and national defense expenditures are likely to be lower than 1% of GDP. If $w_{12} < 0$, the two nations are hostile to each other, and the game is a game of strategic complements. In this case the nations are hurt or threatened by the other's defense spending and

defense expenditures will likely exceed 1% of GDP.

In case $w_{12} = 1$, Example 2.2.3 results in the 2-player pure public good model of defense expenditure between allies, while $0 < w_{12} < 1$ gives the symmetric version of the limited substitutability public good model of defense expenditure between allies, developed by Sandler and Hartley (2001). They do not consider the case $w_{12} < 0$. In the subsequent section, we discuss how the set of Nash equilibria of this particular game depends on the parameters t_1, t_2 , and w_{12} in more detail.

Our setup allows the modeling of more intricate relationships between players, as illustrated by the following example.

Example 2.2.4. Let $I = \{1, 2, 3\}$ be a set of three countries deciding on the level of their defense expenditures, $X_i = [0, Z_i]$, and

$$W = \begin{pmatrix} 1 & -1 & 1 \\ -1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix},$$

and $t = (0.03Z_1, 0.03Z_2, 0.01Z_3)$. In this example, Countries 1 and 2 are rivals, but both of them are friendly to Country 3. An example of this type of relationship may be that of Israel, Saudi Arabia, and the U.S. since the 2000s.

Example 2.2.4 and similar examples that feature intransitive relationships between countries cannot be modeled in the spirit of Sandler and Hartley (2001). Our setup therefore reflects more closely the possible intricacies of diplomatic relationships and can be used to model any system of alliances and threats, provided that the relationship between any two nations is symmetric.

Games with strategic substitutes and complements are of great relevance in the economic literature. The game of Example 2.2.3 can be interpreted as a game where two firms choose their output to maximize their profits, with the interaction weight w_{12} deciding whether their products are net substitutes or net complements. An interaction matrix similar to the one of Example 2.2.4 may describe the relationship between two competitor firms producing substitute goods, e.g. plane manufacturers Airbus and Boeing, and a third firm producing a complementary good, e.g. a kerosene supplier Exxon Mobil. Other such examples include gaming consoles, Xbox and Playstation with a game developer EA Sports, or tea companies Lipton and Twinings with a sugar company Südzucker. Weighted network games provide a framework to model any type of relationship structure with any number of companies. As it will be apparent later,

our main results extend to Bertrand and Cournot games of differentiated products and linear demand curves, since weighted network games are best-response equivalent to such games.

We denote the set of weighted network games satisfying Assumption 2.2.2 by \mathcal{G} .

Since $w_{ii} = 1$ for every $i \in I$ and $w_{ij} = w_{ji}$ for every $i, j \in I$, the number of free parameters in W is $n(n-1)/2$. Let $w \in \mathbb{R}^{n(n-1)/2}$ denote the column vector of the upper triangular elements of W . We define the set of parameters $P = \mathbb{R}_{++}^n \times \mathbb{R}^{n(n-1)/2} \times \mathbb{R}^n$. Then, for $(\bar{x}, w, t) \in P$, let $\mathcal{G}(\bar{x}, w, t)$ be the set of weighted network games in \mathcal{G} with upper bounds \bar{x} , interaction weights w , and targets t . A weighted network game in such a set is characterized by the benefit functions f_i and cost parameters c_i .

The properties in Assumption 2.2.2 imply that for every $x_{-i} \in X_{-i}$, $\pi_i(x_i, x_{-i})$ has a unique global maximizer in X_i . For player $i \in I$, let $b_i: X \rightarrow X_i$ denote his best-response function, i.e. $b_i(x) = \operatorname{argmax}_{x_i \in X_i} \pi_i(x_i, x_{-i})$ for every $x \in X$. We now show that for a fixed configuration $(\bar{x}, w, t) \in P$, all games in $\mathcal{G}(\bar{x}, w, t)$ are best-response equivalent, which implies that all games in $\mathcal{G}(\bar{x}, w, t)$ have the same set of Nash equilibria.

Lemma 2.2.5. *Let $(\bar{x}, w, t) \in P$ and let $G \in \mathcal{G}(\bar{x}, w, t)$ be a weighted network game. For every $i \in I$ and $x \in X$ it holds that*

$$b_i(x) = \begin{cases} 0 & \text{if } t_i - \sum_{j \neq i} w_{ij} x_j < 0, \\ t_i - \sum_{j \neq i} w_{ij} x_j & \text{if } t_i - \sum_{j \neq i} w_{ij} x_j \in [0, \bar{x}_i], \\ \bar{x}_i & \text{if } t_i - \sum_{j \neq i} w_{ij} x_j > \bar{x}_i. \end{cases} \quad (2.1)$$

Proof. By differentiation of the payoff function we get

$$\frac{\partial \pi_i(x_i, x_{-i})}{\partial x_i} = f'_i\left(\sum_{j \in I} w_{ij} x_j\right) - c_i.$$

The first order condition of unconstrained maximization is satisfied if $f'_i(\sum_{j \in I} w_{ij} x_j) - c_i = 0$. Using property (3) of f_i this is satisfied if $\sum_{j \in I} w_{ij} x_j = t_i$.

If $t_i - \sum_{j \neq i} w_{ij} x_j \in [0, \bar{x}_i]$, then it follows that $b_i(x) = t_i - \sum_{j \neq i} w_{ij} x_j$. Note that the second order condition of maximization is satisfied due to the concavity of f_i , and, therefore, of π_i .

If $t_i - \sum_{j \neq i} w_{ij} x_j < 0$, then for every $x_i \in X_i$ it holds that $t_i < \sum_{j \in I} w_{ij} x_j$. Invoking properties (2) and (3) of f_i , for every $x_i \in X_i$ we have $f'_i(\sum_{j \in I} w_{ij} x_j) < c_i$, meaning that $\partial \pi_i(x_i, x_{-i}) / \partial x_i$ is uniformly negative. Therefore, π_i is maximized for the lowest possible value of x_i , so $b_i(x) = 0$.

Similarly, if $t_i - \sum_{j \neq i} w_{ij} x_j > \bar{x}_i$, then for every $x_i \in X_i$ it holds that $t_i > \sum_{j \in I} w_{ij} x_j$. Properties (2) and (3) of f_i guarantee that for every $x_i \in X_i$ we have $f'_i(\sum_{j \in I} w_{ij} x_j) > c_i$, and that $\partial \pi_i(x_i, x_{-i}) / \partial x_i$ is uniformly positive. Therefore, π_i is maximized for the highest possible value of x_i , so $b_i(x) = \bar{x}_i$. This concludes the proof. \blacksquare

It is useful to define a player's unconstrained best response, the contribution level a player would choose if instead of $[0, \bar{x}_i]$, the set of available actions were equal to \mathbb{R} . For player $i \in I$ and action profile $x \in X$, let $\hat{b}_i(x) = t_i - \sum_{j \neq i} w_{ij} x_j$ denote this value. Clearly, $\hat{b}_i(x) \neq b_i(x)$ implies that $t_i - \sum_{j \neq i} w_{ij} x_j \notin [0, \bar{x}_i]$ and therefore the actual best response is on the boundary: $b_i(x) \in \{0, \bar{x}_i\}$.

For $i \in I$, we define the numbers \underline{b}_i and \bar{b}_i by $\underline{b}_i = \min_{x_{-i} \in X_{-i}} (t_i - \sum_{j \neq i} w_{ij} x_j)$ and $\bar{b}_i = \max_{x_{-i} \in X_{-i}} (t_i - \sum_{j \neq i} w_{ij} x_j)$. Since the set X_{-i} is compact, both \underline{b}_i and \bar{b}_i are well-defined. It is easily seen that the unconstrained best response of player i always belongs to the interval $[\underline{b}_i, \bar{b}_i]$.

Lemma 2.2.5 shows that for every player $i \in I$ and every action profile $x \in X$ such that $b_i(x) \in (0, \bar{x}_i)$, ceteris paribus changing player j 's action by Δx_j changes player i 's best response by $-w_{ij} \Delta x_j$.

2.3 Nash equilibria

Since this chapter's main focus is on the convergence of adaptive learning processes to the set of Nash equilibria, as a precursor we characterize the relevant properties of this set.

We first show that a weighted network game from Definition 2.2.1 satisfying Assumption 2.2.2 is a best-response potential game (Voorneveld, 2000). A game with set of players I , action space X , and payoff functions $(\pi_i)_{i \in I}$ is a best-response potential game if there exists a function $\phi: X \rightarrow \mathbb{R}$ such that for every $i \in I$ and every $x \in X$ it holds that

$$\operatorname{argmax}_{x_i \in X_i} \pi_i(x_i, x_{-i}) = \operatorname{argmax}_{x_i \in X_i} \phi(x_i, x_{-i}). \quad (2.2)$$

We call ϕ the best-response potential of game (I, X, π) . In the following proposition we show that every weighted network game is a best-response potential game and use ϕ to characterize the set of Nash equilibria. For $(\bar{x}, w, t) \in P$, let $X^*(\bar{x}, w, t)$ denote the set of Nash equilibria of a game in $\mathcal{G}(\bar{x}, w, t)$.

Proposition 2.3.1. *For every $(\bar{x}, w, t) \in P$, it holds that:*

1. Every game $G \in \mathcal{G}(\bar{x}, w, t)$ is a best-response potential game with the potential $\phi: X \rightarrow \mathbb{R}$ defined by

$$\phi(x) = x^\top t - \frac{1}{2} x^\top W x, \quad x \in X.$$

2. $X^*(\bar{x}, w, t) \neq \emptyset$.

3. $x^* \in X^*(\bar{x}, w, t)$ if and only if x^* satisfies the Karush-Kuhn-Tucker (KKT) conditions, i.e. for every $i \in I$, there exist $\lambda_i, \mu_i \in \mathbb{R}_+$ such that

$$t_i - \sum_{j \in I} w_{ij} x_j^* + \lambda_i - \mu_i = 0,$$

$$x_i^* \geq 0, \quad \bar{x}_i \geq x_i^*,$$

$$\lambda_i x_i^* = 0, \quad \mu_i (\bar{x}_i - x_i^*) = 0.$$

The proof of Proposition 2.3.1 can be found in the Appendix.

The best-response potential allows for a simple characterization of the set of Nash equilibria as the solution set of a KKT problem. The existence of a Nash equilibrium immediately follows due to the bounded action space. In case the action space is unbounded, negative interaction weights may cause an infinite increase of best replies, leading to nonexistence.

Proposition 2.3.1 is related to Lemma 1 of Bramoullé et al. (2014), which applies to weighted network games for which ϕ is an ordinal potential in the sense of Monderer and Shapley (1996). Weighted network games in our class \mathcal{G} do not generally admit an ordinal potential. In fact, Example 2.4.5 shows that better-response dynamics can lead to cycles in our set-up, which is incompatible with the existence of an ordinal potential. Proposition 2.3.1 therefore only makes the weaker claim that ϕ is a best-response potential.

The existence of a best-response potential is of particular interest for a wide class of games. Their significance in the aggregative/network games strand of literature was first pointed out by Kukushkin (2004). Ewerhart (2017) studies this question in the class of contest games. It is a pivotal step in our analysis as well for its strong implications on the convergence properties of best-response dynamics.

Since weighted network games are also generalized aggregative games (Jensen, 2010), the existence of the best-response potential also follows from more general results in the literature. In particular, it can be shown that the pseudo-potential of Theorem 1 in Dubey et al. (2006) simplifies into the quadratic function ϕ .

In addition to ϕ being a best-response potential, for network games on unweighted graphs, Bervoets and Faure (2019) show that locally, ϕ serves as an ordinal potential, i.e:

$$\operatorname{sgn}\left(\frac{\partial\pi_i(x_i, x_{-i})}{\partial x_i}\right) = \operatorname{sgn}\left(\frac{\partial\phi(x_i, x_{-i})}{\partial x_i}\right).$$

This property can be easily generalized for weighted network games.

In what follows we use the characterization of Proposition 2.3.1 to derive conditions for the interaction matrix and the vector of targets that guarantee the equilibrium set to be finite. For $(\bar{x}, w, t) \in P$, let $\Xi(\bar{x}, w, t)$ denote the set of solutions (x, λ, μ) to the conditions of Proposition 2.3.1. For $H \subseteq I$, let

$$\begin{aligned} \Xi_H(\bar{x}, w, t) = \{ (x, \lambda, \mu) \in \Xi(\bar{x}, w, t) : \forall i \in H, \lambda_i = \mu_i = 0, \text{ and} \\ \forall i \in I \setminus H, \max\{\lambda_i, \mu_i\} > 0 \}. \end{aligned}$$

In words, $\Xi_H(\bar{x}, w, t)$ denotes the set of solutions to the KKT conditions of Proposition 2.3.1 such that for every player in H neither complementarity condition is binding and for every player outside H exactly one complementarity condition is binding. Note that both complementarity conditions cannot be binding simultaneously.

The set $X_H^*(\bar{x}, w, t)$ is obtained by taking the projection of $\Xi_H(\bar{x}, w, t)$ to the set of action profiles X , $X_H^*(\bar{x}, w, t) = \operatorname{proj}_X \Xi_H(\bar{x}, w, t)$, where proj_X is the projection mapping into X . If $x^* \in X_H^*(\bar{x}, w, t)$, then $x^* \in X^*(\bar{x}, w, t)$ and for every $i \in I \setminus H$ we have $x_i^* \in \{0, \bar{x}_i\}$. It follows that $X^*(\bar{x}, w, t) = \bigcup_{H \subseteq I} X_H^*(\bar{x}, w, t)$. Then, clearly, the set of Nash equilibria $X^*(\bar{x}, w, t)$ is finite if and only if for every $H \subseteq I$, $X_H^*(\bar{x}, w, t)$ is finite.

Let bounds $\bar{x} \in \mathbb{R}_{++}^n$, interaction weights $w \in \mathbb{R}^{n(n-1)}$, and a set $H \subseteq I$ be given. The set of target vectors for which $X_H^*(\bar{x}, w, t)$ is infinite is denoted by

$$T_H = \{t \in \mathbb{R}^n : |X_H^*(\bar{x}, w, t)| = \infty\}.$$

The set \bar{T}_H denotes its closure. Further, let

$$T = \{t \in \mathbb{R}^n : |X^*(\bar{x}, w, t)| = \infty\}$$

denote the set of target vectors that yield infinitely many Nash equilibria and let \bar{T} denote the closure of T .

Lemma 2.3.2. *For every $\bar{x} \in \mathbb{R}_{++}^n$, for every $w \in \mathbb{R}^{n(n-1)/2}$, for every $H \subseteq I$, the set \bar{T}_H has Lebesgue measure zero.*

Proof. Let some $\bar{x} \in \mathbb{R}_{++}^n$, some $w \in \mathbb{R}^{n(n-1)/2}$, and some $H \subseteq I$ be given.

First consider the case $H = \emptyset$. Then, for every $t \in \mathbb{R}^n$ it holds that $X_\emptyset^*(\bar{x}, w, t) \subseteq \prod_{i \in I} \{0, \bar{x}_i\}$, meaning that for every $t \in \mathbb{R}^n$ we have $|X_\emptyset^*(\bar{x}, w, t)| < \infty$. It follows that $T_\emptyset = \bar{T}_\emptyset = \emptyset$.

Now consider the case $H \neq \emptyset$. We show that there exists a set $U_H \subset \mathbb{R}^n$ of Lebesgue measure zero such that $\bar{T}_H \subseteq U_H$.

For every $t \in \mathbb{R}^n$, for every $x^* \in X_H^*(\bar{x}, w, t)$, we have that

$$\begin{aligned} x_i^* &\in \{0, \bar{x}_i\}, & i &\in I \setminus H, \\ t_i - \sum_{j \in H} w_{ij} x_j^* - \sum_{j \in I \setminus H} w_{ij} x_j^* &= 0, & i &\in H. \end{aligned}$$

Let $W_H = (w_{ij})_{i,j \in H}$ denote the submatrix of W that we obtain by removing every row and every column whose index is not contained in H . Further, let $W_{H,-H} = (w_{ij})_{i \in H, j \in I \setminus H}$, $t_H = (t_i)_{i \in H}$, $x_H^* = (x_i^*)_{i \in H}$, and $x_{-H}^* = (x_i^*)_{i \in I \setminus H}$. Now, the previous system of equations can be written in matrix form as

$$W_H x_H^* = t_H - W_{H,-H} x_{-H}^*. \quad (2.3)$$

Therefore, by the Rouché-Capelli theorem, $|X_H^*(\bar{x}, w, t)| = \infty$ implies $\text{rank}(W_H) < |H|$. So $T_H = \bar{T}_H = \emptyset$ whenever $\text{rank}(W_H) = |H|$. Consider the case where $\text{rank}(W_H) < |H|$. For $y \in \prod_{i \in I \setminus H} \{0, \bar{x}_i\}$, let U_H^y be the set of target vectors t such that $t_H - W_{H,-H} y$ belongs to the span of W_H . Notice that U_H^y is an $(n - |H| + \text{rank}(W_H))$ -dimensional vector space and therefore a closed set of Lebesgue measure zero. Let $U_H = \bigcup_{y \in \prod_{i \in I \setminus H} \{0, \bar{x}_i\}} U_H^y$. Since U_H is a union of finitely many closed sets of Lebesgue measure zero, it is also closed and is of Lebesgue measure zero. Notice that $t \in \mathbb{R}^n \setminus U_H$ implies that $t \in \mathbb{R}^n \setminus T_H$, since for every $t \in \mathbb{R}^n \setminus U_H$ the system

$$W_H x_H^* = t_H - W_{H,-H} x_{-H}^*$$

has no solutions in x_H^* . It follows that $T_H \subseteq U_H$. Furthermore, since U_H is closed, we also have $\bar{T}_H \subseteq U_H$. ■

We prove Lemma 2.3.2 for the closure of the set T_H , which implies that the set of target vectors with infinitely many Nash equilibria is not only small in a measure theoretic sense, but also in a topological sense.

The intuition behind Lemma 2.3.2 is that for a fixed subset of players, the set of interior Nash equilibria corresponds to the solution set of a linear system which, generically, has only one solution. The case of infinitely many solutions, and hence, the possibility of infinitely many Nash equilibria obtains only if the rank of the interaction

matrix is not full and the target vector belongs to a vector space parallel to the span of the interaction matrix, which is only the case for a set of target vectors of Lebesgue measure zero.

As an illustration of Lemma 2.3.2 for the case where all targets are very large or very small, i.e. for every $i \in I$ we have $t_i > \max_{x \in X} \sum_{j \in I} w_{ij} x_j$ or $t_i < \min_{x \in X} \sum_{j \in I} w_{ij} x_j$, the set of Nash equilibria is a subset of the corners of the strategy space X , i.e. $X^*(\bar{x}, w, t) \subseteq \prod_{i \in I} \{0, \bar{x}_i\}$, and is therefore finite.

Lemma 2.3.3. *For every $\bar{x} \in \mathbb{R}_{++}^n$, for every $w \in \mathbb{R}^{n(n-1)/2}$, the set \bar{T} has Lebesgue measure zero.*

Proof. We show that there exists a set $U \subseteq \mathbb{R}^n$ of Lebesgue measure zero such that $\bar{T} \subseteq U$.

Let $U = \bigcup_{H \subseteq I} U_H$. Since U is a union of finitely many sets of Lebesgue measure zero, it has Lebesgue measure zero. Since $T = \bigcup_{H \subseteq I} T_H$, and $T_H \subseteq U_H$ for every $H \subseteq I$, it also holds that $T \subseteq U$. Once again, since U is closed, we have $\bar{T} \subseteq U$. ■

Corollary 2.3.4. *For every $\bar{x} \in \mathbb{R}_{++}^n$, for every $w \in \mathbb{R}^{n(n-1)/2}$, for almost every $t \in \mathbb{R}^n$, the weighted network game $G \in \mathcal{G}(\bar{x}, w, t)$ has a finite number of Nash equilibria.*

The generic finiteness of the set of Nash equilibria is illustrated in the following example.

Example 2.3.5. Fix parameters \bar{x} and w_{12} in a weighted network game with two players. We have shown in Corollary 2.3.4 that for almost every $t \in \mathbb{R}^2$ the set $X^*(\bar{x}, w_{12}, t)$ is finite. For every $t \in \mathbb{R}^2$, the set $X_\emptyset^*(\bar{x}, w_{12}, t)$ is trivially finite, and it is easy to see that the sets $X_{\{1\}}^*(\bar{x}, w_{12}, t)$ and $X_{\{2\}}^*(\bar{x}, w_{12}, t)$ are finite in the case of two players.

We therefore only check the interior solutions to the KKT problem of this game as defined in Proposition 2.3.1, i.e. where all Lagrange parameters λ_i, μ_i are zero. In an interior solution $x^* \in X_{\{1,2\}}^*(\bar{x}, w_{12}, t)$, we have $\hat{b}_1(x^*) = x_1^*$ and $\hat{b}_2(x^*) = x_2^*$, therefore

$$\begin{aligned} x_1^* &= t_1 - w_{12} x_2^*, \\ x_2^* &= t_2 - w_{12} x_1^*. \end{aligned}$$

In case w_{12} is not equal to 1 or -1 , rearranging yields

$$\begin{aligned} x_1^* &= \frac{t_1 - w_{12} t_2}{1 - (w_{12})^2}, \\ x_2^* &= \frac{t_2 - w_{12} t_1}{1 - (w_{12})^2}. \end{aligned}$$

Therefore, for every $w_{12} \in \mathbb{R} \setminus \{-1, 1\}$, we have $|X_{\{1,2\}}^*(\bar{x}, w_{12}, t)| \leq 1$. Whether or not the set of interior equilibria is empty depends on whether x^* is an element of X .

If $w_{12} = -1$, it is easy to check that $t_1 = t_2 = 0$ yields infinitely many Nash equilibria x^* with $x_1^* = x_2^*$. There can also be infinitely many Nash equilibria when $t_1 + t_2 = 0$. There are no interior Nash equilibria for different values of t as then the system of best responses is inconsistent. Similarly, if $w_{12} = 1$ then there can only be infinitely many Nash equilibria if $t_1 = t_2$. Indeed, if $\bar{x}_1 + \bar{x}_2 > t_1$, then there are infinitely many Nash equilibria x^* with $x_1^* + x_2^* = t_1$, if $\bar{x}_1 + \bar{x}_2 = t_1$, then there is a unique interior Nash equilibrium, and if $\bar{x}_1 + \bar{x}_2 < t_1$, then there are no interior Nash equilibria.

Our result makes use of the generic uniqueness of interior equilibria, but the network structure allows for the existence of a finite number of corner equilibria. Ballester and Calvó-Armengol (2010), Belhaj et al. (2014) and Allouch (2015) provide results for the uniqueness of Nash equilibrium.

We conclude this section by discussing efficiency properties of equilibria. We consider efficiency in the Pareto sense. Other models consider efficiency in terms of minimizing total efforts/production (Bramoullé and Kranton, 2007; Goyal, 2012), or maximizing total welfare (Bramoullé and Kranton, 2007; Helsley and Zenou, 2014). We first show that a pair of players with a non-zero interaction weight can always jointly deviate from an interior equilibrium to a better action profile. For a subset of players $H \subseteq I$ and $\delta \in \mathbb{R}$, let $\delta^H \in \mathbb{R}^n$ denote the vector such that $\delta_i^H = \delta$ for $i \in H$ and $\delta_i^H = 0$ for $i \in I \setminus H$.

Proposition 2.3.6. *Let $(\bar{x}, w, t) \in P$ be given. Let $x^* \in X^*(\bar{x}, w, t)$ be a Nash equilibrium of a game $G = (I, X, \pi) \in \mathcal{G}(\bar{x}, w, t)$ such that $x_i^* \in (0, \bar{x}_i)$ and $x_j^* \in (0, \bar{x}_j)$. It holds that:*

1. *If for some $i, j \in I$ with $i \neq j$ it holds that $w_{ij} \neq 0$, then there exists $\delta \in (0, \min\{x_i^*, x_j^*, \bar{x}_i - x_i^*, \bar{x}_j - x_j^*\})$ such that $\pi_i(x^*) < \pi_i(x^* + \text{sgn}(w_{ij})\delta^{\{i,j\}})$ and $\pi_j(x^*) < \pi_j(x^* + \text{sgn}(w_{ij})\delta^{\{i,j\}})$.*
2. *If $w \geq 0$ or $w \leq 0$ and, for some $i, j \in I$ with $i \neq j$, it holds that $w_{ij} \neq 0$, then there exists $\delta \in (0, \min\{x_i^*, x_j^*, \bar{x}_i - x_i^*, \bar{x}_j - x_j^*\})$ such that the action profile $x^* + \text{sgn}(w_{ij})\delta^{\{i,j\}}$ is a Pareto improvement over x^* .*

The proof of Proposition 2.3.6 can be found in the Appendix.

Proposition 2.3.6 implies that interior Nash equilibria are not strong Nash equilibria since there are profitable deviations by coalitions of two linked players. Moreover, interior Nash equilibria are not Pareto efficient provided that the interaction weights

are either all non-negative or all non-positive. See Elliott and Golub (2019) for a characterization of efficient Nash equilibria in the non-negative case.

2.4 Cycling of learning processes

Within the framework of weighted network games, we consider learning processes where players update their strategies sequentially. That is, given some initial action profile, one player changes his action, while that of every other player remains the same. Then, another player makes a change under similar circumstances, and so on. We improve upon existing results (Kukushkin, 2004; Dubey et al., 2006), which show acyclicity of best-response dynamics by identifying the conditions under which a general class of learning processes (particularly, the better-response dynamics) may cycle. We find that learning processes may cycle if and only if players move away from or radically overshoot the best response. Since the non-existence of best-response cycles is a necessary but not sufficient condition of the convergence of best-response dynamics (Kukushkin, 2015), this is a crucial step towards our convergence results.

Let \mathbb{N} denote the set of positive integers and let $\mathcal{K} = \{\{1, 2\}, \{1, 2, 3\}, \dots\} \cup \mathbb{N}$ be a collection of index sets. For $K \in \mathcal{K}$, we denote by K^- the set that results from K by leaving out its highest element. Notice that K^- is equal to K if $K = \mathbb{N}$.

Definition 2.4.1. Let some $G \in \mathcal{G}$ and $K \in \mathcal{K}$ be given. A sequence of action profiles $(x^k)_{k \in K}$ is a *path* in the game G if:

1. For each $k \in K^-$ there exists a player i^k such that $x_{-i^k}^{k+1} = x_{-i^k}^k$,
2. There is at least one $k \in K^-$ such that $x^{k+1} \neq x^k$.

If $x_{-i^k}^{k+1} = x_{-i^k}^k$, and $x_{i^k}^{k+1} \neq x_{i^k}^k$, then we call i^k the *updating player* at period k .

As per Definition 2.4.1, a path is a sequence where at most one player has changed his contribution between any two successive action profiles, while there are at least two different action profiles in the sequence.

Definition 2.4.2. Let some $G \in \mathcal{G}$ and $K \in \mathcal{K}$ be given. A path $(x^k)_{k \in K}$ is *best-response compatible* in the game G if for every $k \in K^-$ it holds that:

1. If $x^{k+1} = x^k$, then there exists $i^k \in I$ such that $x_{i^k}^{k+1} = x_{i^k}^k = b_{i^k}(x^k)$.
2. If $x_{i^k}^{k+1} \neq x_{i^k}^k$, then $x_{i^k}^{k+1} = b_{i^k}(x^k)$.

Definition 2.4.3. Let some $G \in \mathcal{G}$ and $K \in \mathcal{K}$ be given. A path $(x^k)_{k \in K}$ is *better-response compatible* in game G if for every $k \in K^-$ it holds that:

1. If $x^{k+1} = x^k$, then there exists $i^k \in I$ such that $x_{i^k}^{k+1} = x_{i^k}^k = b_{i^k}(x^k)$.
2. If $x_{i^k}^{k+1} \neq x_{i^k}^k$, then $\pi_{i^k}(x^{k+1}) > \pi_{i^k}(x^k)$.

Definitions 2.4.2 and 2.4.3 capture two of the simplest and best-known learning processes. In case of a best-response compatible path, each updating player moves to his best available option. In case of a better-response compatible path, updating players are only required to strictly improve their payoffs. Clearly, a best-response compatible path is also a better-response compatible path.

Definition 2.4.4. Let some $G \in \mathcal{G}$ and $K = \{1, \dots, m\} \in \mathcal{K}$ be given. A finite path $(x^k)_{k \in K}$ in the game G is a *cycle* if $x^1 = x^m$.

It is well known that best-response dynamics do not produce cycles in best-response potential games (Voorneveld, 2000), which includes weighted network games by Proposition 2.3.1. Better-response dynamics do not generate cycles in ordinal potential games (Monderer and Shapley, 1996). The following example shows that better-response cycles can occur within weighted network games.

Example 2.4.5. Let $I = \{1, 2\}$, $X_1 = X_2 = [0, 4]$, and $t_1 = t_2 = 1$. Moreover, let the payoff functions be given by

$$\pi_1(x_1, x_2) = 2\sqrt{x_1 + 0.6x_2} - x_1$$

and

$$\pi_2(x_1, x_2) = 2\sqrt{x_2 + 0.6x_1} - x_2.$$

It is easy to check that π_1 and π_2 satisfy the properties laid down in Definition 2.2.1 and Assumption 2.2.2 with $w_{12} = 0.6$, $f_1(z) = f_2(z) = 2\sqrt{z}$, and $c_1 = c_2 = 1$.

Table 2.1 presents a sequence of action profiles that constitutes a better-response cycle for this example.

Note that the changes in player 1's choice of actions between periods 1 and 2 and between periods 3 and 4, as well as those for player 2 between periods 4 and 5 and between periods 6 and 7 are quite large, given the action space. Columns 5 and 6 of Table 2.2 present for each period the distance between the current action and both the best response and the action chosen by the player updating his action. Notice that in periods 1 and 4, the actions chosen are more than twice as far away from the current

k	x_1^k	x_2^k	$\pi_1(x^k)$	$\pi_2(x^k)$
1	0	0.1	0.49	0.53
2	3	0.1	0.50	2.66
3	3	0	0.46	2.68
4	0.1	0	0.53	0.49
5	0.1	3	2.66	0.50
6	0	3	2.68	0.46
7	0	0.1	0.49	0.53

Table 2.1: Actions played and payoffs in the better-response cycle in Example 2.4.5.

action than the best response is, meaning that the updating player, despite the increase in payoffs, has moved farther from his optimal decision than he originally was. We refer to this as extreme overshooting beyond the best response. This chapter's main result is showing that a lack of extreme overshooting is a sufficient and necessary condition of convergence in the class of learning processes centered in the best responses.

k	x_1^k	x_2^k	$b_{i^k}(x^k)$	$ b_{i^k}(x^k) - x_{i^k}^k $	$ x_{i^k}^{k+1} - x_{i^k}^k $	α_k
1	0	0.1	0.94	0.94	3	-2.19
2	3	0.1	0	0.1	0.1	0
3	3	0	1	2	2.9	-0.45
4	0.1	0	0.94	0.94	3	-2.19
5	0.1	3	0	0.1	0.1	0
6	0	3	1	2	2.9	-0.45
7	0	0.1				

Table 2.2: The size of action changes in the better-response cycle of Table 2.1.

As before, for a path $(x^k)_{k \in K}$, let $(i^k)_{k \in K^-}$ denote the updating player in period k if there was a change in the action profile and let it denote any other player if there was not. Furthermore, for $k \in K^-$, let the overshooting coefficient $\alpha_k \in \mathbb{R} \cup \{-\infty, \infty\}$ be defined as

$$\alpha_k = \frac{x_{i^k}^{k+1} - b_{i^k}(x^k)}{x_{i^k}^k - b_{i^k}(x^k)},$$

where we take the convention that in case the denominator is 0, $\alpha_k = -\infty$ if the numerator is negative, $\alpha_k = 0$ if the numerator is 0, and $\alpha_k = +\infty$ if the numerator is

positive. Column 7 of Table 2.2 shows the values of α_k in the better-response cycle of Example 2.4.5.

The coefficient α_k determines the extent of overshooting of the updating player beyond the best response. If there is overshooting, then α_k is negative. There is no overshooting if α_k is positive. If $\alpha_k = 0$ then the updating player moved to the best response. If $\alpha_k \in \{-\infty, \infty\}$ then $x_{i^k}^{k+1} \neq x_{i^k}^k = b_{i^k}(x^k)$, so the player moved away from a best response. If $\alpha_k < -1$, then the new action is farther from the best response relative to the action before the update.

Values of α_k in $(0, 1)$ correspond to a better response, while in case $\alpha_k > 1$ the payoff of the updating player is lower than before. For negative values of α_k , the threshold between better and worse replies depends on the payoff function. Naturally, the possible values that α_k may take depend on \bar{x} .

As suggested by Example 2.4.5, sequences of action profiles that feature extreme overshooting beyond the best response may cycle. We therefore characterize sequences by their extent of overshooting.

Definition 2.4.6. A path $(x^k)_{k \in K}$ in a game $G \in \mathcal{G}$ is α -centered for some $\alpha > 0$ if for every $k \in K^-$ it holds that $|\alpha_k| < \alpha$.

A best-response compatible path is α -centered for every $\alpha > 0$. Furthermore, for every $\alpha > 0$ there exist paths that are better-response compatible and α -centered, but are not best-response compatible. For instance, it is easy to see that every sequence $(\alpha_k)_{k \in K^-}$ such that for every $k \in K^-$, $\alpha_k \in [0, \min\{\frac{\alpha}{2}, \frac{1}{2}\}]$, is both α -centered and better-response compatible.

The restriction of being α -centered on a better-response dynamic captures a form of cautiousness by the players, as they do not engage in updates that take them very far from their optimal choice. For finite values of α , players do not change their action in an α -centered path if they are at their best response, as that would imply $|\alpha_k| = \infty$. Alternatively, α can be thought of as a measure of inaccuracy of best-responding players who are unable to execute their intended strategies.

For the remainder of this chapter we mainly consider α -centered paths with $\alpha \in (0, 1)$. In these paths, every updating player moves closer to his current best response.

We define the overshooting coefficient $\hat{\alpha}_k$ similar to α_k , replacing the best-response function b with the unconstrained best-response function \hat{b} . For a path of action profiles $(x^k)_{k \in K}$, we define

$$\hat{\alpha}_k = \frac{x_{i^k}^{k+1} - \hat{b}_{i^k}(x^k)}{x_{i^k}^k - \hat{b}_{i^k}(x^k)}, \quad k \in K^-.$$

The relationship between α_k and $\hat{\alpha}_k$ is summarized in the following lemma.

Lemma 2.4.7. *Let $(x^k)_{k \in K}$ be a path of action profiles in a game $G \in \mathcal{G}$. The following statements hold for every $k \in K^-$:*

- (i) $\alpha_k \neq \hat{\alpha}_k$ implies $b_{i^k}(x^k) \in \{0, \bar{x}_{i^k}\}$.
- (ii) $\alpha_k \in (0, 1)$ implies $\hat{\alpha}_k \in (0, 1)$.
- (iii) $\alpha_k \in (-1, 0)$ implies $\alpha_k = \hat{\alpha}_k$.
- (iv) $\alpha_k = 0$ implies $0 \leq \hat{\alpha}_k \leq 1$.
- (v) $\hat{\alpha}_k = 1$ implies $x^{k+1} = x^k$.

Proof. (i). If $b_{i^k}(x^k) \in (0, \bar{x}_{i^k})$, then it holds that $b_{i^k}(x^k) = \hat{b}_{i^k}(x^k)$ and thus $\alpha_k = \hat{\alpha}_k$.
(ii). We only need to consider the case $\alpha_k \neq \hat{\alpha}_k$. By (i) we have $b_{i^k}(x^k) \in \{0, \bar{x}_{i^k}\}$. Take the case $b_{i^k}(x^k) = 0$. Then it holds that $\hat{b}_{i^k}(x^k) < 0$, so $x_{i^k}^{k+1} < x_{i^k}^k$ due to $0 < \alpha_k < 1$, and thus

$$0 < \alpha_k = \frac{x_{i^k}^{k+1}}{x_{i^k}^k} < \frac{x_{i^k}^{k+1} - \hat{b}_{i^k}(x^k)}{x_{i^k}^k - \hat{b}_{i^k}(x^k)} = \hat{\alpha}_k < 1.$$

The case $b_{i^k}(x^k) = \bar{x}_{i^k}$ follows from similar arguments.

(iii). Since $\alpha_k \in (-1, 0)$, we have

$$\text{sgn}(x_{i^k}^k - b_{i^k}(x^k)) = -\text{sgn}(x_{i^k}^{k+1} - b_{i^k}(x^k)) \neq 0.$$

Therefore, it must hold that $b_{i^k}(x^k) \in (0, \bar{x}_{i^k})$, otherwise $x_{i^k}^{k+1}$ would not be in X_{i^k} . It follows that $b_{i^k}(x^k) = \hat{b}_{i^k}(x^k)$.

(iv). Once again, we only need to discuss the case $\alpha_k \neq \hat{\alpha}_k$, so $b_{i^k}(x^k) \in \{0, \bar{x}_{i^k}\}$. Consider the case $b_{i^k}(x^k) = 0$. We have that $\hat{b}_{i^k}(x^k) < 0$ and $x_{i^k}^{k+1} = 0$ since $\alpha_k = 0$. It holds that

$$0 = \alpha_k \leq \frac{-\hat{b}_{i^k}(x^k)}{x_{i^k}^k - \hat{b}_{i^k}(x^k)} = \hat{\alpha}_k \leq 1.$$

The case $b_{i^k}(x^k) = \bar{x}_{i^k}$ follows from similar arguments.

(v). In case $\hat{\alpha}_k = 1$, we have

$$x_{i^k}^{k+1} - \hat{b}_{i^k}(x^k) = x_{i^k}^k - \hat{b}_{i^k}(x^k),$$

so $x_{i^k}^{k+1} = x_{i^k}^k$. ■

In the following proposition we show the relation between the value of $\hat{\alpha}_k$ and changes in the value of the potential as defined in Proposition 2.3.1. This relationship will prove crucial in our convergence analysis.

Proposition 2.4.8. *Let a game $G \in \mathcal{G}$ be given.*

1. *Let $(x^k)_{k \in K^-}$ be a path of action profiles such that for every $k \in K^-$, $\hat{\alpha}_k \in \mathbb{R}$. Then it holds that*

$$\phi(x^{k+1}) - \phi(x^k) = \frac{1}{2}(1 - \hat{\alpha}_k)(1 + \hat{\alpha}_k)(\hat{b}_{i^k}(x^k) - x_{i^k}^k)^2, \quad k \in K^-.$$

2. *The game has no 1-centered cycles.*

Proof. 1. Using the definition of ϕ gives

$$\phi(x^{k+1}) - \phi(x^k) = \sum_{i \in I} (x_i^{k+1} - x_i^k) t_i - \frac{1}{2} \sum_{i \in I} x_i^{k+1} \left(\sum_{j \in I} w_{ij} x_j^{k+1} \right) + \frac{1}{2} \sum_{i \in I} x_i^k \left(\sum_{j \in I} w_{ij} x_j^k \right).$$

Using the symmetry of the interaction matrix W and taking advantage of the fact that x^{k+1} is the successor of x^k in a path, we substitute $x_{-i^k}^k = x_{-i^k}^{k+1}$ to get

$$\phi(x^{k+1}) - \phi(x^k) = (x_{i^k}^{k+1} - x_{i^k}^k) t_i - \frac{1}{2} ((x_{i^k}^{k+1})^2 - (x_{i^k}^k)^2) - (x_{i^k}^{k+1} - x_{i^k}^k) \sum_{j \neq i^k} w_{i^k j} x_j^k.$$

Factoring out $x_{i^k}^{k+1} - x_{i^k}^k$ yields

$$\phi(x^{k+1}) - \phi(x^k) = (x_{i^k}^{k+1} - x_{i^k}^k) \left[(t_{i^k} - \frac{1}{2}(x_{i^k}^{k+1} + x_{i^k}^k)) - \sum_{j \neq i^k} w_{i^k j} x_j^k \right].$$

Substituting $\hat{b}_{i^k}(x^k) = t_{i^k} - \sum_{j \neq i^k} w_{i^k j} x_j^k$ gives

$$\begin{aligned} \phi(x^{k+1}) - \phi(x^k) &= (x_{i^k}^{k+1} - x_{i^k}^k) (\hat{b}_{i^k}(x^k) - \frac{1}{2}(x_{i^k}^{k+1} + x_{i^k}^k)) \\ &= \frac{1}{2} (x_{i^k}^{k+1} - x_{i^k}^k) [\hat{b}_{i^k}(x^k) - x_{i^k}^{k+1} + \hat{b}_{i^k}(x^k) - x_{i^k}^k]. \end{aligned}$$

Finally, substituting $x_{i^k}^{k+1} = (1 - \hat{\alpha}_k) \hat{b}_{i^k}(x^k) + \hat{\alpha}_k x_{i^k}^k$ and $\hat{b}_{i^k}(x^k) - x_{i^k}^{k+1} = \hat{\alpha}_k (\hat{b}_{i^k}(x^k) - x_{i^k}^k)$ gives

$$\phi(x^{k+1}) - \phi(x^k) = \frac{1}{2} (1 - \hat{\alpha}_k) (1 + \hat{\alpha}_k) (\hat{b}_{i^k}(x^k) - x_{i^k}^k)^2.$$

Suppose that $(x^k)_{k \in K^-}$ is a 1-centered cycle. For every $k \in K^-$ it holds by Lemma 2.4.7 that $\hat{\alpha}_k \in (-1, 1]$ and therefore by the first statement we have $\phi(x^{k+1}) - \phi(x^k) \geq 0$.

By Definition 2.4.1, each path has at least one pair of successive action profiles that are different. Let $k' \in K^-$ be such that $x^{k'+1} \neq x^{k'}$. Since the path $(x^k)_{k \in K^-}$ is 1-centered, Lemma 2.4.7 implies $|\hat{\alpha}_{k'}| < 1$, and therefore by the first statement we have $\phi(x^{k'+1}) - \phi(x^{k'}) > 0$. Together with the fact that $\phi(x^{k+1}) - \phi(x^k) \geq 0$ for every $k \in K^-$, we obtain a contradiction to $(x^k)_{k \in K^-}$ being a 1-centered cycle. \blacksquare

Proposition 2.4.8 says that in a path of action profiles, the change of the potential is only determined by the magnitude of $\hat{\alpha}_k$. Each time the updating player gets closer to his unconstrained best response by his update, the value of the potential increases, and each time he gets further from the unconstrained best response, the value of the potential decreases. By this property we show that if every update moves the updating player closer to his best response, then better-response cycles cannot exist. This also implies the non-existence of best-response cycles. Notice that the cycle in Example 2.4.5 is not 1-centered, hence Proposition 2.4.8 is not applicable. Furthermore, notice that for $\alpha > 1$, cycling is possible in an α -centered path. For example, if we have $i^1 = i^2$ and $\alpha_1 = \alpha_2 = -1$, then (x^1, x^2, x^3) constitutes a cycle. This means that $\alpha \leq 1$ is a necessary and sufficient condition for the non-existence of α -centered cycles.

2.5 Convergence results

With the cycling behavior of the learning processes established we present the main results of this chapter, the convergence of a general class of learning processes.

Clearly, processes that are not 1-centered will not converge in general, although there are examples of processes and networks where this is possible, e.g. the better-response dynamic in case of the empty network.

Intuition suggests that for 1-centered processes, where each player moves closer to his best response, convergence is guaranteed. However, a player may get farther away from his best response through the updates of the other players, meaning that subsequent updates for any given player are not necessarily smaller in magnitude than previous ones. As a first step to establish convergence, we show that the distance between consecutive elements of any α -centered path with $\alpha < 1$ converges to zero and that the distance between the current action and the best response to it approaches zero for an updating player.

Proposition 2.5.1. *Let $(x^k)_{k \in \mathbb{N}}$ be an α -centered path in a game $G \in \mathcal{G}$ such that $\alpha < 1$. It holds that:*

1. $\lim_{k \rightarrow \infty} \|x^{k+1} - x^k\|_2 = 0.$
2. $\lim_{k \rightarrow \infty} |b_{i^k}(x^k) - x_{i^k}^k| = 0.$

Proof. 1. We use the fact that $x_{i^k}^{k+1} = (1 - \hat{\alpha}_k)\hat{b}_{i^k}(x^k) + \hat{\alpha}_k x_{i^k}^k$ to obtain

$$\|x^{k+1} - x^k\|_2^2 = (x_{i^k}^{k+1} - x_{i^k}^k)^2 = (1 - \hat{\alpha}_k)^2 (\hat{b}_{i^k}(x^k) - x_{i^k}^k)^2.$$

Applying Proposition 2.4.8 gives

$$\|x^{k+1} - x^k\|_2^2 = 2 \frac{1 - \hat{\alpha}_k}{1 + \hat{\alpha}_k} (\phi(x^{k+1}) - \phi(x^k)).$$

Since the path is α -centered with $\alpha < 1$, by Lemma 2.4.7 we have $-\alpha < \hat{\alpha}_k \leq 1$. It follows that

$$\|x^{k+1} - x^k\|_2^2 \leq 2 \frac{1 + \alpha}{1 - \alpha} (\phi(x^{k+1}) - \phi(x^k)). \quad (2.4)$$

By the first part of Proposition 2.4.8 we have that the sequence $(\phi(x^k))_{k \in \mathbb{N}}$ is monotonically increasing. Furthermore, since ϕ is continuous and the set X is compact, the sequence $(\phi(x^k))_{k \in \mathbb{N}}$ is also bounded, and hence it is convergent, so $\phi(x^{k+1}) - \phi(x^k) \rightarrow 0$ as $k \rightarrow \infty$. Since the right-hand side of (2.4) converges to zero, it follows that $\|x^{k+1} - x^k\|_2^2 \rightarrow 0$ as $k \rightarrow \infty$. This implies the statement.

2. Suppose it does not hold that $\lim_{k \rightarrow \infty} |b_{i^k}(x^k) - x_{i^k}^k| = 0$. Then the sequence $(i^k, x^k)_{k \in \mathbb{N}}$ has a converging subsequence (i^{k^ℓ}, x^{k^ℓ}) with limit (i, x) such that $|b_i(x) - x_i| = \varepsilon > 0$. We distinguish three cases: (a) $\hat{b}_i(x) \in (0, \bar{x}_i)$, (b) $\hat{b}_i(x) \leq 0$, and (c) $\hat{b}_i(x) \geq \bar{x}_i$.

Case (a). $\hat{b}_i(x) \in (0, \bar{x}_i)$.

There is $\ell' \in \mathbb{N}$ such that, for every $\ell \geq \ell'$, $i^{k^\ell} = i$, $b_i(x^{k^\ell}) \in (0, \bar{x}_i)$, and $|b_i(x^{k^\ell}) - x_i^{k^\ell}| \geq \varepsilon/2$. It follows that $\hat{\alpha}_{k^\ell} = \alpha_{k^\ell}$, so by Proposition 2.4.8,

$$\phi(x^{k^\ell+1}) - \phi(x^{k^\ell}) \geq \frac{1}{2}(1 - \alpha_{k^\ell})(1 + \alpha_{k^\ell})\frac{1}{4}\varepsilon^2 > \frac{1}{2}(1 - \alpha)(1 + \alpha)\frac{1}{4}\varepsilon^2, \quad \ell \geq \ell'. \quad (2.5)$$

By Proposition 2.4.8, we have that the sequence $(\phi(x^k))_{k \in \mathbb{N}}$ is monotonically increasing, so the subsequence $(\phi(x^{k^\ell}))_{\ell \in \mathbb{N}}$ is monotonically increasing, and by (2.5) it tends to infinity. This contradicts the fact that the continuous function ϕ has a maximum on the compact set X .

Case (b). $\hat{b}_i(x) \leq 0$.

We have that $b_i(x) = 0$ and $x_i = \varepsilon$. There is $\ell' \in \mathbb{N}$ such that, for every $\ell \geq \ell'$, $i^{k^\ell} = i$, $b_i(x^{k^\ell}) \leq x_i^{k^\ell}$, and $|b_i(x^{k^\ell}) - x_i^{k^\ell}| \geq \varepsilon/2$. If $\hat{b}_i(x^{k^\ell}) \geq 0$, then $\hat{\alpha}_{k^\ell} = \alpha_{k^\ell}$. Otherwise, we have $\hat{b}_i(x^{k^\ell}) < 0$, so $b_i(x^{k^\ell}) = 0$, and

$$0 \leq \hat{\alpha}_{k^\ell} = \frac{x_i^{k^\ell+1} - \hat{b}_i(x^{k^\ell})}{x_i^{k^\ell} - \hat{b}_i(x^{k^\ell})} \leq \frac{x_i^{k^\ell+1} - \underline{b}_i}{x_i^{k^\ell} - \underline{b}_i} = \frac{\alpha_{k^\ell} x_i^{k^\ell} - \underline{b}_i}{x_i^{k^\ell} - \underline{b}_i} \leq \frac{\frac{1}{2}\alpha_{k^\ell}\varepsilon - \underline{b}_i}{\frac{1}{2}\varepsilon - \underline{b}_i} < \frac{\frac{1}{2}\alpha\varepsilon - \underline{b}_i}{\frac{1}{2}\varepsilon - \underline{b}_i}. \quad (2.6)$$

The right-hand side of (2.6), denoted by β , belongs to $(\alpha, 1)$, so it holds that

$$-\alpha < \hat{\alpha}_{k^\ell} \leq \beta, \quad \ell \geq \ell'.$$

By Proposition 2.4.8, we have that

$$\phi(x^{k^\ell+1}) - \phi(x^{k^\ell}) \geq \frac{1}{2}(1 - \alpha_{k^\ell})(1 + \alpha_{k^\ell})\frac{1}{4}\varepsilon^2 > \frac{1}{2}(1 - \beta)(1 + \beta)\frac{1}{4}\varepsilon^2, \quad \ell \geq \ell'. \quad (2.7)$$

By Proposition 2.4.8, we have that the sequence $(\phi(x^k))_{k \in \mathbb{N}}$ is monotonically increasing, so the subsequence $(\phi(x^{k^\ell}))_{\ell \in \mathbb{N}}$ is monotonically increasing, and by (2.7) it tends to infinity. This contradicts the fact that the continuous function ϕ has a maximum on the compact set X .

Case (c). $\hat{b}_i(x) \geq \bar{x}_i$.

We can derive a contradiction along similar lines as in Case (b).

Since all three cases lead to a contradiction, we conclude that $\lim_{k \rightarrow \infty} |b_{ik}(x^k) - x_{ik}^k| = 0$. ■

The first statement of Proposition 2.5.1 follows from the monotonicity and therefore the convergence of the values of the potential along an α -centered path, by applying Proposition 2.4.8 to translate differences in the value of the potential to distances between action profiles. The second statement shows convergence to the best response for all updating players. In order to achieve convergence to a Nash equilibrium, we need convergence to the best response for all players. This can only be achieved if all players update regularly, otherwise nothing guarantees convergence for a player who, for instance, never updates. We therefore define the notion of updating in every ℓ periods, which is going to be the final condition for our main result.

Definition 2.5.2. Player $i \in I$ updates in every ℓ periods in a path of action profiles $(x^k)_{k \in \mathbb{N}}$ in a game $G \in \mathcal{G}$ if for every $k \in \mathbb{N}$ there exists $k' \in \{k, \dots, k + \ell - 1\}$ such that either $[x_i^{k'} \neq x_i^{k'+1}]$ or $[x_i^{k'} = x_i^{k'+1} \text{ and } x_i^{k'} = x_i^{k'+1} = b_i(x^{k'})]$.

A player satisfies Definition 2.5.2 if in every length ℓ segment of the path there is an action profile at which he updated or there exists a pair of successive action profiles that are identical and the player is at his best response. This condition is quite general, e.g. if $\ell > n$, then it includes the possibility of fixing the order of updates across the players – provided that each player gets a turn.

We are ready to present our main results.

Theorem 2.5.3. Let $(\bar{x}, w, t) \in P$ and let $(x^k)_{k \in \mathbb{N}}$ be an α -centered path in a game $G \in \mathcal{G}(\bar{x}, w, t)$. If $\alpha < 1$ and every player updates in every ℓ periods for some $\ell \in \mathbb{N}$, then every cluster point of $(x^k)_{k \in \mathbb{N}}$ belongs to $X^*(\bar{x}, w, t)$. If, in addition, we have that $|X^*(\bar{x}, w, t)| < \infty$, then there exists $x^* \in X^*(\bar{x}, w, t)$ such that $\lim_{k \rightarrow \infty} x^k = x^*$.

Proof. We first show the first statement. Since every linear function is Lipschitz continuous, the function $b_i: X \rightarrow X_i$ is Lipschitz continuous for every $i \in I$. Denote the Lipschitz constant of b_i by L_i .

Let x be a cluster point of $(x^k)_{k \in \mathbb{N}}$. We prove the result by showing that, for every $i \in I$, for every $\varepsilon > 0$, $|b_i(x) - x_i| < \varepsilon$. We need a technical lemma.

Lemma 2.5.4. *Let $(x^k)_{k \in \mathbb{N}}$ be a path of action profiles in a game $G \in \mathcal{G}$ such that $\lim_{k \rightarrow \infty} \|x^{k+1} - x^k\|_2 = 0$. For every $\varepsilon > 0$, for every $\ell \in \mathbb{N}$, there exists $M \in \mathbb{N}$ such that, for every $m > M$, for every $k \in \{m, \dots, m + \ell - 1\}$, we have $\|x^k - x^m\|_2 < \varepsilon$.*

The proof of this lemma can be found in the Appendix. Let $i \in I$ and $\varepsilon > 0$ be given.

Let $M_1 \in \mathbb{N}$ be such that, for every $m > M_1$, for every $k \in \{m, \dots, m + \ell - 1\}$, we have $\|x^k - x^m\|_2 < \varepsilon/(3 + 2L_i)$. Lemma 2.5.4 guarantees the existence of such an M_1 .

Let $M_2 \in \mathbb{N}$ be such that for every $m > M_2$ it holds that $|b_{i^m}(x^m) - x_{i^m}^m| < \varepsilon/(3 + 2L_i)$. The second statement of Proposition 2.5.1 guarantees the existence of such an M_2 .

Let $m > \max\{M_1, M_2\}$ be such that $\|x^m - x\|_2 < \varepsilon/(3 + 2L_i)$. Such an m must exist, since x is a cluster point of the sequence $(x^k)_{k \in \mathbb{N}}$.

If player i updates in every ℓ periods, then there exists $k' \in \{m, \dots, m + \ell - 1\}$ such that $|b_i(x^{k'}) - x_i^{k'}| < \varepsilon/(3 + 2L_i)$, where we use that $m > M_2$. Since $m > M_1$ as well, it holds that $\|x^{k'} - x^m\|_2 < \varepsilon/(3 + 2L_i)$, and by the choice of m we have $\|x - x^m\|_2 < \varepsilon/(3 + 2L_i)$. In particular, it follows that $|x_i^{k'} - x_i^m| < \varepsilon/(3 + 2L_i)$ and $|x_i^m - x_i| < \varepsilon/(3 + 2L_i)$. By the triangle inequality we get

$$|b_i(x^{k'}) - x_i| \leq |b_i(x^{k'}) - x_i^{k'}| + |x_i^{k'} - x_i^m| + |x_i^m - x_i| < \frac{3\varepsilon}{3 + 2L_i}.$$

Also, $\|x - x^m\|_2 < \varepsilon/(3 + 2L_i)$ and $\|x^m - x^{k'}\|_2 < \varepsilon/(3 + 2L_i)$ imply that $\|x - x^{k'}\|_2 < 2\varepsilon/(3 + 2L_i)$. Using the Lipschitz continuity of b_i , we get

$$|b_i(x) - b_i(x^{k'})| < \frac{2L_i\varepsilon}{3 + 2L_i}.$$

Summing up, we have

$$|b_i(x) - x_i| \leq |b_i(x) - b_i(x^{k'})| + |b_i(x^{k'}) - x_i| < \frac{2L_i\varepsilon}{3 + 2L_i} + \frac{3\varepsilon}{3 + 2L_i} = \varepsilon.$$

This concludes the first part of the proof.

For the second part let Y denote the non-empty set of cluster points of $(x^k)_{k \in \mathbb{N}}$. The first statement implies that every element of Y is a Nash equilibrium. We therefore

only have to show that Y is a singleton. We know that the set Y is finite, since $Y \subseteq X^*(\bar{x}, w, t)$ and the set $X^*(\bar{x}, w, t)$ is finite by assumption.

Let some $y \in Y$ be given. Since the set Y is finite, there exists $\varepsilon > 0$ such that for every $x \in X \setminus \{y\}$ with $\|x - y\|_2 \leq \varepsilon$ it holds that $\phi(x) - \phi(y) < 0$. Take $\varepsilon > 0$ sufficiently small such that the set

$$D(y) = \{x \in X : \frac{\varepsilon}{2} \leq \|x - y\|_2 \leq \varepsilon\}$$

is non-empty. Since $D(y)$ is also compact, the number $\underline{\phi} = \max_{x \in D(y)} \phi(x)$ is well-defined. Note that $\phi(y) > \underline{\phi}$.

Since ϕ is continuous and $y \in Y$, we have that $\lim_{k \rightarrow \infty} \phi(x^k) = \phi(y)$. So there exists $M_1 \in \mathbb{N}$ such that for every $k > M_1$ it holds that $\phi(x^k) > \underline{\phi}$. Furthermore, since $\lim_{k \rightarrow \infty} \|x^{k+1} - x^k\|_2 = 0$ by Proposition 2.5.1, it holds that there exists $M_2 \in \mathbb{N}$ such that for every $k > M_2$ we have $\|x^{k+1} - x^k\|_2 < \varepsilon/2$.

Let $m > \max\{M_1, M_2\}$ be such that $\|x^m - y\|_2 < \varepsilon/2$. Such an m must exist due to the fact that $y \in Y$. We argue that for every $k > m$ we have $\|x^k - y\|_2 < \varepsilon/2$. Suppose to the contrary that there exists $k > m$ with $\|x^k - y\|_2 \geq \varepsilon/2$ and let k be the smallest such number. Since $k > M_2$ and $\|x^{k-1} - y\|_2 < \varepsilon/2$, we have $\|x^k - y\|_2 < \varepsilon$, hence $x^k \in D(y)$ and $\phi(x^k) \leq \underline{\phi} < \phi(x^m)$, contradicting the fact that the sequence $(\phi(x^k))_{k \in \mathbb{N}}$ is non-decreasing.

We have shown that for every $\varepsilon > 0$ sufficiently small, there exists $m \in \mathbb{N}$ such that for every $k > m$ it holds that $\|x^k - y\|_2 < \varepsilon/2$. It follows that y is the only cluster point of $(x^k)_{k \in \mathbb{N}}$. ■

Theorem 2.5.3 combines the results in Propositions 2.4.8 and 2.5.1. It identifies sufficient conditions for one-sided learning processes to converge. The conditions are as follows: first, every update must take the updating player closer to his current best response and, second, players must update regularly. Theorem 2.5.3 and Proposition 2.5.1 together imply that any such process will spend an arbitrarily long time in any ε -range of any cluster point of the process. Finally, if we have that the set of Nash equilibria is finite, we get convergence to a single Nash equilibrium point. The latter condition holds generically as stated in Corollary 2.3.4.

We conclude this section by an example illustrating the tightness of these sufficient conditions.

Example 2.5.5. Consider the case with no strategic interaction, $w = 0$, and interior target values, for every $i \in I$, $0 < t_i < \bar{x}_i$. Then a game $G \in \mathcal{G}(\bar{x}, w, t)$ has a single Nash equilibrium, $x^* = t$. Since the Nash equilibrium set is finite, Theorem 2.5.3 applies. It

is easy to see that to achieve convergence to the Nash equilibrium, the $\alpha < 1$ condition cannot be weakened even in this simple case.

Letting $\alpha = 1$ allows for $\lim_{k \rightarrow \infty} \alpha_k = 1$, which means that the conditions of Proposition 2.5.1 are no longer satisfied. In this case, the distance to the best reply, which is equal to the target value t_i , is no longer converging to zero. We have no convergence to the Nash equilibrium.

2.6 Conclusion

In this chapter we consider weighted network games, a class of games with a very wide range of applications, where direct, pairwise player interactions are described by a matrix of weights. We show that these games are best-response potential games and, generically, have a finite set of Nash equilibria.

Pairs of linked players can always benefit from jointly deviating in an interior equilibrium. Two players whose contributions are strategic substitutes of each other can jointly increase their actions to increase their payoffs, while players whose contributions are strategic complements can jointly decrease their actions to improve their payoffs. In case all players' actions are strategic substitutes or all players' actions are strategic complements, such deviations lead to Pareto improvements. Therefore, in general, equilibria are neither strong nor efficient.

We study a large class of better-response learning processes. The convergence properties of these processes are determined by their centering parameter, which indicates to what extent players can overshoot their best responses. If players move closer to the best response at each update, as is the case for best-response dynamics and better-response dynamics with a centering parameter of one, then the players get arbitrarily close to the set of Nash equilibria and converge to a single Nash equilibrium whenever the set of Nash equilibria is finite, which is generically the case. This is due to the fact that the best-response potential is symmetric around the best response of the players, hence moving closer to the best response increases the value of the potential. In the case of better-response dynamics with unrestricted overshooting, it is shown that cycles may arise.

The restrictions on overshooting that guarantee convergence to a Nash equilibrium in the general case are the same as in a trivial game with no strategic interaction. The reason for this is that the best-response potential can be shown to increase whenever an updating player moves closer to his best response, irrespective of the values of the

interaction weights. Our results hence identify a rich class of learning processes that produce Nash equilibria, including cautious better-reply dynamics.

Topics that are left unexplored in this chapter include asymmetric interaction weights, which is the topic of Chapter 4. Another interesting topic is the issue of inefficiency of equilibria. These inefficiencies may disappear in different – possibly more centralized – classes of learning processes. Finally, the beliefs that shape the updates themselves are left unmodeled and unexplored. These topics are open for future research.

2.A Appendix: Proofs

The Appendix contains some of the more straightforward proofs, or proofs with close analogues in already existing literature.

Proposition 2.3.1. *For every $(\bar{x}, w, t) \in P$, it holds that:*

1. *Every game $G \in \mathcal{G}(\bar{x}, w, t)$ is a best-response potential game with the potential $\phi: X \rightarrow \mathbb{R}$ defined by*

$$\phi(x) = x^\top t - \frac{1}{2} x^\top W x, \quad x \in X.$$

2. $X^*(\bar{x}, w, t) \neq \emptyset$.

3. $x^* \in X^*(\bar{x}, w, t)$ if and only if x^* satisfies the Karush-Kuhn-Tucker (KKT) conditions, i.e. for every $i \in I$, there exist $\lambda_i, \mu_i \in \mathbb{R}_+$ such that

$$t_i - \sum_{j \in I} w_{ij} x_j^* + \lambda_i - \mu_i = 0,$$

$$x_i^* \geq 0, \quad \bar{x}_i \geq x_i^*,$$

$$\lambda_i x_i^* = 0, \quad \mu_i (\bar{x}_i - x_i^*) = 0.$$

Proof. We show that for every $x \in X$ and every $i \in I$ it holds that

$$\operatorname{argmax}_{x_i \in X_i} \pi_i(x_i, x_{-i}) = \operatorname{argmax}_{x_i \in X_i} \phi(x_i, x_{-i}).$$

1. The left-hand side of the equality above equals $b_i(x)$. For the right-hand side, notice that $x^\top W$ is the row vector of consumption levels of each player, $x^\top W = (\sum_{j \in I} w_{1j} x_j, \dots, \sum_{j \in I} w_{nj} x_j)$, using the fact that W is symmetric. Multiplication by x gives

$$x^\top W x = x_1 \sum_{j \in I} w_{1j} x_j + \dots + x_n \sum_{j \in I} w_{nj} x_j.$$

Differentiating $\phi(x_i, x_{-i})$ by x_i leads to

$$\frac{\partial \phi(x_i, x_{-i})}{\partial x_i} = t_i - \frac{1}{2} \left(2x_i + \sum_{j \neq i} w_{ij} x_j + \sum_{j \neq i} w_{ji} x_j \right) = t_i - \sum_{i \in I} w_{ij} x_j,$$

where the last equality uses the symmetry of W .

Setting the derivative of ϕ with respect to x_i equal to zero gives the extreme point $x_i = t_i - \sum_{j \neq i} w_{ij} x_j$, and as long as $t_i - \sum_{j \neq i} w_{ij} x_j \in [0, \bar{x}_i]$, it is the unique maximum, since the second derivative is $-w_{ii} = -1$.

If $t_i - \sum_{j \neq i} w_{ij} x_j < 0$, then, since $x_i \geq 0$, the first derivative of ϕ with respect to x_i is uniformly negative on $[0, \bar{x}_i]$ hence the unique maximum is achieved for the minimal contribution, $x_i = 0$.

Similarly, if $t_i - \sum_{j \neq i} w_{ij} x_j > \bar{x}_i$, then the first derivative is uniformly positive, meaning that, in $[0, \bar{x}_i]$, the unique maximum is achieved for the maximal contribution, $x_i = \bar{x}_i$.

2. Let $b : X \rightarrow X$ be the function such that its component $i \in I$ is equal to b_i , the best-response function of player i . Since X is non-empty, compact, and convex, and b is continuous, the existence of an x^* such that $b(x^*) = x^*$ is guaranteed by Brouwer's fixed-point theorem.

3. Since every game in $\mathcal{G}(\bar{x}, w, t)$ is a best-response potential game with potential ϕ , every Nash equilibrium satisfies the stated KKT conditions.

Since, for every $i \in I$, for every $x \in X$, it holds that

$$\frac{\partial^2 \phi(x)}{\partial x_i^2} = -1 < 0,$$

every point satisfying the KKT conditions yields a Nash equilibrium. ■

Proposition 2.3.6. *Let $(\bar{x}, w, t) \in P$ be given. Let $x^* \in X^*(\bar{x}, w, t)$ be a Nash equilibrium of a game $G = (I, X, \pi) \in \mathcal{G}(\bar{x}, w, t)$ such that $x_i^* \in (0, \bar{x}_i)$ and $x_j^* \in (0, \bar{x}_j)$. It holds that:*

1. *If for some $i, j \in I$ with $i \neq j$ it holds that $w_{ij} \neq 0$, then there exists $\delta \in (0, \min\{x_i^*, x_j^*, \bar{x}_i - x_i^*, \bar{x}_j - x_j^*\})$ such that $\pi_i(x^*) < \pi_i(x^* + \text{sgn}(w_{ij})\delta^{\{i,j\}})$ and $\pi_j(x^*) < \pi_j(x^* + \text{sgn}(w_{ij})\delta^{\{i,j\}})$.*
2. *If $w \geq 0$ or $w \leq 0$ and, for some $i, j \in I$ with $i \neq j$, it holds that $w_{ij} \neq 0$, then there exists $\delta \in (0, \min\{x_i^*, x_j^*, \bar{x}_i - x_i^*, \bar{x}_j - x_j^*\})$ such that the action profile $x^* + \text{sgn}(w_{ij})\delta^{\{i,j\}}$ is a Pareto improvement over x^* .*

Proof. 1. We first discuss the case $w_{ij} > 0$.

Since $x_i^* \in (0, \bar{x}_i)$ and $x_j^* \in (0, \bar{x}_j)$, it holds that $\sum_{k \in I} w_{ik} x_k = t_i$ and $\sum_{k \in I} w_{jk} x_k = t_j$. Hence, for $\delta > 0$ such that $x_i^* + \delta \leq \bar{x}_i$ and $x_j^* + \delta \leq \bar{x}_j$, we have

$$\frac{\pi_i(x^* + \delta^{i,j}) - \pi_i(x^*)}{\delta} = \frac{f_i(t_i + \delta(1 + w_{ij})) - f_i(t_i)}{\delta} - c_i.$$

Since f_i is concave, we have

$$f_i(t_i + \delta) \leq f_i(t_i + \delta(1 + w_{ij})) - \delta w_{ij} f'_i(t_i + \delta(1 + w_{ij})).$$

Therefore, we can write

$$\frac{\pi_i(x^* + \delta^{i,j}) - \pi_i(x^*)}{\delta} \geq \frac{f_i(t_i + \delta) - f_i(t_i)}{\delta} - c_i + w_{ij} f'_i(t_i + \delta(1 + w_{ij})).$$

Let $\varepsilon_i = \min_{x \in X} f'_i(\sum_{k \in I} w_{ik} x_k)$. Since f'_i is a continuous function, its minimum over the compact set X is well-defined. Notice that Assumption 2.2.2 guarantees that $\varepsilon_i > 0$. Thus,

$$\frac{\pi_i(x^* + \delta^{i,j}) - \pi_i(x^*)}{\delta} \geq \frac{f_i(t_i + \delta) - f_i(t_i)}{\delta} - c_i + w_{ij} \varepsilon_i.$$

Also due to the continuity of f'_i , the term $(f_i(t_i + \delta) - f_i(t_i))/\delta - c_i$ converges to zero as δ goes to zero. Hence, for sufficiently small positive δ , we have $\pi_i(x^* + \delta^{i,j}) - \pi_i(x^*) > 0$. The same argument applies to agent j .

The case $w_{ij} < 0$ follows from very similar arguments. ■

2. We first consider the case where $w_{ij} > 0$. As per Proposition 2.3.6, there exists $\delta \in (0, \min\{x_i^*, x_j^*, \bar{x}_i - x_i^*, \bar{x}_j - x_j^*\})$ such that $\pi_i(x^* + \delta^{i,j}) > \pi_i(x^*)$ and $\pi_j(x^* + \delta^{i,j}) > \pi_j(x^*)$. Since $w \geq 0$ it follows that for every other player $h \in I \setminus \{i, j\}$ we have $f_h(\sum_{k \in I} w_{hk} x_k^* + \delta^{i,j}) \geq f_h(\sum_{k \in I} w_{hk} x_k^*)$, while his own action did not change, and therefore $\pi_h(x^* + \delta^{i,j}) \geq \pi_h(x^*)$, meaning that players i and j increasing their action by δ yields a Pareto improvement.

The case $w \leq 0$ follows from similar arguments. ■

Lemma 2.5.4. *Let $(x^k)_{k \in \mathbb{N}}$ be a path of action profiles in a game $G \in \mathcal{G}$ such that $\lim_{k \rightarrow \infty} \|x^{k+1} - x^k\|_2 = 0$. For every $\varepsilon > 0$, for every $\ell \in \mathbb{N}$, there exists $M \in \mathbb{N}$ such that, for every $m > M$, for every $k \in \{m, \dots, m + \ell - 1\}$, we have $\|x^k - x^m\|_2 < \varepsilon$.*

Proof. Let some $\varepsilon > 0$ and some $\ell \in \mathbb{N}$ be given. The statement obviously holds for $\ell = 1$, so consider the case $\ell > 1$.

For every $\delta > 0$ there exists $M_\delta \in \mathbb{N}$ such that for every $m > M_\delta$ we have $\|x^{m+1} - x^m\|_2 < \delta$, since $\lim_{k \rightarrow \infty} \|x^{k+1} - x^k\|_2 = 0$. We take $\delta = \varepsilon/(\ell - 1)$ and consider an arbitrary $m > M_\delta$.

Then, by the triangle inequality, for every $k \in \{m, \dots, m + \ell - 1\}$ we can write

$$\|x^k - x^m\|_2 \leq \|x^k - x^{k-1}\|_2 + \dots + \|x^{m+1} - x^m\|_2 < \frac{k-m}{\ell-1} \varepsilon \leq \varepsilon.$$

Therefore, M_δ is a suitable candidate for M . ■

Chapter 3

Myopic and farsighted players in the local public goods game

“Forethought we may have, undoubtedly, but not foresight.”

Napoleon Bonaparte

This chapter¹ studies the exploitation of a myopic population by a single farsighted player in a network game. Our model contributes to the growing literature of local public goods games by considering multiple levels of player sophistication. By framing the game as a Markov Decision Process we show the existence and payoff-uniqueness of optimal strategies in every network structure. Every stationary equilibrium strategy profile of this game has a set of rest points called absorbing effort profiles. Myopic players have a propensity for a rising inequality in production, creating rich free-riders and poor producers. This property of the game results in the convergence of equilibrium play to an absorbing effort profile for every network structure and every initial state. The farsighted player benefits from new links that connect him to more myopic neighbors or the breaking of a link running between two of his neighbors.

¹This chapter is based upon Bayer, Herings, and Peeters (2019). A part of this chapter was written during my visit at Aix-Marseille School of Economics between March 1 and June 1, 2018 sponsored by the Graduate School of Business and Economics of Maastricht University. I would like to thank Péter Csóka, and László Kóczy, for comments on previous versions of this chapter, as well as Yann Bramoullé for feedback and suggestions.

3.1 Introduction

Economic agents interacting on the same market have vast differences in their sophistication. The sources of these differences are many, they include asymmetries in technology (an innovator against followers), goals (a speculator against investors), or experience (a new entrant against incumbents), or various other cognitive and physical limitations. When an agent with a higher level of sophistication is pitted against a population of naive agents, the former type may use its advantages to exploit the latter, if by such actions it can achieve higher gains. Depending upon the market in question and on the situation, exploitation may manifest as free-riding, breaking social contracts, manipulation, and obstructing or destroying the coordination of other players. Many models focus on homogeneous player sophistication and employ tools such as the Nash equilibrium. These tool are efficient, but conceptually valid only if every agent is sufficiently naive. Therefore, the predictions of these models are often not robust to the inclusion of a sophisticated agent. As the induced behavior, exploitation, tends to be destructive to social norms as well, its analysis is motivated by pure theoretical, behavioral, and sociological factors.

In addition to sophistication, agents also differ in their spatial properties. These properties – which may refer to physical space such as geographical position, or an abstract space such as online connectivity – define the ways by which each agent interacts with the rest of the population. Simple graphs provide a very rich but at the same time a highly efficient tool to model spatial asymmetries. In this chapter we combine tools of game theory in order to capture the strategic nature of exploitation with tools of network theory to examine the ways exploitative behavior is dependent upon the position of the sophisticated player.

The specific model in which our analysis is framed is the private provision of public goods game, also known as the local public goods game, introduced by Bramoullé and Kranton (2007). The direct predecessor of this model is Ballester et al. (2006)’s model of externalities. The class of local public goods games has an established theoretical literature as well as a wide range of applicability. The latter includes applications in industrial organization, by a model of distribution of R&D costs by interlinked firms (König et al., 2014), sociology, by a paper on crime and social networks (Calvó-Armengol and Zenou, 2004) as well as a spatial model of peer effects (Blume et al., 2010), and political economy, by a model on the distribution of defense expenditures within allied nations (Sandler and Hartley, 1995) and its follow-up, on the distribution of counter-terrorism expenditures between allies (Sandler and Hartley, 2007). Goyal

(2012) and Jackson and Zenou (2014) offer comprehensive overviews of related models.

The case of homogeneous player rationality and ability in this gameclass is well-understood. Bramoullé and Kranton (2007) shows existence of Nash equilibria and that equilibria with free riders are present in any non-trivial network. The follow-up paper, Bramoullé et al. (2014) shows existence conditions of stable Nash equilibria with respect to continuous best-response dynamics. Subsequent literature focuses on the asymptotic behavior of basic learning processes. It has been shown that best-response dynamics always converge to a single Nash equilibrium (Bervoets and Faure, 2019) and that convergence may be achieved by a learning process with very little requirements on player sophistication (Bervoets et al., 2016).

In this chapter we study the effects of heterogeneous player sophistication in the above gameclass. Through this effort we provide an insight into the behavior of economic agents who are differentiated both by their abilities and their spatial properties. To our knowledge the paper on which this chapter is based is the first to do so. In the spirit of Ellison (1997) we discuss the inclusion of a single player of increased sophistication into a population of naive agents. In games of strategic complements such as coordination games, the addition of the more sophisticated player improves efficiency, as coordination shifts towards a better Nash equilibrium via strategic teaching by the sophisticated player (Camerer et al., 2002). In contrasts, our game, a game of (local) strategic substitutes, induces a manipulative behavior by the sophisticated player with no clear gains in efficiency. By modeling sophistication as a form of farsightedness, i.e. the ability to discount future gains rather than optimizing for a single period, we create a player whose interest lies in the maximization of his neighbors' efforts into creating the local public good, while minimizing his own efforts.

The main component of local public goods games is the underlying network, describing the nature of interactions between the players, with one player designated as the farsighted player. All players, regardless of type, may decide on the amount of effort they exert into creating a specialized good. Every player consumes the total amount of goods created in their neighborhood and pays the costs of own effort. To keep our analysis focused on the players' spatial properties and on the farsighted player's attempts at exploitation, we assume that players are symmetric in their target level of consumption.

As we are interested in the process of exploitation, we examine the game described above in a dynamic setting which takes place in discrete time. In every period a randomly selected player has a chance to revise his exerted effort. Once he does, period payoffs are awarded, and a new random player receives a revision opportunity.

Period payoffs are discounted by the farsighted player but not by the myopic players who therefore maximize their instantaneous payoffs. By applying classic results of the theory of Markov Decision Processes, we show the existence of optimal stationary strategies of the farsighted player for every network structure, given the myopic players' equilibrium behavior. Therefore, the existence of stationary subgame perfect equilibria of the implied stochastic game is guaranteed. Moreover, by the theory of best-response potential games (Voorneveld, 2000), we get that the above game played by a population of myopic best-responders converges to a rest point, a Nash equilibrium of the static game.

In this chapter we show the existence of such rest points, called absorbing effort profiles, for every stationary subgame perfect equilibrium. The addition of the farsighted player means that the rest points will not necessarily coincide with static Nash equilibria, as they would do for a homogeneously myopic population. The stability of the absorbing profiles requires that each myopic player's exerted efforts be a best response to the other players', and that the farsighted player has no profitable way of advancing the game into a state that is better for him. Furthermore, in our main convergence result we show that for every network, every starting state, and every equilibrium play the process converges to such an absorbing profile with probability 1.

Our results allow for a straightforward description of the ways in which the inclusion of a farsighted player alters key aspects of equilibrium play. The presence of the farsighted player does not weaken the convergence properties of the game, but changes the set of rest points. There are two mechanisms working to this effect: (1), Nash equilibria which are easy to improve upon by the farsighted player will not be stable against farsighted deviations, and (2), some effort profiles that are stable for the myopic players may be desirable for the farsighted player even if he is not at his static best response. As a result the set of absorbing effort profiles is neither a subset nor a superset of the Nash equilibria of the static game.

A further key effect of including the farsighted player is that, unlike a game played by myopic players, equilibrium play may produce cycles between action profiles, as the farsighted player may try to revisit a previous action profile hoping for a more favorable random draw in the order of revising players. However, the probability of the game being locked in an infinite cycle is zero. This, as well as the convergence of the game, is guaranteed by a crucial property of the gameclass, a tendency towards increased free-riding behavior. This effect is propagated by the local substitutability of efforts between neighboring players. An increase in one player's efforts will decrease his neighbor's myopic best-response efforts. Similarly, a neighbor's decrease will increase

a player's myopic best-response efforts. As we will show in this chapter, the result of these mutually reinforcing effects is strong enough to break any cycles that the farsighted player might try to implement, and thus leads to convergence.

Throughout most of this chapter we keep the underlying network structure fixed, but in applications the spatial properties of the players are liable to change, either by the players' own choice, or by external factors. For this reason we consider comparative statics of different networks. Evaluating networks can inform the farsighted player's choice on which links to form and which to sever to increase his utility in a network formation setting or advise a social planner who aims to minimize the possibilities of exploitation in a game by restructuring the network. We find that adding a link to connect the farsighted player to an additional neighbor, or removing a link between two of his neighbors leads to an increase in the farsighted player's evaluations.

The combination of these two results confirms the intuition that being the center of a star network is optimal for the farsighted player since it allows for the full exploitation of the entire myopic population. Moreover, it shows that there is a sequence of single changes (the addition or removal of one link) that lead to the star network and each step improves on the farsighted player's evaluation of the network. However, we also show that not every sequence of changes that leads to the star network is profitable in every step, meaning that networks may contain stumbling blocks that hinder the farsighted player's attempts of reaching the optimal network structure.

To sum up, this chapter contributes to the network literature and to the game theory literature by building a dynamic model of exploitation of a myopic population by a single farsighted player on a network. We show existence of equilibrium play, existence of stable effort profiles with respect to equilibrium play, and convergence to these stable effort profiles. Finally, we consider how changes in the network affect the farsighted player's long-run utility.

The chapter proceeds as follows. Section 3.2 introduces local public goods games and introduces the dynamic and stochastic components of the model. Section 3.3 shows the existence of absorbing profiles of efforts. Section 3.4 contains our results regarding the propensity of rising inequality between players under equilibrium play. Section 3.5 contains our main convergence result, showing that equilibrium play progresses the game to an absorbing profile of actions. In Section 3.6 we examine how the farsighted player's evaluations of a network are affected by adding or severing a link. Section 3.7 concludes.

3.2 The model

We now introduce the static version of local public goods game with the set of primitives, $(n, G, (f_i)_{i \in I}, (c_i)_{i \in I}, d)$. We then go on to extend this model into a recurrent setting taking place in discrete time with additional primitive δ .

Let $I = \{1, \dots, n\}$ denote the set of players with $n \geq 2$. Player 1 is called the farsighted player. All players other than 1 are called myopic. Let the set $I \setminus \{1\}$ be denoted by I_{-1} .

A network structure, characterized by the $n \times n$ binary matrix G , will represent local interaction. For $i, j \in I$, $g_{ij} = 1$ means that players i and j are connected, while $g_{ij} = 0$ means that they are not. We assume $g_{ii} = 1$, and we assume that connections run both ways, i.e. for every $i, j \in I$ we have $g_{ij} = g_{ji}$. For a player $i \in I$, the set $N_i(G) = \{j \in I \setminus \{i\} : g_{ij} = 1\}$ is called the set of i 's neighbors, or i 's neighborhood. For $J \subseteq I$ let $N_J(G) = \{i \in I \setminus J : g_{ij} = 1 \text{ for some } j \in J\}$ denote the neighborhood of a set of players J . Finally, for $i \in I$ let $\bar{N}_i(G) = N_i(G) \cup \{i\}$, and for $J \subseteq I$ let $\bar{N}_J(G) = N_J(G) \cup J$. Throughout Sections 3.2, 3.3, 3.4, and 3.5, we assume a fixed network structure and hence shorten these notations to N_i , N_J , \bar{N}_i , and \bar{N}_J , respectively.

The set of possible efforts of player i is $X_i = \{0, 1/d, 2/d, \dots\}$, with positive integer d . Let $X = \prod_{i \in I} X_i$ denote the set of effort profiles, and let $X_{-i} = \prod_{j \in I \setminus \{i\}} X_j$. For $i \in I$ and $x \in X$, the i th component of the effort profile, x_i is called the effort exerted by player i . For an effort profile $x \in X$ and a set of players $J \subseteq I$ we denote by $x(J) = \sum_{i \in J} x_i$ the sum of the players' efforts. The value $x(\bar{N}_i)$ is called player i 's consumption.

The players are engaged in a local public goods game with payoff functions $\pi_i : X \rightarrow \mathbb{R}$ that are given by

$$\pi_i(x) = f_i(x(\bar{N}_i)) - c_i x_i,$$

where $c_i \in \mathbb{R}^+$, and the functions $f_i : \{0, 1/d, 2/d, \dots\} \rightarrow \mathbb{R}^+$ satisfy the following properties:

1. $f_i(0) = 0$,
2. f_i is monotonically increasing,
3. $f_i(y) - f_i(y - 1/d) \geq f_i(y + 1/d) - f_i(y)$ for all $y \in \{1/d, 2/d, \dots\}$,
4. $f_i(1) - f_i(1 - 1/d) > c_i/d$, and $f_i(1 + 1/d) - f_i(1) < c_i/d$.

5. f_i is bounded.

For $i \in I$ and $x \in X$, the value $f_i(x(\bar{N}_i))$ is called the benefits enjoyed by player i . The function f_i is called the benefit function of player i , while the parameter c_i is the cost of one unit of effort for player i . Let $\pi(x)$ denote the vector of payoffs $(\pi_1(x), \dots, \pi_n(x))$ in effort profile x .

The interpretation of the game characterized by $\Gamma = (I, X, (\pi_i)_{i \in I})$ is the following. Similarly to the game in Chapter 2, each player decides how much to contribute to the local public good. After choosing their efforts, the players reap the benefits from their own efforts, as well as from the efforts of their immediate neighbors, and the value of the benefits is given by the benefit functions f_i . In addition, all players incur linear costs for their own efforts, exerting unit level of efforts comes with a cost of c_i for player i . The five properties of the benefit function respectively mean (1) the benefits of consuming zero units of the good are zero, (2) monotonicity, (3) concavity, (4) the existence of an optimal effort level, above which the cost of producing any additional amount of the public good is larger than the benefits of its consumption for each player, and (5) boundedness. These properties are standard in the discussion of local public good games. The optimal effort levels are the same for all players and, by (4), their value is normalized to unity.

We are interested in the dynamic behavior of a farsighted player acting against a set of myopic players in a recurrent version of Γ , denoted by Γ' . The dynamics take place in discrete time, indexed by $t = 1, 2, \dots$. The game starts with an arbitrary profile of efforts, x^0 . In every period t , an active player i^t is chosen via a fixed probability distribution from all players. Similarly to Chapter 2's updating process, if player i^t holds a revision opportunity, then the profile x^t is given as a concatenation of his choice, denoted by z^t , and every other player's effort in period $t - 1$. Instantaneous payoffs, $\pi_i(x^t)$ are awarded in period t . In period $t + 1$ a new active player is selected, and so on. A state is therefore characterized by a vector of efforts, and the active player. The set of states is denoted by $\Omega = X \times I$. For a player $i \in I$, an effort level $x_i \in X_i$, and an effort profile $x' \in X$ the effort profile (x_i, x'_{-i}) is the one that differs from x' in only the i th coordinate, with player i playing x_i .

Fix $T \geq 1$. For time periods $1 \leq t \leq T$, states $\omega^t = (x^{t-1}, i^t) \in \Omega$, and choices $z^t \in X_{i^t}$, a sequence $h^T = (\omega^1, z^1, \omega^2, z^2, \dots, \omega^{T-1}, z^{T-1}, \omega^T)$ such that for every $1 \leq t \leq T$ it holds that

$$x_{-i^t}^t = x_{-i^t}^{t-1} \text{ and } x_{i^t}^t = z^t \quad (3.1)$$

is called a T -long history. Let H^T denote the set of T -long histories, and $H = \bigcup_{t \in \mathbb{N}} H^t$ denote the set of finite histories. For a player i , an effort profile x , and a state ω , let H_i^T , H_x^T , and H_ω^T denote the set of T -long histories for which the last active player, i^T is i , the last effort profile, x^{T-1} is x , and the last state, ω^T is ω , respectively. Define $H_i = \bigcup_{t \in \mathbb{N}} H_i^t$, $H_x = \bigcup_{t \in \mathbb{N}} H_x^t$, and $H_\omega = \bigcup_{t \in \mathbb{N}} H_\omega^t$.

The function representing a player's decisions in all possible situations is called a strategy. A strategy s_i of a player $i \in I$, is a function $s_i: H_i \rightarrow X_i$. For a history $h_i^T \in H_i^T$, the value $s_i(h_i^T)$ is the effort chosen by player i after observing history h_i^T . Denote the set of strategies of player i by S_i and let $S = \prod_{i \in I} S_i$.

The transitions between the states are governed by a probability distribution $Q: \Omega \times \{0, 1/d, \dots\} \rightarrow \Delta(\Omega)$. The probability that state ω' follows ω if choice z is made equals $Q(\omega'|\omega, z)$. The marginal distributions of Q with respect to X and I are denoted by Q_X and Q_I , respectively, i.e. for every $\omega \in \Omega$ and $z \in \{0, 1/d, \dots\}$ it holds that $Q_X(x|\omega, z) = \sum_{i \in I} Q(x, i|\omega, z)$, and $Q_I(i|\omega, z) = \sum_{x \in X} Q(x, i|\omega, z)$. Due to the consistency condition (3.1), for every $x, x' \in X$, $i \in I$ and $z \in X_i$ we have

$$Q_X(x'|(x, i), z) = \begin{cases} 1, & \text{if } x' = (z, x_{-i}) \\ 0, & \text{otherwise.} \end{cases}$$

Hence, any assumption made on Q_I pins down Q itself. For now we assume that the next active player is selected uniformly from the set of all players at all times.

Assumption 3.2.1. For every $i \in I$, every $\omega \in \Omega$, and every $z \in \{0, 1/d, \dots\}$ we assume that $Q_I(i|\omega, z) = 1/n$.

Assumption 3.2.1 provides simplicity for our framework. It constitutes a deviation from Chapter 2's model, where no explicit assumptions were made on the assignment of revision opportunities. As it will be apparent later on, the results in this chapter hold for a more general class of probability distributions, requiring only that for each period, every player should have at least some $\varepsilon > 0$ probability of becoming the active player.

All players receive instantaneous payoffs in all periods, regardless of activity and history. In period $t \geq 1$, player $i \in I$ receives instantaneous payoff $\pi_i(x^t)$. For strategy profile s and a T -long history $h \in H^T$, let $u_1^t(h, s)$ denote the expected instantaneous payoff that player 1 receives in period $T + t - 1$. For example, $u_1^1(h, s) = \pi_1(s_{iT}(h), x_{-iT}^{T-1}) = \pi_1(x^T)$. Player 1's utility is given as the discounted expected instantaneous payoffs, with discount factor $\delta \in (0, 1)$ as follows:

$$U_1(h, s) = \sum_{t=1}^{\infty} \delta^{t-1} u_1^t(h, s)$$

as player 1's utility function.

A myopic player $i \in I_{-1}$ cares only about his instantaneous payoffs, so for $x \in X$ and $h \in H^T$, his utility function is defined by

$$U_i(h, s) = \pi_i(s_{iT}(h), x_{-iT}).$$

The main equilibrium concept of dynamic games is the Subgame Perfect Equilibrium (SPE), which is defined as follows:

Definition 3.2.2 (SPE). Strategy profile s^* is an SPE of Γ' if for every player $i \in I$, every history $h \in H$, and every strategy $s_i \in S_i$ it holds that

$$U_i(h, s^*) \geq U_i(h, (s_i, s_{-i}^*)).$$

In an SPE in every period the farsighted player maximizes his discounted expected payoffs of the future, while the myopic players maximize instantaneous payoffs.

The strategy $b_i \in S_i$ for which $b_i(h^t) = \max\{0, 1 - x^{t-1}(N_i)\}$ is called the best-responder strategy of player i . Clearly, if s^* is an SPE, then $s_i^* = b_i$ for every $i \in I_{-1}$.

Note that the best-responder strategy depends only on the last state observed. Strategies that satisfy this criterion are called stationary strategies. More formally, a strategy $s_i \in S_i$ is called stationary if there is a function $s'_i: X \rightarrow X_i$ such that for every $x \in X$, and every $h \in H_{(x,i)}$ it holds that $s_i(h) = s'_i(x)$ (note that the set X is isomorphic with $H_{(\cdot,i)}^0$). The value $s'_i(x)$ is player i 's choice at effort profile x . We identify the set of stationary strategies of player i with the set of maps $\tilde{S}_i = \{s'_i: X \rightarrow X_i\}$, and the set of stationary strategy profiles with $\tilde{S} = \prod_{i \in I} \tilde{S}_i$. The main equilibrium concept used in this chapter is the Stationary Subgame Perfect Equilibrium (SSPE), defined as follows:

Definition 3.2.3 (SSPE). A strategy profile s^* is an SSPE of Γ' if it is an SPE and for every $i \in I$, the strategy s_i^* is stationary.

Again, if s^* is an SSPE of Γ' , then for every $i \in I_{-1}$ we have $s_i^* = b_i$. In Proposition 3.2.4 we show that an SSPE exists in this game. Denote by $(s_1, b_{-1}) \in S$ the profile of strategies where player 1 plays s_1 and every other player is a best-responder.

Proposition 3.2.4. *The game Γ' has an SSPE.*

Proof. Consider a Markov Decision Process (MDP) with set of states Ω . We assume that in every state (x, i) for which $i \neq 1$, nature chooses $b_i(x)$. With that assumption in place, player 1 is left as the only strategic player. A policy of his is a function $p: H_1 \rightarrow X_1$, whereas a stationary policy is a function $p': X \rightarrow X_1$. The process that drives the transition to a new state from ω to ω' , given choice z is $Q(\omega'|\omega, z)$.

Since Q is Markovian, this is indeed an MDP. Let the reward after history h given policy p be given as

$$V_p(h) = U_1(h, (p, b_{-1})).$$

Denote this MDP by Γ'_1 and let p^* be an optimal stationary policy of Γ'_1 . Then, it holds that for every $h \in H$ we have

$$p^* \in \operatorname{argmax}_p V_p(h).$$

Hence, for every h we have $p^* \in \operatorname{argmax}_{s_1 \in S_1} U_1(h, (s_1, b_{-1}))$, and therefore $(p^*, b_{-1}) \in \tilde{S}$ is an SSPE. ■

Let the set of SSPEs be denoted by S^* . Proposition 3.2.4 makes use of the fact that the myopic players' equilibrium behavior is independent of that of their opponents. Therefore, we can model the farsighted player's decision problem as an MDP with the myopic players treated as parts of nature. An optimal stationary policy in this MDP is a best response by the farsighted player to the myopic players' equilibrium strategies in game Γ' . Together, an optimal stationary policy and the myopic best-responses form an SSPE of Γ' . Note that if δ is small, we have a unique SPE/SSPE, in which every player is a best-responder. Furthermore, it can be shown that for a general δ every SPE yields the same expected utility for player 1:

Proposition 3.2.5. *For any two SPE strategy profiles $r, s \in S$ of game Γ' it holds that for every $h \in H$ we have*

$$U_1(h, r) = U_1(h, s).$$

Proof. If s, r are both SPE strategy profiles, then s_1 and r_1 are both optimal (potentially non-stationary) policies of Γ'_1 . Therefore, it holds for every h that

$$V_{s_1}(h) = V_{r_1}(h).$$

■

Propositions 3.2.4 and 3.2.5 establish the existence and the payoff-uniqueness of the farsighted player's best strategy for any network structure. Note, however, that the myopic players' payoffs may not be the same in each SSPE.

In the remaining of this chapter we restrict our attention to stationary strategies and will use SSPE as our equilibrium concept, since no player can realize gains in utility by choosing a non-stationary strategy.

Furthermore, we will pay special attention to the case where player 1 is patient, i.e. δ is close to 1, as in this case player 1 is most prompted to play more sophisticated strategies than his opponents. A close examination of this case reveals in which situations a farsighted agent is able to manipulate the game in order to end up in a position where he can free-ride on his neighbors' efforts.

3.3 Existence of absorbing effort profiles

This section contains the definition of an equilibrium concept of effort profiles, called absorbing effort profiles. We then go on to show existence.

Definition 3.3.1 (Absorbing effort profile). An effort profile $x \in X$ is *absorbing* for stationary strategy profile $s \in \tilde{S}$ if for every player $i \in I$ it holds that

$$s_i(x) = x_i.$$

In other words, if x is an absorbing effort profile of s , then if the game reaches x , it will never move out of it. Let the set of absorbing effort profiles of s be denoted by $A(s)$.

We define two standard concepts of graph theory, cliques and independent sets.

Definition 3.3.2 (Cliques). A non-empty set of players $M \subseteq I$ is called a clique, if for every $i, j \in M$ we have $g_{ij} = 1$.

Notice that every single player constitutes a clique, as does every pair of neighbors.

Definition 3.3.3 (Independent sets). Let $J \subseteq I$. The set of players $M \subseteq J$ is called an *independent* set of J if for every $i, j \in M$ with $i \neq j$ it holds that $g_{ij} = 0$. The set M is called a *maximal* independent set of J if it is independent and $J \subseteq \overline{N}_M$. The set M is called a *maximum* independent set of J if it is independent and there is no other independent set M' of J for which $|M'| > |M|$.

Notice that for every $J \subseteq I$, every maximum independent set of J is also a maximal independent set of J .

We now present the main result of this section, the existence of absorbing effort profiles for SSPE strategies. Even stronger, we identify a profile, x^* , that is absorbing for every SSPE.

Let M_1 be a maximum independent set of N_1 . Let M be a maximal independent set of I such that for every $i \in M_1$ we have $i \in M$. Such an M can be easily constructed for a given M_1 , by adding a player who is not in \overline{N}_{M_1} , then adding a player outside the neighborhood of M_1 and the new player, and so on, until we run out of players. Let $x^* \in X$ be defined as $x_i^* = 1$ if $i \in M$ and $x_i^* = 0$ if $i \notin M$. Notice that $N_1 \neq \emptyset$ implies $1 \notin M$ and $N_1 = \emptyset$ implies $1 \in M$.

Proposition 3.3.4. *For every SSPE s^* , and any x^* as defined above, it holds that $x^* \in A(s^*)$.*

The proof of Proposition 3.3.4 requires a number of technical steps. We first introduce the notion of reachability between effort profiles.

Definition 3.3.5 (Reachability). Let $x, x' \in X$ and $s \in \tilde{S}$. We say that x' is *reachable* from x in s if either $x = x'$, or there exists a sequence of effort profiles satisfying equation (3.1) in which x precedes x' (not necessarily as an immediate predecessor) and the updates are made in accordance with strategy profile s , i.e. $z^t = s_{i^t}(x^{t-1})$.

If x' is reachable from x in s then we write $x \preceq_s x'$. We write $x \sim_s x'$ if $x \preceq_s x'$ and $x' \preceq_s x$ both hold. We write $x \prec_s x'$ if $x \preceq_s x'$ and $x' \not\preceq_s x$. Clearly, \sim_s is an equivalence relation. An equivalence class of \sim_s containing x is called a recurring set in s if there exists no $x' \in X$ for which $x \prec x'$. The recurring set in s that contains the effort profile $x \in X$ is denoted by $[x]_s$. If $y \in [x]_s$ for some x then we say that y is recurring. Finally, let $R_s(x) = \{x' \in X : x \preceq_s x'\}$ denote the set of effort profiles that are reachable from x in s .

We now introduce the notion of pillar sets as an upper bound of the total contribution of a subset of myopic players. The idea is the following: We pick a myopic player with positive effort, called a pillar, whose neighborhood without the farsighted player has a total effort of at most one. We remove the player and his myopic neighbors from the network. We find another such player and repeat until we run out of players with positive effort. The pillar set contains the players selected by this process. The name pillar is chosen for the players' property of supporting at most 1 unit of effort.

Let $I_{-1}^+(x) = \{i \in I_{-1} : x_i > 0\}$, and $I_{-1}^0(x) = \{i \in I_{-1} : x_i = 0\}$, denote the myopic contributors and free-riders in effort profile x , respectively.

Definition 3.3.6 (Pillar set). Let $J \subseteq I_{-1}$ and $x \in X$ be given. The set $M \subseteq J$ is called a *pillar set* of J in x if there exists some enumeration of the elements of M , $(i^k)_{k=1}^{|M|}$, such that

- M is a maximal independent set within $J \cap I_{-1}^+(x)$.
- for every $k \in \{1, \dots, |M|\}$ we have $x((J \cap \overline{N}_{i^k}) \setminus \bigcup_{k'=1}^{k-1} \overline{N}_{i^{k'}}) \leq 1$.

We denote the collection of pillar sets of J in effort profile x by $\mathcal{M}_x(J)$. In the special case $J \cap I_{-1}^+ = \emptyset$, $M = \emptyset$ is a pillar set.

A pillar set of a subset of players for a given effort profile provides an upper bound on the total effort levels of all players in that set. Clearly, if a pillar set M exists for an effort profile $x \in X$ and a subset of players $J \subseteq I_{-1}$, then the properties of Definition 3.3.6 guarantee that $x(J) \leq |M|$ holds. Note that this upper bound is usually not tight. The notion of pillars is demonstrated by an example.

Example 3.3.7. Let $I = \{1, \dots, 7\}$, and consider the network structure in which every myopic player is a neighbor of the farsighted one, and the myopic players are located on a circle. Consider the effort profile $x = (0.1, 0.4, 0.4, 0.3, 0.4, 0.4, 0.1)$ (see Figure 3.1). Then, $\{2, 4, 6\}$ is a pillar set of I_{-1} in x , as all myopic players exert positive effort in x , M is a maximal independent set of I_{-1} , and, if we take the sequence $(2, 4, 6)$, it holds that $x(\{2, 3, 7\}) = 0.9 \leq 1$, $x(\{4, 5\}) = 0.7 \leq 1$, and $x(\{6\}) = 0.4 \leq 1$. This set is of cardinality 3, and hence we should have $x(I_{-1}) \leq 3$, which is of course true, as $x(I_{-1}) = 2$. Notice that $\{2, 5\}$ is not a pillar set of I_{-1} , since $x(\{4, 5, 6\}) = 1.1 > 1$.




Figure 3.1: The links between the myopic players of Example 3.3.7.

By Definition 3.3.6, two properties immediately follow.

1. **Inclusion.** For $x \in X$ and sets $J, J' \subseteq I_{-1}$ with $J \subseteq J'$, if $M \in \mathcal{M}_x(J')$ with $M \subseteq J$, then $M \in \mathcal{M}_x(J)$.

2. **Addition.** For $x \in X$ and sets $J, J' \subseteq I_{-1}$, if $M \in \mathcal{M}_x(J)$ and $M' \in \mathcal{M}_x(J')$, and if $M \cup M'$ is an independent set, then $M \cup M' \in \mathcal{M}_x(J \cup J')$.

Let $X'_{-1} \subset X_{-1}$ denote the set of effort profiles of myopic players that accommodate a pillar set for every $J \subseteq I_{-1}$. Let $X' = X_1 \times X'_{-1}$. Notice that X'_{-1} is a finite set. The existence of a pillar set of I_{-1} in an effort profile does not necessarily imply the existence of a pillar set in a $J \subseteq I_{-1}$. We now show that if every myopic player is a best-responder, as they are in equilibrium, the game will enter the set X' with positive probability, and once there, it cannot leave it.

Lemma 3.3.8. *Let $s \in \tilde{S}$ be given such that $s_{-1} = b_{-1}$.*

1. *Reachability:* For every $x \in X \setminus X'$ it holds that $R_s(x) \cap X' \neq \emptyset$.
2. *No exit:* For every $x \in X'$ it holds that $R_s(x) \subseteq X'$.

Proof. 1. We show that once every myopic player has been active at least once, the resulting effort profile must be in X' . Let $((x^0, 2), \dots, (x^{n-2}, n), x^{n-1}) \in \Omega^{n-1} \times X$ be a sequence satisfying (3.1) such that $x^0 \in X \setminus X'$. The proof constructs a pillar set of x^{n-1} for an arbitrary J .

Fix $J \subseteq I_{-1}$ and consider the player in J with the highest index number with positive efforts in x^{n-1} . If no myopic player in J has positive efforts, the empty set is a pillar set of J in x^{n-1} . Otherwise let this player be denoted by j_1 . Then, we have $x^{j_1-1}(\overline{N}_{j_1}) = 1$. By the choice of j_1 we have $1 = x^{j_1-1}(\overline{N}_{j_1}) \geq x^{n-1}(J \cap \overline{N}_{j_1})$. Remove the myopic players in \overline{N}_{j_1} from the game, and let j_2 denote the myopic player in $J \setminus \overline{N}_{j_1}$ with the highest index number of the remaining players with positive efforts in x^{n-1} . If there is no such player, then $M = \{j_1\}$ is a suitable pillar set of J in x^{n-1} . Otherwise, by the choice of j_2 , it holds that $1 = x^{j_2-1}(\overline{N}_{j_2}) \geq x^{n-1}((J \cap \overline{N}_{j_2}) \setminus \overline{N}_{j_1})$. Remove the myopic players belonging to \overline{N}_2 and repeat. If there are no more myopic players with positive efforts in $J \setminus (\overline{N}_{j_1} \cup \overline{N}_{j_2})$, then $M = \{j_1, j_2\}$ is a suitable pillar set of J in x^{n-1} . Otherwise we go on as before with a j_3 . Since there are only finitely many myopic players, this construction will terminate after at most $n - 1$ steps, resulting in pillar set of J in x^{n-1} . Since such a construction can be performed for any choice of J , it holds that $x^{n-1} \in X'$.

2. Suppose that the statement is false. Then, there exists a sequence of states, $((x^0, i^1), \dots, (x^T, i^{T+1})) \in \Omega^{T+1}$, that satisfies (3.1) such that $x^0 \in X'$ and $x^T \in X \setminus X'$. Then, there must exist $J \subseteq I_{-1}$ such that there is no pillar set of J in x^T . We show that, no matter the choice of J , a pillar set can always be constructed.

We use an induction argument. First, consider the case $|J| = 1$, i.e. $J = \{j\}$. If $x_j^T = 0$, then the empty set is a pillar set. If $0 < x_j^T \leq 1$, then $\{j\}$ is a pillar set. The case $x_j^T > 1$ is impossible, since $j \in \{i^1, \dots, i^T\}$ implies $x_j^T = [0, 1]$, while $j \notin \{i^1, \dots, i^T\}$ means $x_j^T = x_j^0$, hence the pillar set of $\{j\}$ in x^0 is also a pillar set in x^T .

Suppose $|J| = k$ for $k > 1$ and for every $1 \leq k' < k$ it holds that every subset of $I_{-1}^+(x^T)$ with a cardinality of k' has a pillar set in x^T . If there exists $j \in J$ with $x_j^T = 0$, then take the set $J' = J \setminus \{j\}$. Since a pillar set only includes contributing players, by the induction hypothesis there exists a pillar set of J' in x^T , which is clearly a pillar set of J in x^T as well.

If $J \subseteq I_{-1}^+(x^T)$ and if $J \cap \{i^1, \dots, i^T\} = \emptyset$, then for every $j \in J$ we have $x_j^0 = x_j^T$, hence the existence of a pillar set of J in x^0 implies the existence of a pillar set of J in x^T .

If $J \subseteq I_{-1}^+(x^T)$ and if $J \cap \{i^1, \dots, i^T\} \neq \emptyset$, then let $j_1 \in J$ denote the last player in the sequence $(i^t)_{t=1}^T$ who is also a member of J . By this choice, we must have $x^T(\overline{N}_{j_1} \cap J) \leq 1$. Consider the set of players, $J' = J \setminus \overline{N}_{j_1}$. By the induction hypothesis, there exists an $M' \subseteq J'$ which is a pillar set of J' in x^T . Then, since $M' \cap N_{j_1} = \emptyset$, by the Addition property, $M = \{j_1\} \cup M'$ is a pillar set of J' in x^T . ■

Remark 3.3.9. By the proof of Lemma 3.3.8, two statements follow.

1. Notice that not only is the set X' reachable from any effort profile, it is reachable via at most $n - 1$ updates, hence it will be reached almost surely for any starting state.
2. For a fixed $x_1 \in X_1$ it also holds that once the game reaches $\{x_1\} \times X'_{-1}$, it cannot leave this set unless player 1 updates. The proof is identical to that of the 'No exit' part of Lemma 3.3.8.

As a consequence of Lemma 3.3.8, for any strategy profile s such that $s_{-1} = b_{-1}$, and any recurrent set $[x]_s$ we must have $[x]_s \subseteq X'$.

Let $X^* = \{x \in X: \forall i \in I_{-1} \ x_i = b_i(x)\}$ and $X_+^* = \{x \in X: \forall i \in I_{-1}^+(x) \ x_i = b_i(x)\}$ denote the sets of effort profiles where every myopic player is at his best response, and every contributing myopic player is at best response, respectively. Notice that for every $s_1 \in \tilde{S}_1$ we have $A(s_1, b_{-1}) \subseteq X^* \subseteq X_+^*$.

We now prove Proposition 3.3.4.

Proof of Proposition 3.3.4. Since every myopic player is a best-responder in equilibrium, for every SSPE strategy profile s^* , every $x \in X^*$, and every $i \in I_{-1}$ we have that

$s_i^*(x) = x_i$. Let $X^1 = \operatorname{argmax}_{x' \in X'} \pi_1(x')$. Due to the second point of Lemma 3.3.8, for every s^* and $x \in X^1 \cap X^*$ we have $s_1^*(x) = x_1$, as $\pi_1(x)$ is the maximum instantaneous payoff that can be reached in X' . Recall the definitions of x^* and M_1 . Clearly, $x^* \in X^*$. We now show that x^* is an element of X^1 as well.

If $N_1 \neq \emptyset$, then, since the cardinality of every pillar set of N_1 is an upper bound of the total effort in N_1 , and no pillar set of N_1 can have a bigger cardinality than M_1 , by the fact that $x^*(N_1) = |M_1|$ we have that $x^* \in \operatorname{argmax}_{x' \in X'} x'(N_1)$. By the fact that $x_1^* = 0$, and by the assumptions made on π_1 , it follows that $x^* \in X^1$ as well. This means that $x^* \in A(s^*)$ for any SSPE s^* .

If $N_1 = \emptyset$, then for any effort profile $x \in X$ it holds that $x(N_1) = 0$, and the payoff of player 1 depends only on x_1 . By the assumptions placed on π_1 , it follows that $s_1^*(x) = 1$ for every $x \in X$, hence $x^* \in A(s^*)$ for any SSPE s^* . ■

Proposition 3.3.4 identifies an effort profile which is both acceptable for the myopic players, as they are playing best responses, and for the farsighted player, who gets the maximum instantaneous payoff attainable from this position. By this construction we show not only that the set of absorbing effort profiles is non-empty for every SSPE, but also that it has a non-empty intersection with the set of static Nash equilibria of the game. However, it is not true that the set of absorbing effort profiles would equal the set of Nash equilibria, as shown by the following example.

Example 3.3.10. Let $I = \{1, 2, 3\}$, $d = 2$, $c_i = 1$ for $i \in I$. Let the network G be given by $g_{12} = g_{23} = 1$, and $g_{13} = 0$. Consider the effort profile $x^0 = (0, 0.5, 0.5)$ (see figure 3.2).




Figure 3.2: The network and initial effort profile of Example 3.3.10. Player 1 can choose between remaining under his target consumption, or disturbing the equilibrium between players 2 and 3 and prompting a response.

Clearly, for a strategy $s_1 \in \tilde{S}_1$ such that $s_1(x^0) = 0$, it holds that $x^0 \in A(s_1, b_{-1})$, whereas we clearly have $x^0 \notin X^*$. We show next that by choosing δ and f_1 appropriately, $s_1(x^0) = 0$ is the equilibrium play of player 1. By playing 0 at x^0 , player 1 will receive the payoff stream $f_1(0.5)$ indefinitely.




Figure 3.3: If player 1 increases his effort in Example 3.3.10, then he faces a permanent decrease in efforts from player 2. Player 3 provides for both 2 and 3, and player 1 cannot change that.

For $s_1 \in \tilde{S}_1$ with $s_1(x^0) > 0$, the only effort profiles we can reach from x^0 are the ones where players 2 and 3 play 0.5 and 0.5, or 0 and 1, respectively. Moreover, with probability 1, we will reach a profile where the latter is the case, since at any point when player 1 has a positive effort level, if players 2 and 3 receive revision opportunities in this order, they will set their efforts to 0 and 1 respectively. From this point onwards, by the assumptions on f_1 , every period's instantaneous payoff that player 1 receives is at most $f_1(1) - 1$. Hence, if $f_1(1) - 1 > f_1(0.5)$ and if δ is high enough, player 1's equilibrium choice in profile x^0 is $s_1(x^0) = 0$. In this case, the set of absorbing effort profiles of the equilibrium strategy will include an effort profile that is not a Nash equilibrium of the static game.

Notably, Example 3.3.10 also illustrates that the best profile for the farsighted player, $x^* = (0, 1, 0)$, is not attainable from the $(0, 0.5, 0.5)$ starting effort profile. Even so, an absorbing profile has been reached, which, as shown in Section 3.5, is generally true for all games and all initial profiles.

3.4 Rising inequality between myopic players

In this section we concentrate on the myopic players. In particular, we focus on the progression of the game when it is played by a set of best-responding players revising in sequence. We show that such dynamics allow for a rise of inequality in the efforts exerted towards the local public good.

Since the farsighted player can gain utility by exploiting as many of his myopic neighbors as possible, the results of this section provide insight to his opportunities and limitations. Therefore, we will find it particularly useful to set up our convergence theorem. First we show that an effort profile where all myopic players are at their best replies is reachable from any effort profile.

Lemma 3.4.1. *Let $s \in \tilde{S}$ be such that $s_{-1} = b_{-1}$. For every $x \in X$ there exists an $x' \in R_s(x)$ such that $x' \in X^*$ and $x'_1 = x_1$.*

Lemma 3.4.1 follows by the fact that local public goods games are best-response potential games (Voorneveld, 2000). As shown in Chapter 2, the function

$$\phi(x) = \sum_{i \in I} x_i - \frac{1}{2} x^\top G x$$

is a best-response potential of a static local public goods game with continuous action sets. Notice that for every $(x, i) \in \Omega$ it holds that $\operatorname{argmax}_{x_i \in \mathbb{R}_+} \phi(x_i, x_{-i}) \in X_i$, since $x(N_i) \in \{0, 1/d, \dots\}$, and hence we have

$$\operatorname{argmax}_{x_i \in X_i} \pi_i(x_i, x_{-i}) = \operatorname{argmax}_{x_i \in X_i} \phi(x_i, x_{-i})$$

by the same arguments as in Chapter 2, meaning that ϕ is a best-response potential of Γ .

By Remark 3.3.9, if player 1 is never selected to become active, the game can reach the finite set $\{x_1\} \times X'_{-1}$ by at most $n - 1$ updates, with every update increasing the value of ϕ . Furthermore, as the game cannot leave $\{x_1\} \times X'_{-1}$ and the value of ϕ cannot decrease, an effort profile will be reached, where ϕ cannot increase any further by unilateral deviations. At that point, every myopic player will be at his best response.

Lemma 3.4.1 means that should the farsighted player refrain from changing his own effort – whether by his own choice or by the fact that he is not selected to become active for a time – the myopic players will reach a partial equilibrium.

We now present a lemma on rising inequality between a linked pair of myopic players. Given an effort profile such that one member of the pair has a higher consumption than the other, the effort profile where the former player exerts no effort while the latter one provides for the former is always reachable.

Lemma 3.4.2 (Rising inequality between pairs). *Let $s \in \tilde{S}$ be such that $s_{-1} = b_{-1}$. Let $x \in X$ and $i, j \in I_{-1}$ be given such that $g_{ij} = 1$ and $x(\overline{N}_j) > x(\overline{N}_i)$. Then there exists $y \in R_s(x)$ that can be reached via updates only by players i and j such that $y_j = 0$, $y_i = \max\{0, 1 - x(N_i) + x_j\}$.*

Proof. If $x_j = 0$, let player i revise and denote the resulting profile y , which satisfies the statement. If $x_j > 0$, then let players j and i revise once in sequence and denote the resulting profile by x' . Clearly, $x'(\overline{N}_j) > x'(\overline{N}_i)$ holds, as only i and j updated. Furthermore, we have $x'_i = b_i(x')$. If $x'_i = 0$, then we have $x'(N_i) \geq 1$, and hence $x'(N_j) > 1$, meaning that $x'_j = b_j(x') = 0$, and the statement holds for $y = x'$.

If $x'_i > 0$, then we must have $x'(\overline{N}_j) > x'(\overline{N}_i) = 1$. Consider a sequence of states $((x^0, i^1), (x^1, i^2), \dots, (x^{2d-1}, i^{2d})) \in \Omega^{2d}$ such that $x^0 = x'$, $i^t = j$ if t is odd and $i^t = i$

if t is even, and for every $t \geq 1$ we have $x^t = (b_{it}(x^{t-1}), x_{-it}^{t-1})$. Since $x'(\bar{N}_j) > 1$ we have $x_j^1 < x_j^0$, and $x^1(\bar{N}_i) = 1 - x_j^0 + x_j^1$, hence player i will increase his effort by the same amount as player j 's decrease, i.e., $x_i^2 - x_i^1 = x_j^0 - x_j^1$. Since $x^2(\bar{N}_j) > x^2(\bar{N}_i) = 1$, the same argument can be repeated. After at most $2d$ steps, player j will reach 0 and player i will have taken over all of player j 's efforts. ■

Lemma 3.4.2 is illustrated in the following example.

Example 3.4.3. Consider again the network structure and effort profile of Example 3.3.7 (see Figure 3.1). Since $x(\bar{N}_2) = 1 < x(\bar{N}_3) = 1.2$, the conditions of Lemma 3.4.2 apply. It is easy to see that letting players 3 and 2 best-respond to x in sequence leads to the profile $x^2 = (0.1, 0.6, 0.2, 0.3, 0.4, 0.4, 0.1)$, with $x^2(\bar{N}_2) = 1 < x^2(\bar{N}_3) = 1.2$. Doing so again gives $x^4 = (0.1, 0.8, 0, 0.3, 0.4, 0.4, 0.1)$, satisfying $x_2^4 = 1 - x^0(N_2) + x_3^0 = 0.8$ and $x_3^4 = b_3(x^4) = 0$.

Repeated application of Lemma 3.4.2 allows it to be generalized in Lemma 3.4.4 to one player and several neighbors with higher consumption.

Lemma 3.4.4 (Rising inequality in neighborhoods). *Let $s \in \tilde{S}$ be such that $s_{-1} = b_{-1}$. Let $x \in X$, $i \in I_{-1}$, and $J \subseteq I_{-1} \cap N_i$ be given such that for every $j \in J$ it holds that $x(\bar{N}_j) > x(\bar{N}_i)$. Then there exists $y \in R_{s^*}(x)$ that can be reached via updates only by players in $\{i\} \cup J$ such that $y(J) = 0$ and $y_i = \max\{0, 1 - x(N_i) + x(J)\}$.*

Proof. Fix $j \in J$. As shown in the proof of Lemma 3.4.2, a sequence of alternating best-response updates by players i and j allows the game to reach an action profile x' for which we have $x'_j = 0$. We now show that for every $j' \in J \setminus \{j\}$ it still holds that $x'(\bar{N}_{j'}) > x'(\bar{N}_i)$. We have

$$x'(\bar{N}_{j'}) = x'(\bar{N}_{j'} \setminus \{i, j\}) + x'_i + g_{jj'}x'_j = x(\bar{N}_{j'}) + (x'_i - x_i) + g_{jj'}(x'_j - x_j),$$

and

$$x'(\bar{N}_i) = x'(\bar{N}_i \setminus \{i, j\}) + x'_i + x'_j = x(\bar{N}_i) + (x'_i - x_i) + (x'_j - x_j).$$

Since $x'_j - x_j \leq 0$ and by $g_{jj'} \in \{0, 1\}$ we have $(g_{jj'} - 1)(x'_j - x_j) \geq 0$, hence we have $x'(\bar{N}_{j'}) - x'(\bar{N}_i) \geq x(\bar{N}_{j'}) - x(\bar{N}_i) > 0$.

Therefore, the same method used in the proof of Lemma 3.4.2 can be repeated until every player in J exerts an effort of 0, and player i takes over the efforts of every $j \in J$. ■

Lemma 3.4.4 is a direct consequence of Lemma 3.4.2, as differences in consumptions between two neighboring players are either preserved or magnified when the player with

lower consumption takes over the efforts of another of his neighbors. The two lemmas together describe a progression towards increasing inequality in efforts between myopic players, given an inequality in consumption.

We now introduce two notions to describe sets of contributing myopic players.

Definition 3.4.5 (Islands, factions, partnerships). Let $x \in X$. A connected, non-empty set of myopic players $L \subseteq I_{-1}$ is called an *island* in x if

- for every $j \in L$ it holds that $x_j > 0$,
- for every $i \in N_L \setminus \{1\}$ it holds that $x_i = 0$.

A possibly disconnected L satisfying the same properties is called an *archipelago*, consisting of unions of islands. Furthermore, a connected $F \subseteq I_{-1}$ is called a *faction* in x if F is an island and for every $j \in F$ it holds that $x_j = b_j(x)$. Similarly, an archipelago that satisfies the same properties is called a *partnership*.

Let the collection of islands in x be denoted by $\mathcal{L}(x)$, the collection of factions by $\mathcal{F}(x)$, and the collection of partnerships by $\mathcal{B}(x)$.

An island is a set of contributing myopic players and whose myopic neighbors outside the set are free riders. A faction is an island where every player is at his best response, and therefore has a consumption of exactly 1. Archipelagos are unions of islands, while partnerships are unions of factions. Clearly, every myopic player in every effort profile is either a member of an island or a free rider. For any $s_1 \in \tilde{S}_1$, in every effort profile in either of the three sets, X^* , X_+^* , $A(s_1, b_{-1})$, every myopic player is either a member of a faction or a free rider.

Notice that given a profile $x \in X$, for every partnership $B \in \mathcal{B}(x)$ it holds that $M \in \mathcal{M}_x(B)$ if and only if M is a maximal independent set of B .

The following lemma shows an important property, that every island can give rise to a partnership through a series of best-responses.

Lemma 3.4.6 (From islands to partnerships). *Let $x \in X$ and let s be given such that $s_{-1} = b_{-1}$. Let $L \in \mathcal{L}(x)$ be such that there is a player $i \in L$ for which $x(N_i \cap \{1\}) < 1$. Then, there exists an $x' \in R_s(x)$ with updates only by players belonging to L , and a partnership $B \in \mathcal{B}(x')$ such that $B \subseteq L$ and for every $j \in L$ we have that $x'_j = b_j(x')$*

Proof. Similarly to the proof of Lemma 3.4.1 suppose that only the players in L are allowed to move. As before, by the fact that this is a best-response potential game, together with the results of Voorneveld (2000) guarantee that the players will reach an

action profile $x' \in X$ where for every $j \in L$ it holds that $x'_j = b_j(x')$. The condition $x(N_i \cap \{1\}) < 1$ for at least one $i \in L$ guarantees that at least one member of L will have a positive effort level. It follows that there must be a $B \subseteq L$ that is a partnership in x' . ■

The main result of this section shows that in recurring effort profiles no member of a partnership and its neighborhood will change his efforts if the partnership lies outside player 1's neighborhood.

Proposition 3.4.7 (Members of distant partnerships never change their efforts). *Let $s \in \tilde{S}$ be such that $s_{-1} = b_{-1}$, let $x \in X$, and let x be a recurrent profile in s . Then, for every $B \in \mathcal{B}(x)$ such that $B \cap N_1 = \emptyset$, for every $j \in \overline{N}_B$ and every $x' \in [x]_s$ it holds that $x_j = x'_j$.*

Proof. Suppose that $|B| = 1$, i.e. B is a one-player faction, $\{j\}$. It must hold that $x_j = 1$. Then, no neighbor of j can increase his efforts before j decreases his own. However, j cannot decrease his own efforts before a neighbor of his increases. Hence, no effort profile can be reached from x where any player in $\overline{N}_B = \overline{N}_j$ has a different effort level.

Now let $m > 1$ and suppose that the statement holds for every partnership with fewer than m members. Let $|B| = m$ and suppose that we have an $x' \in [x]_s$ and a $j \in \overline{N}_B$ such that $x'_j \neq x_j$. Since the members of B are playing their best-responses in x , every history through which x' can be reached from x must include an effort profile y and a player $i \in N_B$ such that $y_i = b_i(y) > 0$ but $y_j = x_j$ for every $j \in \overline{N}_B \setminus \{i\}$.

It must hold that $y(\overline{N}_i) = 1$, as $y_i = b_i(y) > 0$. Furthermore, it must be true that $B \setminus N_i \neq \emptyset$, otherwise $y(N_i) \geq 1$ and $b_i(y) = 0$, a contradiction. Moreover, for every $j \in N_i \cap B$ we have $y(\overline{N}_j) = x(\overline{N}_j) + y_i > 1$. Hence, by Lemma 3.4.4 there exists $y' \in R_s(y)$ such that $y'(N_i \cap B) = 0$. Notice that i is not in the neighborhood of $B \setminus N_i$. Hence, the remainder of the original partnership is an archipelago in y' .

By Lemma 3.4.6 there exists $y'' \in R_s(y')$ and a partnership $B' \subseteq B \setminus N_i$ in y'' . However, since $y'' \in [x]_s$, and since $|B'| \leq |B \setminus N_i| < m$ we must have $x_j = y''_j$ for every $j \in \overline{N}_{B'}$. This contradicts the assumption that B is a partnership of m players in x and that x is recurrent. ■

The intuition behind Proposition 3.4.7 is as follows: Any one-player faction outside player 1's neighborhood will remain one indefinitely, as it is simply an expert with an effort of 1, surrounded by free riders. Neither will change their effort levels before the other does so. A two-player faction, if it breaks can lead to a one-player partnership,

and therefore it either cannot break in a recurrent set of effort profiles, or a profile with a two-player faction cannot be recurrent. By induction the same argument can be extended to partnerships of more players.

3.5 Reaching a single effort profile

The main result of this section is the following:

Theorem 3.5.1. *There exists a $\bar{\delta} < 1$ such that for every $\delta \in (\bar{\delta}, 1)$ and every SSPE strategy profile s^* it holds that every recurring set is a singleton, and for any given starting state, ω , one of the recurring sets is reached with probability 1.*

Naturally, Theorem 3.5.1 implies that every recurrent effort profile belongs to X^* , furthermore some element of X^* is reached almost surely from any starting state.

The proof of Theorem 3.5.1 needs a number of steps. We introduce a binary relation on effort profiles regarding their pillar sets.

Definition 3.5.2 (Pillar-dominance). Let $x, y \in X'$, and $J \subseteq I_{-1}$ be given. We say that y *pillar-dominates* x in J , denoted by $x \triangleleft_J y$, if

$$\min_{M_x \in \mathcal{M}_x(J)} |M_x| \leq \min_{M_y \in \mathcal{M}_y(J)} |M_y|.$$

Notice that $x \triangleleft_J y$ holds automatically if $J \cap I_{-1}^+(x) = \emptyset$.

The idea behind the notion of pillar-dominance is to introduce an upper bound on the effort levels in x with information regarding y . For example, if we have $x \triangleleft_{N_1} y$, and y is an experts and free riders Nash equilibrium with ℓ experts in N_1 , then we know that $x(N_1) \leq \ell$. The whole proof idea of Theorem 3.5.1 makes use of this relation. Crucially, notice that \triangleleft_J is a transitive relation.

We now turn to the first in what will be a series of statements that set up the proof of Theorem 3.5.1. The proof strategy is as follows: Starting from a non-singleton recurring set under an SSPE we construct an effort profile where all myopic players are at best response, while at the same time it is no worse for the farsighted player than any other in the recurring set, and we also provide a corresponding strategy for player 1 by which he is guaranteed to reach it, which leads to contradiction.

To achieve this, given an SSPE s^* , and a non-singleton recurring set $[x]_{s^*}$ we need a profile $y \in [x]_{s^*}$ with four properties:

1. $y \in X^*$,

2. $x^{\max} \triangleleft_{N_1} y$,
3. every faction in y is either completely inside or completely outside of N_1 ,
4. every faction in y inside N_1 is a clique,

where $x^{\max} \in \operatorname{argmax}_{x' \in [x]_{s^*}} \pi_1(x')$ denote one of the best profiles for player 1 of the recurring set.

Properties 1, 3, and 4 together describe an effort profile where, for some $y_1 \in [0, 1]$, the total contribution of the farsighted player's neighbors is $y(N_1) = m(1 - y_1)$, with m denoting the number of contributing cliques in player 1's neighborhood. Property 2 ensures that $m \geq x^{\max}(N_1)$. The proof strategy is to construct a strategy for player 1 which leads him to an effort profile where his consumption is m while his efforts are 0 with probability 1. If such a strategy exists, then s_1^* cannot be part of an SSPE.

In what follows we first present a technical lemma, then prove that such a y can indeed be reached, always adding one property at a time. Finally, we prove Theorem 3.5.1. For $i \in I_{-1}$, let $L(i, x)$ denote the island that i belongs to, if any.

For $x \in X$ let $\bar{L}(x) = \{i \in I_{-1}^+(x) : L(i, x) \text{ is not a clique}\}$ denote the set of contributors for whom the islands they belong to are not cliques.

Lemma 3.5.3. *Fix $s \in \tilde{S}$ such that $s_{-1} = b_{-1}$. Let $x \in X$ be a recurrent effort profile in s such that $x_1 < 1$. Let $i \in I_{-1}$ be such that $x(\bar{N}_i) < 1$. There exists $y \in X_+^* \cap R_s(x)$ that can be reached by updates only by myopic players such that*

1. $y_i = b_i(y)$,
2. $x \triangleleft_{N_1} y$, and
3. $|\bar{L}(x)| \geq |\bar{L}(y)|$.

The proof can be found in the Appendix.

The next lemma states that from any recurring effort profile the game reaches a state where every myopic player is at best response, the ensuing effort profile pillar-dominates the original one, and the number of myopic contributors who are not in cliques does not increase. This establishes properties 1 and 2 that we will use for the proof of Theorem 3.5.1, and builds towards property 4.

Lemma 3.5.4. *Fix $s \in \tilde{S}$ such that $s_{-1} = b_{-1}$. If there exists an $x \in X$ with $|[x]_s| > 1$, and $x_1 < 1$, then there exists an action profile $y \in X^* \cap R_s(x)$, that can be reached via updates only by myopic players, such that $x \triangleleft_{N_1} y$ and $|\bar{L}(x)| \geq |\bar{L}(y)|$.*

The proof can be found in the Appendix. Lemma 3.5.4 relies on the repeated use of Lemma 3.5.3 to advance the game into further effort profiles that both pillar-dominate the original, ensure that all myopic contributors are at best response, and that the number of contributors not in cliques does not increase. The fact that the game is a best-response potential game guarantees that the myopic players will eventually stop revising and they all reach their best responses. We demonstrate the meaning of Lemma 3.5.4 with an example.

Example 3.5.5. Consider again the network of Example 3.3.7. Suppose that the profile $x = (0.1, 0.4, 0.4, 0.3, 0.4, 0.4, 0.1)$ (see Figure 3.1) is recurring. Then, the effort profile we get by allowing player 2 to usurp players 3 and 7 in succession is also recurring (shown in Figures 3.4 and 3.5). Notice that in the new profile, player 4 can also usurp player 5 which would leave player 4 with an effort of 0.9. Subsequently, player 6 can move to his best response, also exerting 0.9 effort, creating three one-player factions (Figure 3.6). Every player is at best response, every contributor is part of a one-player faction that is also a clique, and the only pillar set is $\{2, 4, 6\}$. Hence, the resulting profile pillar-dominates x .




Figure 3.4: The links between the myopic players and the effort profile of Example 3.5.5 after player 2 usurps player 3's efforts. Now rising inequality applies for players 2 and 7.

Notice that in Example 3.5.5 property 3 is vacuous and we got property 4 for free, neither of which is guaranteed to happen in the general case. Further examples (Examples 3.5.7 and 3.5.9) will be used to illustrate such cases. By its construction, Example 3.5.5 also demonstrates why showing that recurring sets include an effort profile with all four properties is a crucial step for proving Theorem 3.5.1. Once the farsighted player becomes active, it is clear that his best strategy is to move to zero,




Figure 3.5: The links between the myopic players and the effort profile of Example 3.5.5 after player 2 usurps player 3 and 7's efforts in succession.




Figure 3.6: The links between the myopic players and the effort profile of Example 3.5.5. Rising inequality allows the formation of an effort profile where the farsighted player's problem is greatly simplified.

which will be followed by players 2, 4, and 6 moving to 1, hence the best profile for player 1 is reached, and x could not have been recurring under an SSPE.

As mentioned, generally properties 1 and 2 do not imply properties 3 and 4. As a next step we show that once we reached a profile that satisfies 1 and 2, then we can reach a profile that satisfies 1, 2, and 3.

Lemma 3.5.6. *Fix $s \in \tilde{S}$ such that $s_{-1} = b_{-1}$. If there exists an $x \in X$ with $||x|_s| > 1$ and $x_1 < 1$, then there exists an action profile $y \in [x]_s \cap X^*$, such that $x \triangleleft_{N_1} y$ and for every $F \in \mathcal{F}(y)$ it holds that either $F \subseteq N_1$ or $F \cap N_1 = \emptyset$, and $|\bar{L}(x)| \geq |\bar{L}(y)|$.*

The proof can be found in the Appendix. The idea behind Lemma 3.5.6 is to start with an effort profile where all myopic players are at best response, which is possible

due to Lemma 3.5.4. If there is a faction with members both inside and outside of N_1 , then we check how the farsighted player reacts. If he does not revise his efforts, the recurrent set must be a singleton. If he decreases, we apply Lemma 3.5.4 to get back to a state with all myopic players at their best responses, with the farsighted player now having a lower contribution. If he increases, then a faction with members both inside and outside his neighborhood will splinter into smaller ones, with one splinter left completely outside his neighborhood. Thus, a contradiction is reached by Proposition 3.4.7, as the original profile could not have been recurrent. This idea is demonstrated by a simple example.

Example 3.5.7. Consider the line network, and the initial effort profile, x^0 of Example 3.3.10 (shown in Figure 3.2). If this profile is recurring, then the recurring set is a singleton: If the farsighted player would not deviate, i.e. $s_1(x^0) = 0$, we have $[x^0]_s = \{x^0\}$, a singleton. If the farsighted player would deviate, i.e. $s_1(x^0) > 0$, then by the rising inequality between pairs lemma (Lemma 3.4.2) we can reach a profile where player 3 is an expert and player 2 is a free rider. There is no way back to x . Therefore, it can be easily seen that any candidate non-singleton recurring set in the line network must contain at least one profile in which player 3 is a free rider, meaning that the only partnership is completely inside N_1 .

Finally, we add property 4. We show that a profile can be reached such that every faction inside the farsighted player's neighborhood is a clique while also satisfying properties 1, 2, and 3. For $x \in X$ let $\mathcal{F}_1(x) = \{F \in \mathcal{F}(x) : F \subseteq N_1\}$ denote the collection of factions that are entirely inside player 1's neighborhood.

Lemma 3.5.8. *Fix $s \in \tilde{S}$ such that $s_{-1} = b_{-1}$ and for every $x \in X$ it holds that $s_1(x) \leq 1$. If there exists an $x \in X$ with $[x]_s > 1$ and $x_1 < 1$, then there exists an action profile $y \in [x]_s \cap X^*$, such that $x \triangleleft_{N_1} y$, for every $F \in \mathcal{F}(y)$ we either have $F \subseteq N_1$ or $F \cap N_1 = \emptyset$, and for every $F \in \mathcal{F}_1(y)$ it holds that F is a clique.*

The proof can be found in the Appendix. The idea behind Lemma 3.5.8 is to start from a profile that satisfies the first three properties, and then, similarly to Lemma 3.5.6 we check what the farsighted player's choice would be. If he does not revise, then again, we are in a singleton recurring set. If he does, then we break up an original faction into smaller ones. Once the factions are small enough, i.e. they have one, two, or three players, they are guaranteed to be cliques. We demonstrate the idea with an example.

Example 3.5.9. Consider once more the network of Example 3.3.7 with the following effort profile: $x = (0.1, 0.4, 0.4, 0.1, 0.4, 0.4, 0.1)$ (Figure 3.7). Since $x \in X^*$, we have no inequality to exploit with the rising inequality lemmas.




Figure 3.7: The links between the myopic players and the effort profile of Example 3.5.9. All players are at best response but the only faction is not a clique. There is no inequality, none of the myopic players can usurp any other.

The only way for the game to escape x is by the farsighted player's deviation. If he does not deviate, the recurring set containing x must be a singleton. If he does, then, depending on the order of players revising, inequalities can appear that we can exploit to fragment the faction $\{2, 3, 4, 5, 6, 7\}$ into smaller ones.

Suppose that player 1 increases his effort by 0.2 to exert 0.3, and that this change is followed by player 3's revision (a decrease to 0.2), then player 7's (a decrease to 0, since he cannot go below that). Figure 3.8 shows the resulting effort profile which we call x' .

Notice that $x'(\bar{N}_2) = 0.9 < x'(\bar{N}_3) = 1 < x'(\bar{N}_7) = 1.1$, and hence the rising inequality in neighborhoods lemma is now applicable for player 2 and his only contributing myopic neighbor, 7. Hence, from x' we can reach an effort profile in which player 2 is a one-player faction with an effort of 0.7, and players 3 and 7 are both free riders. Following this, it is easy to see that the profile $y = (0.3, 0.7, 0, 0.7, 0, 0.7, 0)$ is also reachable by player 4 usurping player 5 followed by player 6 moving to his best reply. Note that any increase of player 1 allows the emergence of the same three one-player factions via similar methods.

If player 1 decreases instead of increasing, we can also reach the same outcome. Suppose that player 1 decreases his effort to 0, and his revision is followed by player 2's (an increase to 0.5) and then by player 4's (an increase to 0.2). Figure 3.9 shows this effort profile which we call x'' .




Figure 3.8: The links between the myopic players and the effort profile of Example 3.5.9, after an increase by player 1, and two subsequent decreases by players 3 and 7. Rising inequality applies for player 2 and 3.




Figure 3.9: The links between the myopic players and the effort profile of Example 3.5.9, after a decrease by player 1, and two subsequent increases by players 2 and 4. Rising inequality applies for player 2 and 3.

Since $x''(\bar{N}_2) = 1 < x''(\bar{N}_3) = 1.1$, the rising inequality between pairs lemma applies for players 2 and 3. Moreover, player 2 can usurp player 7 as well, again forming a one-player faction. This can be followed by player 4 usurping 5, and then 6 moving to his best response to reach another effort profile with three one-player factions, $(0, 1, 0, 1, 0, 1, 0)$

In the following proof we assume that an SSPE strategy profile can permit non-singleton recurrent sets, which makes the four lemmas applicable, then arrive at a contradiction.

Proof of Theorem 3.5.1. Assume that we have an $x \in X$ and an $s^* \in S^*$ for which $|[x]_{s^*}| > 1$. Let $x^{\max} \in \operatorname{argmax}_{x' \in [x]_{s^*}} \pi_1(x')$ denote one of the best effort profiles for

player 1. Note that since $|[x]_{s^*}| > 1$, player 1 has to move at least once. Since $N_1 = \emptyset$ leads to a trivial contradiction, we assume $N_1 \neq \emptyset$.

If $I_{-1}^+(x^{\max}) = \emptyset$, then we must have $x^{\max}(N_1) = 0$, and for every $i \in N_1$ we have $x^{\max}(N_i) \geq 1$, otherwise any neighbor revising would improve player 1's instantaneous payoff. Clearly, by allowing all myopic players outside N_1 to reach their best responses as well we can reach $x^* \in X^*$ with $x_1^* = x_1^{\max}$, $x^*(N_1) = 0$, and $\pi_1(x^*) = \pi_1(x^{\max})$. If this is the best that player 1 can achieve, then in an SSPE we must have $s_1^*(x^*) = x_1^{\max} = x_1^* = 1$, since, in absence of any efforts from his neighbors, any change in his own effort would decrease his own instantaneous payoffs. This means that the absorbing set is a singleton, a contradiction.

If $I_{-1}^+(x^{\max}) \neq \emptyset$, then let $M \in \mathcal{M}_{x^{\max}}(N_1)$. Then, for every $x' \in [x]_{s^*}$ we have

$$\pi_1(x') \leq f_1(|M|).$$

Furthermore, since the farsighted player changes his efforts at least once, changing his instantaneous payoff, we must have a $x' \in [x]_{s^*}$ with $\pi_1(x') < f_1(|M|)$. We will show that there exists $y \in [x]_{s^*}$ for which $\pi_1(y) \geq f_1(|M|)$ and $y \in X^*$, with $y_1 = 0$. This leads to contradiction, as under these circumstances, playing $s_1^*(y) > 0$ would not be an SSPE strategy, and we have $[y]_{s^*} = \{y\}$.

Let $y^0 \in [x]_{s^*}$ denote the effort profile such that $y^0 \in X^*$, for every $F \in \mathcal{F}(y^0)$ with $F \cap N_1 \neq \emptyset$ it holds that $F \subseteq N_1$ and F is a clique, and there does not exist a $y' \in [x]_{s^*}$ with the same properties for which $|\mathcal{F}_1(y')| > |\mathcal{F}_1(y^0)|$. Let $|\mathcal{F}_1(y^0)| = m^0$. By Lemma 3.5.8, such a y^0 must exist, and $m^0 \geq |M|$. In other words, y^0 is a profile where m^0 cliques of myopic players contribute in player 1's neighborhood with each clique's total effort equaling $1 - y_1^0$.

Consider the following strategy $s_1 \in \tilde{S}_1$ for $x' \in X$:

$$s_1(x') = \begin{cases} \max\{x'_1 - 1/d, 0\} & \text{if } x' \in X^*, \\ x'_1 & \text{otherwise.} \end{cases}$$

We now show that if, instead of s_1^* , player 1 adopts s_1 as soon as the game reaches the effort profile y^0 , the game will reach an effort profile where player 1's instantaneous payoff is at least $f_1(m^0)$, and all myopic players are at best response. If $y_1^0 = 0$, this is already achieved, hence $\pi_1(y^0) = f_1(m^0) \geq f_1(|M|)$, meaning that in any SSPE, s^* , we must have $s_1^*(y^0) = 0$, contradicting $|[x]_{s^*}| > 1$.

Let $y' = (s_1(y^0), y_{-1}^0)$. What we show next is that for every $i \in N_1$ such that $y'_i = 0$ and $b_i(y') > y'_i$, \bar{N}_i contains the entirety of one clique of myopic players of positive

effort and no other myopic players of positive effort. If \bar{N}_i includes no contributing myopic players, then letting exactly one player in every faction $F \in \mathcal{F}_1(y^0)$ move to his best response after y' , and finally letting i do the same would produce $m^0 + 1$ factions in N_1 , contradicting the choice of y^0 .

If for some $F \in \mathcal{F}_1(y^0)$ we have $\bar{N}_i \cap F \neq \emptyset$, but $F \not\subseteq \bar{N}_i$, then take $j \in F \setminus \bar{N}_i$. Since $j \in F$ we have $y'(\bar{N}_j) = 1 - 1/d$, and since i is not at his best response, but was before player 1 decreased, we have $y'(\bar{N}_i) = 1 - 1/d$ as well. Now, let y'' denote the profile we get from y' by letting players i and j move successively to best response. Then, for every $j' \in F \cap \bar{N}_i$ we have $y''(\bar{N}_{j'}) = 1 + 1/d$, while $y''(\bar{N}_i) = y''(\bar{N}_j) = 1$. Hence, by Lemma 3.4.4, there exists a profile $\hat{y} \in R_s(y^0)$, with $\hat{y}(F \cap N_i) = 0$. From here, by letting players i, j , and one player from every other clique $F' \in \mathcal{F}_1(y^0)$, $F' \neq F$ move to best response we get a profile with $m^0 + 1$ cliques of myopic players contributing, contradicting the choice of y^0 .

Hence, if $i \in N_1$ with $b_i(y') > y'_i$, then \bar{N}_i contains the entirety of one clique. It cannot contain any other myopic contributors, since for every $F \in \mathcal{F}_1(y^0)$ we have $y'(F) = 1 - y_1^0$, meaning that $y'((\bar{N}_i \cap F) \cup \{1\}) = 1 - 1/d$. If i had any other contributing myopic neighbors, he would have a consumption of at least one.

Therefore, once player 1 adopts the strategy s_1 , every time a myopic player updates, a faction is created that is also a clique. If $y_1^0 > 0$, then let $s = (s_1, b_{-1})$, and let $y^1 \in R_s(y^0) \cap X^*$ be such that $y_1^1 = y_1^0 - 1/d$. Notice that $y^0 \notin R_s(y^1)$, because player 1 never increases his efforts while playing s_1 . As factions cannot split, nor can they merge, y^1 will have m^0 cliques of myopic players contributing with every clique's effort totaling $1 - y_1^1 = 1 - y^0 + 1/d$. The same arguments can be repeated for y^2 , and so on, until we have $y_1^k = 0$ for some $k \leq d$. By continuing along this strategy, player 1 slowly decreases his efforts, while the cliques of myopic players in his neighborhood increase. If we let y denote any profile where player 1's effort is 0, and the neighboring cliques' efforts are 1, then clearly we have $\pi_1(y) = f_1(m^0) \geq f_1(|M|)$. Hence, the instantaneous payoff $f_1(m^0)$ is reached with certainty.

With that we achieve contradiction, as a player with a high enough δ will choose s_1 instead of s_1^* , meaning that s^* is not an SSPE. ■

We conclude this section by a technical result regarding the limit of expected utilities for a given starting state.

Proposition 3.5.10. *There exists a $\bar{\delta}$ such that for every $\delta \in (\bar{\delta}, 1)$, every $\omega \in \Omega$, and every SSPE profile s^* , $\lim_{t \rightarrow \infty} u_1^t(\omega, s^*)$ exists.*

Proof. Let $\omega \in \Omega$ and an SSPE s^* be given. Let $Q^t(x)$ denote the probability that the game is in effort profile x in period t if the game starts in state ω and players follow strategy profile s^* . By Theorem 3.5.1 it follows that with probability 1, the game proceeds to an absorbing effort profile. By the fact that no absorbing profile can be reached from any other, for every $x \in X$ the limit $\lim_{t \rightarrow \infty} Q^t(x)$ exists and can only be positive for $x \in A(s^*)$.

Hence

$$\begin{aligned} \lim_{t \rightarrow \infty} u_1^t(\omega, s^*) &= \lim_{t \rightarrow \infty} \sum_{x \in X} Q^t(x) \pi_1(x) = \\ \lim_{t \rightarrow \infty} \sum_{x \in A(s^*)} Q^t(x) \pi_1(x) &= \sum_{x \in A(s^*)} \lim_{t \rightarrow \infty} Q^t(x) \pi_1(x) \end{aligned}$$

also exists. ■

3.6 Lower and upper bounds on long-run payoffs

In this section we compare networks based on the opportunities they offer. Given a network, the objective is to identify which connections are valuable for the farsighted player, and which connections are harmful. The most direct application of this analysis is in models where network formation is an option, i.e. the farsighted player or a planner can form and sever links for a cost. Such is the case in most social networks, computer networks, and many biological networks.

In Section 3.5 we have shown that equilibrium strategies lead to convergence to an absorbing effort profile in case the farsighted player is patient. The utility, i.e. the discounted sum of expected instantaneous payoffs, of these players is almost completely decided by the instantaneous payoffs they receive in the absorbing effort profiles they can reach by their chosen SSPE strategy. However, since the set of SSPE strategies depends on the network structure and on the initial effort profile of the players in a convoluted way, we conduct comparative statics of networks based on lower and upper bounds of the long-run payoffs that the farsighted player can achieve.

The formal definitions are as follows. For a network G let $\bar{\delta}(G)$ denote threshold value of δ for which Theorem 3.5.1 holds.

Definition 3.6.1. Let $\delta > \bar{\delta}(G)$ and let

- $\underline{u}(G) = \inf_{\omega \in \Omega, s^* \in S^*} \lim_{t \rightarrow \infty} u_1^t(\omega, s^*)$.
- $\bar{u}(G) = \sup_{\omega \in \Omega, s^* \in S^*} \lim_{t \rightarrow \infty} u_1^t(\omega, s^*)$.

Notice that both values exist as for every SSPE strategy profile Theorem 3.5.1 ensures that some absorbing effort profile is reached in the long run. The values $\underline{u}(G)$ and $\bar{u}(G)$ are lower and upper bounds on the long-run expected payoff of the farsighted player who knows only the network G , his position in it, but not the initial state, and thus has not yet chosen an equilibrium strategy. Clearly, $\underline{u}(G) \leq \bar{u}(G)$ for every $G \in \mathcal{G}$. For a patient farsighted player as assumed by Theorem 3.5.1 and Definition 3.6.1, the long-run expected payoff almost completely determines the utility of the farsighted player and hence $\underline{u}(G)$ and $\bar{u}(G)$ are close estimates on the infimum and supremum of $U_1(\omega, s^*)$, respectively. More precisely, for every ε there exists a $\bar{\delta}$ such that for every $\delta > \bar{\delta}$ we have $|\underline{u}(G) - (1 - \delta) \inf_{\omega \in \Omega, s^* \in S^*} U_1(\omega, s^*)| < \varepsilon$ and $|\bar{u}(G) - (1 - \delta) \sup_{\omega \in \Omega, s^* \in S^*} U_1(\omega, s^*)| < \varepsilon$.

By using the two bounds we can make qualitative comparisons between networks from the perspective of the farsighted player. Without further information on the initial state of the game, every farsighted player should place his evaluation of a network within an ε -range of the interval between the two bounds.

Let $\mathcal{G} = \{G \in \mathbb{R}^{n \times n} : \forall i, j \in I \quad g_{ij} = g_{ji}, g_{ii} = 1\}$. The main results of this section concern the comparative statics of the bounds as links are added to the network. We first show that the two bounds coincide in networks where the farsighted player is linked to every myopic opponent. Let $\mathcal{G}_1 = \{G \in \mathcal{G} : \forall i \in I_{-1} \quad g_{1i} = 1\}$ denote the set of such networks with n players.

Let $\mathcal{M}_1^{ind}(G)$ denote the set of independent sets of myopic players in the farsighted player's neighborhood in network G .

Proposition 3.6.2. *There exists $\bar{\delta}$ such that for every $G \in \mathcal{G}_1$ and every $\delta > \bar{\delta}$ it holds that*

$$\underline{u}(G) = \bar{u}(G) = f_1\left(\max_{M \in \mathcal{M}_1^{ind}(G)} |M|\right).$$

The proof of Proposition 3.6.2 can be found in the Appendix. It constructs a strategy that achieves a long-run payoff equal to the network's upper bound for every initial state. If the farsighted player is sufficiently patient, then his SSPE strategy must also achieve this long-run payoff. This strategy consists of two simple stages: (1) establishing *dependence* by providing for the myopic neighbors, i.e. exerting an effort of 1 until every neighbor decreases to 0, and (2) *withdrawal* of efforts and forcing an increase, i.e. moving to 0 effort. The steps (1) and (2) may be repeated if not enough players increase after the withdrawal stage.

The complexity of this strategy is rather low. As an immediate implication of the simplicity of this strategy, the values of the upper and lower bounds may be reached

even without the ability to compute SSPE strategies. We will generalize this point to the entire class of networks towards the end of this section. Note however, that the above strategy is not necessarily an SSPE itself, as there may be better strategies that guarantee the same outcome, e.g. by reaching the optimum sooner.

For now we turn to our comparison results. We first introduce the measures of networks by which we will characterize the values of the upper and lower bounds, respectively. We start with the upper bound.

Definition 3.6.3. Let $\bar{e}: \mathcal{G} \rightarrow \mathbb{N}$ be given by

$$\bar{e}(G) = \max_{M \in \mathcal{M}_1^{ind}(G)} |M|.$$

The value $\bar{e}(G)$ is the cardinality of a maximum independent set of the farsighted player's neighborhood in network G . It is the maximum number of players that the farsighted player can fully exploit simultaneously without any of the myopic players preferring to deviate. We now show a strictly monotone relationship between $\bar{e}(G)$ and $\bar{u}(G)$

Proposition 3.6.4. *For every $G \in \mathcal{G}$ it holds that*

$$\bar{u}(G) = \begin{cases} f_1(\bar{e}(G)) & \text{if } \bar{e}(G) > 0, \\ f_1(1) - c_1 & \text{if } \bar{e}(G) = 0. \end{cases}.$$

Proof. Recall the effort x^* from Proposition 3.3.4. It is a Nash equilibrium of the static game Γ such that players in a maximum independent set $M \in \mathcal{M}_1^{ind}$ exert an effort of 1 and players in $\bar{N}_1 \setminus M$ exert 0. Hence, it maximizes player 1's instantaneous payoffs over the set of effort profiles X' , meaning that $\bar{u}(G) = f_1(|M|) = f_1(\bar{e}(G))$, if $N_1 \neq \emptyset$, while $\bar{u}(G) = f_1(1) - c_1$ if $N_1 = \emptyset$. \blacksquare

Proposition 3.6.4 shows that the value of the upper bound is very easily calculated as the cardinality of a maximum independent set of the farsighted player's neighborhood. Hence, it can be done so even without knowledge of the entire network.

We now provide a similar characterization of $\underline{u}(G)$. We need a technical definition.

Definition 3.6.5 (Closed partnership). For $x \in X$, a partnership $B \in \mathcal{B}(x)$ is *closed* if for every $i \in N_B \setminus \{1\}$ it holds that $x(\bar{N}_i \cap B) \geq 1$.

A closed partnership is a partnership whose members provide at least 1 unit of effort for each of their free-riding neighbors. Note that a closed partnership may be a union of other, potentially not closed, factions or partnerships. Let the collection of closed partnerships in effort profile x be denoted by $\tilde{\mathcal{B}}(x)$.

Definition 3.6.6 (Blocking profile). The set $\tilde{X} = \{x \in X^*: x_1 = 1, \text{ and } \forall j \in I_{-1} \setminus N_1 \exists B \in \tilde{\mathcal{B}}(x) \text{ with } j \in \overline{N}_B\}$ is called the set of *blocking profiles*.

The action profiles included in \tilde{X} are the ones for which no player in $I_{-1} \setminus N_1$ reacts to player 1's changes, as all are either inside of or are neighbors of a closed partnership. Notice that \tilde{X} is always non-empty: setting a maximal independent set of $I \setminus \overline{N}_1$ and the farsighted player's exerted efforts to 1 and every other player's to 0 satisfies the definition. Furthermore, note that $x \in X^*$ with $x_1 = 1$ alone does not guarantee that every myopic player outside N_1 will be in a closed partnership or in the neighborhood of a closed partnership.

For $\tilde{x} \in \tilde{X}$ let $I(\tilde{x}) = \{i \in I: \text{ there is no } B \in \tilde{\mathcal{B}}(\tilde{x}) \text{ such that } i \in \overline{N}_B\}$ denote the players who are not blocked by a closed partnership. Note that $I(\tilde{x}) \subseteq \overline{N}_1$ as all other players are blocked by a closed partnership by Definition 3.6.6, but $I(\tilde{x})$ is not generally equal to \overline{N}_1 .

Definition 3.6.7. Let the function $\underline{e}: \mathcal{G} \rightarrow \mathbb{N}$ be given by

$$\underline{e}(G) = \min_{\tilde{x} \in \tilde{X}} \max_{M \in \mathcal{M}_1^{ind}(G|_{I(\tilde{x})})} |M|.$$

The value $\underline{e}(G)$ in Definition 3.6.7 gives the minimum number of players that the farsighted player can exploit no matter what the initial state is. The reason for this is the following. Since in a blocking effort profile \tilde{x} all players outside of $I(\tilde{x})$ are getting more than 1 unit of effort from some members of a closed partnership, we consider the subgraphs of G where these players are removed. From such a subgraph we take the cardinality of the maximum independent set, $\max_{M \in \mathcal{M}_1^{ind}(G|_{I(\tilde{x})})} |M|$, since, as seen previously in Propositions 3.3.4 and 3.6.4, getting 1 effort from the maximum independent set of his neighborhood is the best instantaneous payoff the farsighted player can hope for. In addition, this payoff is attainable, as we will show in the next proposition. By taking a minimum over effort profiles in \tilde{X} the value $\underline{e}(G)$ gives the exact number of units of the public good that is attainable for even the worst choice of \tilde{x} .

Similarly to Proposition 3.6.4 we characterize $\underline{u}(G)$ by $\underline{e}(G)$ and show that they are strictly monotonous transformations of each other.

Proposition 3.6.8. *For every $G \in \mathcal{G}$ it holds that*

$$\underline{u}(G) = \begin{cases} f_1(\underline{e}(G)) & \text{if } \underline{e}(G) > 0, \\ f_1(1) - c_1 & \text{if } \underline{e}(G) = 0. \end{cases}$$

The proof of Proposition 3.6.8 can be found in the Appendix.

By Proposition 3.6.8, the value of the lower bound hinges on the number of independent myopic neighbors in player 1's neighborhood that cannot be fully supported by closed partnerships outside player 1's neighborhood. Every such player is liable to exploitation by the farsighted player. As all players outside the farsighted player's neighborhood are in closed partnerships, Proposition 3.6.2 can be applied to ensure that even an unsophisticated farsighted player takes full advantage of a maximum independent set of unblocked myopic neighbors. Taking the minimum over all effort profiles with closed partnerships gives a lower bound of the farsighted player's possibilities.

Similarly to Proposition 3.6.2, the strategy to reach a long-run payoff that equals the lower bound is very simple. To guarantee this value, a patient farsighted player must switch between the dependence stage and the withdrawal stage, alternating between playing 1 and 0.

We now present the main comparison results. For a network G with $g_{ij} = 0$, let $G^{ij} = G + \{i, j\}$. G^{ij} is the network we get by adding the link $\{i, j\}$ to G .

Theorem 3.6.9. *If we add the link $\{i, j\}$ to the network G with $g_{ij} = 0$, the following comparisons hold:*

1. *If $i = 1$, we have $\bar{u}(G) \leq \bar{u}(G^{ij})$ and $\underline{u}(G) \leq \underline{u}(G^{ij})$. In addition, if $j \notin N_{N_1}$ we have $\bar{u}(G) < \bar{u}(G^{ij})$.*
2. *If $i, j \in N_1$, we have $\bar{u}(G) \geq \bar{u}(G^{ij})$ and $\underline{u}(G) \geq \underline{u}(G^{ij})$.*
3. *Otherwise, $\bar{u}(G) = \bar{u}(G^{ij})$.*

Proof. By Propositions 3.6.4 and 3.6.8, and by the strict monotonicity of f_1 , it is sufficient to show the analogue of Theorem 3.6.9 after replacing \bar{u} with \bar{e} and \underline{u} with \underline{e} in all three statements. The substitution also makes use of $f_1(1) > f_1(1) - c_1$ for the cases $\underline{e}(G) = 0$ and $\bar{e}(G) = 0$. Let \tilde{X}^{ij} be defined in the same manner as \tilde{X} for the network G^{ij} .

1. The first inequality follows due to the fact that the cardinality of the maximum independent set of the farsighted player's neighborhood cannot decrease by the arrival of a new neighbor. For the second inequality notice that $1 \in \{i, j\}$ implies that $\tilde{X} \supseteq \tilde{X}^{ij}$, since every closed partnership in $I_{-1} \setminus N_1(G^{ij})$ is a closed partnership in $I_{-1} \setminus N_1(G)$. On the other hand, for a fixed $\tilde{x} \in \tilde{X}^{ij}$ we have $\mathcal{M}_1(G|_{I(\tilde{x})}) \subseteq \mathcal{M}_1(G^{ij}|_{I(\tilde{x})})$. Both of these facts imply $\underline{e}(G) \leq \underline{e}(G^{ij})$. The final inequality is due to the fact that the maximum independent set of the farsighted player's neighborhood necessarily increases if we add a node without any links to the others.

2. As before, the first inequality follows very simply by the fact that the cardinality of a maximum independent set in a network cannot increase by the addition of a new link. For the second inequality, $i, j \in N_1$ implies $\tilde{X} = \tilde{X}^{i,j}$, since for every $\tilde{x} \in \tilde{X}$ we have $\tilde{x}(N_1) = 0$ and $\tilde{x}_1 = 1$. Therefore we only need to consider what happens to the maximum independent set of $N_1 \cap I(\tilde{x})$ for a fixed \tilde{x} , whose cardinality cannot increase by the addition of another link.
3. Naturally, any link that does not change $G|_{N_1(G)}$ does not affect its maximum independent set and hence, does not change $\bar{e}(G)$. ■

Theorem 3.6.9 fully characterizes the qualitative effects of an additional link on the upper bound. A new link is either beneficial or neutral to the farsighted player if it connects him to a new neighbor, harmful or neutral if it connects two of his neighbors, and neutral in all other cases. The first two statements are upheld very similarly for the lower bound. In all other cases the effect of a new link on the lower bound is ambiguous. This will be demonstrated by Examples 3.6.10 and 3.6.11.

The comparison results concerning the upper bound are achieved as straightforward consequences of its characterization as the cardinality of a maximum independent set of the farsighted player's neighborhood. Adding a player into the neighborhood increases the number of players whose efforts are enjoyed by the farsighted player, although depending on the connections the upper bound may not change. In case the new neighbor shares no connection with any of the old ones, the upper bound is guaranteed to increase. This has the implication that farsighted players prefer to link to distant players in the network rather than neighbors of his neighbors. Another way to improve his position is by weakening the connectedness of his existing neighborhood, as any severed link has a chance to increase the upper bound. All other single changes in the network are inconsequential on the upper bound, although this does not apply to a series of single changes.

Changes in the network affecting the neighborhood of the farsighted player have the same effect on the lower bound. Connecting two neighbors of the farsighted player has no effects on the set of effort profiles with closed partnership outside the farsighted player's neighborhood but it can increase the cardinality of maximum independent sets of the remaining players, increasing the lower bound. Adding a new neighbor reduces the possibilities to create closed partnerships while potentially increasing the cardinality of maximum independent sets, both effects work to increase the lower bound. In all other cases the new link changes the set blocking profiles in a way that can lead to an increase or to a decrease in the lower bound. This is demonstrated by the following

two examples which make up the remainder of this section.

Example 3.6.10. Consider a network of five players $\{1, 2, 3, 4, 5\}$ with the following links: $\{\{1, 2\}, \{2, 3\}, \{1, 4\}, \{4, 5\}\}$. In this network, the value of the minimum exploitation function is 0, and hence the lower bound is $f_1(1) - c_1$. Figure 3.10 shows the network and an effort profile with two closed partnerships which block both myopic players from exerting positive effort.




Figure 3.10: The network blocking effort profile of Example 3.6.10. Players 2 is blocked by player 3's effort while player 4 is blocked by player 5. Player 1's best option is to exert an effort of 1.

Adding the link $\{3, 5\}$ eliminates this blocking effort profile as players 3 and 5 both exerting an effort of 1 is no longer stable for myopic deviation. Hence, at least one of players 2 and 4 will always be free to provide for player 1, hence the value of the lower bound increases to $f_1(1)$. The resulting network with a new blocking profile is shown in Figure 3.11




Figure 3.11: The network blocking effort profile of Example 3.6.10 after connecting the two non-neighbors of player 1. Players 2 and 4 cannot be blocked at the same time.

Now consider connecting a player in the neighborhood of the farsighted player with one outside, e.g. adding the link $\{2, 5\}$. Then, the profile in which player 5 plays 1 and all other myopic players play 0 once again blocks both neighbors of the farsighted

player. Hence, adding a link between two non-neighbors can benefit the farsighted player, while adding a link between a neighbor and a non-neighbor may harm him.

Example 3.6.10 shows two simple cases through which the lower bound can be influenced. Linking the farsighted player's neighbors to non-neighbors can make it easier to achieve blocking effort profiles due to more links, while linking two non-neighbors to each other can make it harder, as closed partnerships will be harder to find.

While the above reasoning holds for many networks, there are examples where connecting a neighbor to a non-neighbor, or deleting a link between two non-neighbors both have the potential to increase the lower bound if the original blocking profile is sensitive for any changes in the network. This is shown in the following example.

Example 3.6.11. Consider a game with 17 players. The farsighted player is linked to 10 myopic players, none of which have links between them. The remaining 6 myopic players are grouped into two cliques of three players, but there are no links between the cliques. One of the farsighted player's neighbors has no other connections. The 9 other neighbors are connected to three other players each: one player from one clique and two players of the other. No two of these 9 players have the same set of neighbors.

In this network, the profile where the 6 clique members are playing $1/3$ each blocks the 9 neighbors, and this is the only blocking profile that will block all of them, while the 10th cannot be blocked in any way, hence the value of the lower bound is $f_1(1)$. Figure 3.12 shows this network and blocking profile.

Suppose that we remove a link one of the three cliques, the one with which each of the 9 neighbors has one connection. We thereby form a line of three players. As a result, the blocking profile of Figure 3.12 will no longer be admissible, since the clique will not be a partnership. The new worst blocking profile is one in which the center player of the line is a free rider, the other two contribute 1, while the other clique has two free riders and one player contributing 1. The value of the lower bound therefore increases to $f_1(2)$. The new blocking profile is shown on Figure 3.13.

Finally, suppose that we take the original network and add a link between the 10th neighbor to a member of the clique with which the 9 neighbors have one connection. Then, the original blocking profile of Figure 3.12 will again be inadmissible as the equal $1/3$ - $1/3$ - $1/3$ distribution only gives the 10th neighbor $1/3$ of effort, hence the clique is not a closed partnership. The worst blocking profile is the one where each clique has a single player with 1 effort, including the one to whom the 10th neighbor is now connected to. In this profile the farsighted player is no longer able to exploit the 10th player, but gains two players from his 9 other neighbors, who are now connected to




Figure 3.12: The network and the worst blocking effort profile of Example 3.6.11. The dotted ellipsis and lines indicates the neighborhood of the farsighted player.

trios of zeros. Hence the value of the lower bound increases to $f_1(2)$ again. This profile is shown on the right side of Figure 3.10.

As a consequence of Theorem 3.6.9, the star network with the farsighted player in the middle maximizes both the lower and the upper bound, and they are equal to $f_1(n-1)$ by Proposition 3.6.2. However, as shown by Example 3.6.10, deleting a link between two myopic players who are not neighbors of the farsighted player may decrease the lower bound. Furthermore, Example 3.6.11 shows that adding a link between two myopic players can dramatically increase the value of the lower bound. These two facts serve as potential stumbling blocks for a farsighted player who is trying to reach the star network. However, if linking to new players is a costless option, then the first statement of Theorem 3.6.9 guarantees that the farsighted player will link to all myopic players. Additionally, if the farsighted player can sever links running between myopic players, then the second statement implies that the star network will be reached by




Figure 3.13: The network and a worst effort profile of Example 3.6.11 with one link from one of the cliques removed. The number of neighbors that can be exploited increases to 2.

severing all links between his neighbors.

3.7 Conclusion

In this chapter we consider the private provision of local public goods game, a gameclass with a wide array of applicability and a strong body of theoretical contributions. We build a model of exploitation of a myopic population by a single farsighted player with the goal of analyzing exploitative behavior based on the spatial properties of the population. Our chapter is the first to consider a theoretical model combining the spatial asymmetries with various levels of player ability.

We consider a dynamic and stochastic setting such that in every time period a randomly selected player is given a chance to revise his exerted effort into the local public good. With this dynamic, the game is known to converge to a Nash effort profile.




Figure 3.14: The network and a worst effort profile of Example 3.6.11 with one link added. The number of neighbors that can be exploited increases to 2.

The farsighted player's optimal strategies correspond to optimal policies of a Markov Decision Process, and hence, stationary subgame perfect equilibria of this game exist for every network structure.

Our next question was the issue of convergence. We find that every network and every SSPE permits the existence of absorbing effort profiles that are stable for deviations given the strategies of the players. Moreover, we show that the game dynamics between myopic players propagate a rising inequality in the exerted efforts of the population. Unless stable partnerships are able to form, then myopic players with lower total efforts in their neighborhood will increase their efforts, while players with high total efforts in their neighborhood will lower theirs. As these two effects are mutually reinforcing each other, we observe a tendency towards a rising inequality of production of the local public good within the myopic population. This benefits the farsighted player if the high-effort players end up in his neighborhood and harms him otherwise.

As a result of the rising inequality, self-sustaining partnerships may form with high-effort players in the center and free-riders around them. These autarchic partnerships may become closed as the surrounding free-riders may dampen any ripple effect in the network (including the ones caused by the farsighted player) from reaching the central players with high efforts. In this chapter we show that no matter the network structure or the starting effort profile, due to the rise of inequality, an optimally playing farsighted player will cause the game to eventually settle in an absorbing state.

Finally, we analyze the farsighted player's minimum and maximum available long-run payoff in a given network as a way of evaluating and comparing different network structures. We provide characterizations for both bounds by the use of maximum independent sets of (some subset of) the farsighted player's neighborhood, and show that the value of the lower bound is reachable from any starting state by a very simple farsighted strategy. We show that adding a link that connects the farsighted player to a new neighbor or removing a link that runs between two of his neighbors improves both his minimum and maximum evaluation of the network. From this we conclude that the star network is the best structure for the farsighted player for a fixed set of players, but there are cases when adding a link between two myopic players, particularly, two myopic players outside the farsighted player's neighborhood, can increase his minimum and maximum evaluation.

3.A Toolkit: Buffers

This section introduces necessary concepts for the proofs of the statements in Section 3.5. Notice that the proof of Proposition 3.4.7 is achieved mostly via the application of Lemmas 3.4.2 and 3.4.4. In this section we extend Lemma 3.4.4 along the lines of the following example:

Example 3.A.1. Consider the game and the initial effort profile of Example 3.3.7. We have $x(\overline{N}_2) = x(\overline{N}_7) = 1 < x(\overline{N}_3) = 1.2$ (see Figure 3.1). Rising inequality between pairs applies for players 2 and 3 but not for players 2 and 7. However, after player 2 usurps player 3 (figure 3.15), player 7 becomes richer than player 2.

This is due to the fact that players 3 and 7 are not neighbors, hence player 7 gains from player 2's increase but does not lose from player 3's decrease. Hence, player 2 can usurp player 7 as well (figure 3.16).

Once player 2 usurps both of his neighbors, he becomes a one-player island (and faction). The set of free-riding players $\{3, 7\}$ acts as a buffer between the two islands,




Figure 3.15: The links between the myopic players and the effort profile of Example 3.A.1 after player 2 usurps player 3's efforts. Now rising inequality applies for players 2 and 7.




Figure 3.16: The links between the myopic players and the effort profile of Example 3.A.1 after player 2 usurps player 7's efforts as well. Player 2 becomes a one-player faction, and is separated from the island $\{4, 5, 6\}$ by a buffer of players, $\{3, 7\}$.

$\{2\}$ and $\{4, 5, 6\}$: as long as the farsighted player does not move away from his effort of 0.1, no matter how the efforts change in the latter, the former is not affected.

Naturally, once a second player has been usurped, the same argument can be repeated on higher levels, allowing for additional players with equal consumption to be taken over. We now introduce the idea formally.

Definition 3.A.2 (Buffer set). Let $x^0 \in X$ and $i \in I_{-1}$ be given. Let $D_i^0(x^0) = \{j \in N_i \cap I_{-1}^+(x^0) : x^0(\overline{N}_j) > x^0(\overline{N}_i)\}$. For $k \geq 1$ let $x^k \in X$ and $D_i^k(x^0) \subseteq N_i$ be defined recursively as follows:

$$x_j^k = \begin{cases} 0 & \text{if } j \in D_i^{k-1}(x^0), \\ \max\{0, 1 - x^{k-1}(N_i \setminus D_i^{k-1}(x^0))\} & \text{if } j = i, \\ x_j^{k-1} & \text{otherwise.} \end{cases}.$$

$$D_i^k(x^0) = D_i^{k-1}(x^0) \cup \{j \in N_i \cap I_{-1}^+(x^k) : x^k(\bar{N}_j) > x^k(\bar{N}_i)\} = D_i^{k-1}(x^0) \cup D_i^0(x^k).$$

Once $D_i^k(x^0) = D_i^{k-1}(x^0)$ for some $k \geq 1$, the set $D_i(x^0) = D_i^k(x^0)$ is called the *buffer set* of player i in effort profile x^0 .

The set $D_i^k(x^0)$ denotes the set of players whose efforts can be usurped by player i starting from effort profile x^0 , as per the conditions of Lemma 3.4.4. The effort profile x^{k+1} , is the one we get once the efforts of players in $D_i^k(x^0)$ have been usurped. By Lemma 3.4.4, for every $s \in \tilde{S}$ with $s_{-1} = b_{-1}$ we have $x^{k+1} \in R_s(x^k)$ with updates only from player i and elements of $D_i^k(x^0)$. Hence, $D_i(x^0)$ is the set of all players whose efforts can be usurped iteratively by player i , starting from effort profile x^0 . Notice that $D_i(x^0)$ may contain neighbors of i who do not have a higher consumption than i in the effort profile x^0 . Such is the case in Example 3.A.1, with player 7, who is a member of $D_2(x)$, but not of $D_2^0(x)$.

Lemma 3.A.3. *Let $x^0 \in X$ and $J \subseteq I_{-1}^+(x^0)$ be such that for every $i' \in N_J \cap I_{-1}^+(x^0)$ it holds that $x^0(\bar{N}_{i'}) > \min_{j' \in J} x^0(\bar{N}_{j'})$. For any $k \geq 0$, $i \in \operatorname{argmin}_{j' \in J} x^0(\bar{N}_{j'})$, and $j \in N_i \cap I_{-1}^+(x^0) \setminus D_i^k(x^0)$ it holds that*

$$j \in D_i^{k+1}(x^0) \setminus D_i^k(x^0) \Leftrightarrow D_i^k(x^0) \not\subseteq N_j.$$

Typically we take J to be an island, in which case the condition for all contributing myopic neighbors to consume more than the player with the lowest consumption is vacuous.

Proof. By Definition 3.A.2, $j \in (N_i \cap I_{-1}^+(x^0)) \setminus D_i^k(x^0)$ means $x^k(\bar{N}_j) \leq x^k(\bar{N}_i)$. By $i \in \operatorname{argmin}_{j' \in J} x^0(\bar{N}_{j'})$ together with the assumptions in Lemma 3.A.3 we have $x^k(\bar{N}_j) \geq x^k(\bar{N}_i)$, hence $x^k(\bar{N}_i) = x^k(\bar{N}_j)$. Furthermore, also by Definition 3.A.2, it holds that

$$x^{k+1}(\bar{N}_i) - x^k(\bar{N}_i) = x_i^{k+1} - x_i^k - x^k(D_i^k(x^0)).$$

\Rightarrow : Consider the contrapositive. If we had $D_i^k(x^0) \subseteq N_j$, then by $j \in N_i$ we have

$$x^{k+1}(\bar{N}_j) - x^k(\bar{N}_j) = x_i^{k+1} - x_i^k - x^k(D_i^k(x^0)),$$

meaning that $x^k(\bar{N}_j) = x^k(\bar{N}_i)$ leads to $x^{k+1}(\bar{N}_j) = x^{k+1}(\bar{N}_i)$, hence $j \notin D_i^{k+1}(x^0)$.

\Leftarrow : $D_i^k(x^0) \not\subseteq N_j$ means that

$$x^{k+1}(\overline{N}_j) - x^k(\overline{N}_j) = x_i^{k+1} - x_i^k - x^k(D_i^k(x^0) \cap N_j) > x_i^{k+1} - x_i^k - x^k(D_i^k(x^0)),$$

meaning that $x^k(\overline{N}_j) = x^k(\overline{N}_i)$ implies $x^{k+1}(\overline{N}_j) > x^{k+1}(\overline{N}_i)$, hence $j \in D_i^{k+1}(x^0)$. \blacksquare

Lemma 3.A.4. *Let $x \in X$ and a $J \subseteq I_{-1}^+(x)$ be given such that for every $i' \in N_J \cap I_{-1}^+(x)$ it holds that $x(\overline{N}_{i'}) > \min_{j' \in J} x(\overline{N}_{j'})$. Let $i, j \in \operatorname{argmin}_{j' \in J} x(\overline{N}_{j'})$ such that $g_{ij} = 1$. We have two statements:*

1. *if $j \notin D_i(x)$ and $i \notin D_j(x)$, then $D_i(x) = D_j(x)$.*
2. *if $j \notin D_i(x)$ and $i \in D_j(x)$, then $D_i(x) \subset D_j(x)$.*

Proof. We show that if $j \notin D_i(x)$, then $D_i(x) \subseteq D_j(x)$, which directly implies statement 2, and, by applying symmetry, we get statement 1.

By Lemma 3.A.3, if $j \notin D_i(x)$, then for every $k \geq 0$, $D_i^k(x) \subseteq N_j$, otherwise we would have $j \in D_i^{k+1}(x)$. Furthermore, it is also true that $D_i^0(x) \subseteq D_j^0(x)$ since for every $i' \in D_i^0(x)$ we have $x(\overline{N}_j) = x(\overline{N}_i) < x(\overline{N}_{i'})$, hence $i' \in D_j^0(x)$.

We show that, for any $k \geq 0$, $D_i^k(x) \subseteq D_j^k(x)$ implies $D_i^{k+1}(x) \subseteq D_j^{k+1}(x)$. Suppose the contrary, i.e. we have $D_i^k(x) \subseteq D_j^k(x)$ but $D_i^{k+1}(x) \not\subseteq D_j^{k+1}(x)$. Then there exists $j' \in D_i^{k+1}(x)$ such that $j' \notin D_j^{k+1}(x)$. Since $j' \notin D_i^k(x)$, we have $j' \notin D_i^0(x)$, hence it holds that $x(\overline{N}_{j'}) = x(\overline{N}_i)$. By $i \in J$ and by the assumption made on J we have $j' \in J$ and $j' \in \operatorname{argmin}_{i' \in J} x(\overline{N}_{i'})$. Since $D_i^{k+1}(x) \subseteq N_j$ we have $j' \in N_j$ as well. Hence, by Lemma 3.A.3, $j' \in D_i^{k+1}(x)$ implies $D_i^k(x) \not\subseteq N_{j'}$, while $j' \notin D_j^{k+1}(x)$ implies $D_j^k(x) \subseteq N_{j'}$, contradicting $D_i^k(x) \subseteq D_j^k(x)$. \blacksquare

Lemma 3.A.5. *Let $x \in X$ and a $J \subseteq I_{-1}^+(x)$ be given such that for every $i' \in N_J \cap I_{-1}^+(x)$ it holds that $x(\overline{N}_{i'}) > \min_{j' \in J} x(\overline{N}_{j'})$. For every $i^0 \in \operatorname{argmin}_{j' \in J} x(\overline{N}_{j'})$ there exists a non-empty connected $C \subseteq \operatorname{argmin}_{j' \in J} x(\overline{N}_{j'})$ such that*

1. *for every $i \in C$ it holds that $D = D_i(x) = N_C \cap I_{-1}^+(x)$,*
2. *$D_{i^0}(x) \subseteq D$.*

Furthermore, if there exists a $j \in N_{i^0} \cap I_{-1}^+(x)$ with $x(\overline{N}_{i^0}) < x(\overline{N}_j)$, then the resulting buffer, $D = N_C \cap I_{-1}^+(x)$ is non-empty.

Proof. The proof is constructive. Let $C^0 = \{i^0\}$ and let $D^0 = D_{i^0}(x)$. For $k \geq 0$ we construct a C^k and a D^k such that for every $i, i' \in C^k$ it holds that $D_i(x) = D_{i'}(x) =$

D^k , $x(\overline{N}_i) = x(\overline{N}_{i'}) = \min_{j' \in J} x(\overline{N}_{j'})$ and $D_{i^0}(x) \subseteq D^k$. Consider the following construction for C^{k+1} and D^{k+1} .

If we have $N_{C^k} \cap I_{-1}^+(x) = D^k$, then we terminate. If not, then there must exist a $j \in (N_{C^k} \cap I_{-1}^+(x)) \setminus D^k$. Notice that $j \notin D^k$ implies $x(\overline{N}_j) = \min_{j' \in J} x(\overline{N}_{j'})$, meaning that $j \in J$ and $j \in \operatorname{argmin}_{j' \in J} x(\overline{N}_{j'})$.

Take an arbitrary $j \in (N_{C^k} \cap I_{-1}^+(x)) \setminus D^k$ and an $i^k \in C^k \cap N_j$. If $i^k \notin D_j(x)$, then we take $C^{k+1} = C^k \cup \{j\}$. Then, by the first part of Lemma 3.A.4 we have $D_j(x) = D_{i^k}(x)$. Hence, if we take $D^{k+1} = D^k$, then if point 2 is satisfied for C^k and D^k , they are also satisfied for C^{k+1} and D^{k+1} . If $i^k \in D_j(x)$, then we take $C^{k+1} = \{j\}$. By the second part of Lemma 3.A.4 we have $D_j(x) \supset D_{i^k}(x)$. For $D^{k+1} = D_{i^k}(x)$, if point 2 is satisfied for C^k and D^k , it will also be satisfied for C^{k+1} and D^{k+1} .

In both cases, either C^k grows, or D^k does, while D^k can never decrease. Property 2 trivially holds for C^0 , and is preserved by the construction. The process will ultimately terminate with a $\kappa \leq n$ such that we have $D = N_{C^\kappa} \cap I_{-1}^+(x) = D^\kappa$, satisfying property 1 as well.

Clearly, if there exists a $j \in N_{i^0} \cap I_{-1}^+(x)$ such that $x(\overline{N}_{i^0}) < x(\overline{N}_j)$, then $j \in D_{i^0}^0$, meaning that D is non-empty. ■

The idea behind Lemma 3.A.5 is to take a player with a minimum consumption, find his neighbors, then take one of those players, and so on. Lemma 3.A.4 guarantees that we will arrive at a maximal set of players that can be usurped, and that this maximal set covers the entire non-free riding neighborhood.

Definition 3.A.6 (Partnership with buffer). Let $x \in X$ be given. The trio of sets (B, C, D) , with $B, C, D \subseteq I_{-1}$ is called a *partnership with a buffer* in x if the following properties are all satisfied:

- $B \in \mathcal{B}(x)$,
- $B \subseteq C$ and for every $i \in C$ it holds that $x_i = b_i(x)$,
- C is connected,
- $D = N_C \cap I_{-1}^+(x)$,
- for every $(i, j) \in C \times D$ it holds that $g_{ij} = 1$,
- $x(C \setminus B) = x(D) = 0$.

Note that Definition 3.A.6 only requires B and C to be non-empty. A partnership with a buffer consists of a partnership, B nested in a connected subset of myopic players, C which may contain free riders, and a buffer of free-riders, D separating the partnership from the remainder of the myopic player set.

In the following lemma we extend both Lemma 3.4.4 and Lemma 3.4.6.

Lemma 3.A.7. *Let $s \in \tilde{S}$ such that $s_{-1} = b_{-1}$. Let $x \in X$ be a recurrent effort profile such that $x_1 < 1$. Let $J \subseteq I_{-1}^+(x)$ be such that for every $i' \in N_J \cap I_{-1}^+(x)$ it holds that $x(\bar{N}_{i'}) > \min_{j' \in J} x(\bar{N}_{j'})$. For every $i^0 \in \operatorname{argmin}_{j' \in J} x(\bar{N}_{j'})$ there exists an effort profile $y \in R_s(x)$ and a partnership with a buffer (B, C, D) in y such that*

1. $C \subseteq \operatorname{argmin}_{i \in J} x(\bar{N}_i)$, $D_{i^0}(x) \subseteq D$,
2. y can be reached from x via updates only by the players in $C \cup D$.
3. $x \triangleleft_{(C \cup D) \cap N_1} y$.

Proof. By Lemma 3.A.5 there exists a connected $C \subseteq \operatorname{argmin}_{i \in J} x(\bar{N}_i)$ such that for every $i \in C$ we have $N_C \cap I_{-1}^+(x) = D_i(x)$ and $D_{i^0}(x) \subseteq D_i(x)$, satisfying properties 1 and 2.

Recall the definition of x^k from Definition 3.A.2. For every k it is true that $x^k \in R_s(x)$ with updates only by player i and the elements of $D_i^k(x)$. Let $\kappa \geq 0$ be given such that $x^\kappa = x^{\kappa-1}$, and let $y' = x^\kappa$. Thus, we get $y'(D) = 0$. Since $D = N_C \cap I_{-1}^+(x)$ and $y'(D) = 0$, C is an archipelago in y' . By applying Lemma 3.4.6 to each of its islands we reach a profile y with updates only by players belonging to C , such that for every $i \in C$ we have $y_i = b_i(y)$. Let $B = C \cap I_{-1}^+(y)$. Then, it is true that $B \in \mathcal{B}(y)$ and $y(C \setminus B) = 0$, satisfying all required properties of partnerships with buffers.

Now we show property 3. If $B \cap N_1 = \emptyset$, then $x = y$, otherwise B is a partnership in x but not in y , contradicting the recurrence of x by Proposition 3.4.7. If $x = y$, then property 3 is satisfied. Assume $B \cap N_1 \neq \emptyset$. Notice that since $y(C \setminus B) = y(D) = 0$, $M \in \mathcal{M}_y((C \cup D) \cap N_1)$ if and only if $M \in \mathcal{M}_y(B \cap N_1)$. Since x is recurrent, the set $\bar{N}_J \cap I_{-1}^+(x)$ has a pillar set in x , meaning that $\min_{j \in \bar{N}_J \cap I_{-1}^+(x)} x(\bar{N}_j \setminus \{1\}) \leq 1$. By the assumptions on J we must have $\operatorname{argmin}_{j \in \bar{N}_J \cap I_{-1}^+(x)} x(\bar{N}_j \setminus \{1\}) = \operatorname{argmin}_{j \in J} x(\bar{N}_j \setminus \{1\})$. Since for every $i \in B \cap N_1$ it holds that $i \in \operatorname{argmin}_{j \in J} x(\bar{N}_j)$, we also have $i \in \operatorname{argmin}_{j \in J} x(\bar{N}_j \setminus \{1\})$. Hence, for every $i \in B \cap N_1$ we have $x(\bar{N}_i \setminus \{1\}) \leq 1$, meaning that for every $M \in \mathcal{M}_y((C \cup D) \cap N_1)$ we have $M \in \mathcal{M}_x((C \cup D) \cap N_1)$, implying property 3. \blacksquare

Note that by Lemma 3.A.7, if there exists $j \in N_{i^0} \cap I_{-1}^+(x)$ such that $x(\bar{N}_{i^0}) < x(\bar{N}_j)$, then it also holds that $D_{i^0}(x) \subseteq D = N_C \cap I_{-1}^+(x) \neq \emptyset$.

This section's final lemma sets up the proofs of the lemmas in Section 3.5.

Lemma 3.A.8. *Fix $s \in \tilde{S}$ such that $s_{-1} = b_{-1}$. Let $x \in X$ be a recurrent effort profile in s such that $x_1 < 1$. There exists a $y \in X_+^* \cap R_s(x)$ that can be reached by updates only by myopic players such that $x \triangleleft_{N_1} y$ and $|\bar{L}(x)| \geq |\bar{L}(y)|$.*

Proof. First, let $I = \{1, 2\}$ with $g_{12} = 1$. For any initial profile x such that $x_1 < 1$, letting player 2 move to a best response will satisfy the statement for $y = (x_1, 1 - x_1)$, since $1 - x_1 > 0$, and $M = \{2\}$ is a pillar set of $I_{-1} = \{2\}$ in every profile. Clearly, $|\bar{L}(x)| = |\bar{L}(y)| = 0$.

We proceed by induction for the number of players, n . Assume that the statement is satisfied for every network of fewer than n players. If there exists $i \in I_{-1}^0(x)$, then, since myopic players behave identically in the presence of the free-rider i as they would in a game played by the set $I \setminus \{i\}$ on the network $G|_{I \setminus \{i\}}$, we can invoke the induction hypothesis for a game of $n - 1$ players which guarantees the statement. Hence we can assume that $I_{-1}^0(x) = \emptyset$. This means that the property $|\bar{L}(x)| \geq |\bar{L}(y)|$ is vacuous, since the number of contributors who are not in cliques cannot increase.

By Lemma 3.A.7 there exists profile $x' \in R_s(x)$ and a partnership with a buffer, B, C, D in x' such that $C \subseteq \operatorname{argmin}_{i \in I_{-1}} x(\bar{N}_i)$ and $x \triangleleft_{(C \cup D) \cap N_1} x'$. Since $D = N_C \setminus \{1\}$, we can remove $C \cup D$ and apply the induction condition on the set of players $I \setminus (C \cup D)$. Hence, there exists a profile $y \in X_+^* \cap R_s(x')$ for which $x' \triangleleft_{(I_{-1} \setminus (C \cup D)) \cap N_1} y$ and for every $i \in C \cup D$ it holds that $y_i = x'_i$. Any pillar sets of $(C \cup D) \cap N_1$ and of $(I_{-1} \setminus (C \cup D)) \cap N_1$ in y are disjoint due to $D = N_C \setminus \{1\}$ and $y(D) = 0$, hence $x \triangleleft_{N_1} y$ is implied by the Addition property of pillar sets. \blacksquare

3.B Proofs of the lemmas of Section 3.5

The proofs found in this section make extensive use of the concepts outlined in the previous one.

Lemma 3.5.3. *Fix $s \in \tilde{S}$ such that $s_{-1} = b_{-1}$. Let $x \in X$ be a recurrent effort profile in s such that $x_1 < 1$. Let $i \in I_{-1}$ be such that $x(\bar{N}_i) < 1$. There exists $y \in X_+^* \cap R_s(x)$ that can be reached by updates only by myopic players such that*

1. $y_i = b_i(y)$,
2. $x \triangleleft_{N_1} y$, and
3. $|\bar{L}(x)| \geq |\bar{L}(y)|$.

Proof. By Lemma 3.A.8 and by the transitivity of \triangleleft_{N_1} , we can take $x \in X_+^*$. If $x(\overline{N}_i) \geq 1$, then we are done. Otherwise, since $x \in X_+^*$, we must have $x_i = 0$. If $I_{-1}^+(x) \cap N_i = \emptyset$, then by letting player i move to his best response we get a profile that satisfies all properties. If not, then we must have $i \in N_1$, otherwise, by applying Lemma 3.4.4 on player i and the set $N_i \cap I_{-1}^+(x)$ we would get a partnership outside player 1's neighborhood, a contradiction by Proposition 3.4.7. Since, for every $F \in \mathcal{F}(x)$ it holds that $x(F) \geq 1 - x_1$, for any $F \in \mathcal{F}(x)$ it holds that $F \not\subseteq N_i$, otherwise $x(\overline{N}_i) \geq 1$. Hence, there must exist a player $i^0 \in I_{-1}^+(x) \setminus \overline{N}_i$ such that $(N_i \cap N_{i^0}) \cap I_{-1}^+(x) \neq \emptyset$.

Let x' be the profile we get by letting player i move to his best response. Then, for every $j \in N_i \cap I_{-1}^+(x)$ it holds that $x'(\overline{N}_j) > 1$, and hence, every myopic contributor in the neighborhood of $I_{-1}^+(x) \setminus \overline{N}_i$ has a higher consumption than the minimum consumption in $I_{-1}^+(x) \setminus \overline{N}_i$. Furthermore, since $i^0 \in \operatorname{argmin}_{i' \in I_{-1}^+(x) \setminus \overline{N}_i} (\overline{N}_{i'})$, by Lemma 3.A.7 we have a profile $x^1 \in R_s(x')$ and a partnership with a buffer, B^0, C^0, D^0 , in x^1 such that $C^0 \subseteq \operatorname{argmin}_{i' \in I_{-1}^+(x) \setminus \overline{N}_i} x(\overline{N}_{i'})$, $D_{i^0}(x') \subseteq D^0$, and $x' \triangleleft_{(C^0 \cup D^0) \cap N_1} x^1$. Notice that $x_i^1 = x_i' > 0$, since i 's move had to be an increase. Furthermore, $i \notin C^0 \cup D^0$ due to the fact that $C^0 \subseteq I_{-1}^+(x') \setminus \overline{N}_i$ and $D^0 \subseteq N_{C^0}$, meaning that $x' \triangleleft_{(C^0 \cup D^0) \cap N_1} x^1$ implies $x \triangleleft_{(C^0 \cup D^0) \cap N_1} x^1$. Finally we have $x^1(\overline{N}_i) < 1$, since $D_{i^0}^0(x')$ includes some of player i 's neighbors due to $(N_i \cap N_{i^0}) \cap I_{-1}^+(x) \neq \emptyset$, $D_{i^0}^0 \subseteq D^0$, and $x^1(D^0) = 0$.

Let $I^1 = (I_{-1}^+(x) \cup \{i\}) \setminus (C^0 \cup D^0) = I_{-1}^+(x^1) \setminus B^0$. For $k \geq 1$ we will construct the profile $x^{k+1} \in R_s(x^k)$ and a partnership with a buffer B^k, C^k, D^k in x^{k+1} , such that a number of properties are satisfied. Let $I^{k+1} = I^k \setminus (C^k \cup D^k)$. The properties we expect from x^{k+1} and B^k, C^k, D^k are as follows:

- $B^k, C^k, D^k \subseteq I^k$ and $D^k = N_{C^k} \cap I_{-1}^+(x^k) = N_{C^k} \cap I^k$.
- Only the players in $C^k \cup D^k$ update between x^k and x^{k+1} .
- If $i \in I^{k+1}$, then $x_i^{k+1} \geq x_i^k$ and $x^{k+1}(\overline{N}_i) < 1$.
- $x \triangleleft_{(C^k \cup D^k) \cap N_1} x^{k+1}$.

Let $\kappa = \min\{k \geq 1 : I^k = \emptyset\}$ and let $y = x^\kappa$. We will show that such a y satisfies all desired properties. For the construction of x^{k+1} we differentiate two cases.

Case 1: There exists a non-empty $J \subseteq I^k \setminus \{i\}$ such that for every $j \in I^k \cap N_J$ it holds that $x^k(\overline{N}_j) > \min_{j' \in J} x^k(\overline{N}_{j'})$. Then, by Lemma 3.A.7 there exists a profile $x^{k+1} \in R_s(x^k)$ and a partnership with a buffer, $B^k, C^k, D^k \subseteq I^k$ in x^{k+1} such that $C^k \subseteq \operatorname{argmin}_{j \in J} x^k(\overline{N}_j)$, $D^k = N_{C^k} \cap I_{-1}^+(x^k) = N_{C^k} \cap I^k$, and $x^k \triangleleft_{(C^k \cup D^k) \cap N_1} x^{k+1}$. Since $i \notin J$, it holds that $i \notin C^k$. If $i \notin D^k$ as well, then we have $x_i^{k+1} = x_i^k$, while

$x^k(\bar{N}_i) < 1$ implies $x^{k+1}(\bar{N}_i) < 1$ as, by the fact that only the players in $C^k \cup D^k$ have updated and $x^{k+1}(D^k) = 0$, no player in i 's neighborhood has a higher effort in x^{k+1} than in x^k . If $i \in D^k$ then $x_i^{k+1} = 0$ and $i \notin I^{k+1}$. Since for every $i' \in C^k \cup D^k$ we have $x_{i'} = x_{i'}^k$, $x^k \triangleleft_{(C^k \cup D^k) \cap N_1} x^{k+1}$ implies $x \triangleleft_{(C^k \cup D^k) \cap N_1} x^{k+1}$.

Case 2: For every $J \subseteq I^k \setminus \{i\}$ there exists a $j \in N_J \cap I^k$ such that $x^k(\bar{N}_j) \leq \min_{j' \in J} x^k(\bar{N}_{j'})$. This property implies that I^k is connected, as any component of I^k that does not include i would satisfy the previous property. It also means that $i \in \operatorname{argmin}_{j \in I^k} x^k(\bar{N}_j)$, otherwise we have a contradiction for $J = I^k \setminus \{i\}$.

If $I^k = \{i\}$, then let i move to his best response and let x^{k+1} denote the resulting effort profile. Since $x^k(\bar{N}_i) < 1$, we have $x_i^{k+1} > x_i^k$. Furthermore, $\{i\} \in \mathcal{B}(x^{k+1})$. Let $C^k = \{i\}$ and $D^k = \emptyset$. We clearly have $x \triangleleft_{(C^k \cup D^k) \cap N_1} x^{k+1}$.

If $I^k \subseteq \bar{N}_i$, but $I^k \neq \{i\}$, then I^k must be a clique due to $i \in \operatorname{argmin}_{j \in I^k} x^k(\bar{N}_j)$. Therefore, for every $j \in I^k$ we have $x^k(\bar{N}_j) = x^k(\bar{N}_i) < 1$. Let any player from I^k move to best response and let x^{k+1} denote the resulting effort profile. Let $B^k = C^k = I^k$ and $D^k = \emptyset$. Then, every pillar set of $C^k \cup D^k$ in x^{k+1} is a singleton, containing one player of I^k . Since $i \in N_1$ we have $\min_{M \in \mathcal{M}_{x^{k+1}}((C^k \cup D^k) \cap N_1)} |M| = 1$, while $\min_{M \in \mathcal{M}_x((C^k \cup D^k) \cap N_1)} \leq 1$, we also have $x \triangleleft_{(C^k \cup D^k) \cap N_1} x^{k+1}$.

If I^k has members outside of \bar{N}_i , then we first show that $(\operatorname{argmin}_{j \in I^k \setminus \bar{N}_i} x^k(\bar{N}_j)) \cap N_{\bar{N}_i} \neq \emptyset$. Suppose the contrary. Then, $\min_{j \in I^k \setminus \bar{N}_i} x^k(\bar{N}_j) < \min_{j' \in I^k \cap N_{\bar{N}_i}} x^k(\bar{N}_{j'})$. This means that $I^k \setminus \bar{N}_{\bar{N}_i} \neq \emptyset$. Let $J = I^k \setminus \bar{N}_{\bar{N}_i}$. Then, $N_J \cap I^k = N_{\bar{N}_i} \cap I^k$, and therefore, for every $j \in N_J \cap I^k$ it holds that $x^k(\bar{N}_j) > \min_{j' \in J} x^k(\bar{N}_{j'})$, contradicting the condition of reaching Case 2.

Case 2a: There exists an $i' \in (\operatorname{argmin}_{j \in I^k \setminus \bar{N}_i} x^k(\bar{N}_j)) \cap N_{\bar{N}_i}$ and a $j' \in (N_i \cap N_{i'}) \cap I^k$ such that $x^k(\bar{N}_i) = x^k(\bar{N}_{j'}) < x^k(\bar{N}_{i'})$. Then $j' \in \operatorname{argmin}_{j \in I^k} x^k(\bar{N}_j)$. By Lemma 3.A.7 there exists a profile $x^{k+1} \in R_s(x^k)$ and a partnership with a buffer B^k, C^k, D^k such that $C^k \subseteq \operatorname{argmin}_{j \in J} x^k(\bar{N}_j)$, $D^k = N_{C^k} \cap I^k = N_{C^k} \cap I_{-1}^+(x^k)$, $D_{j'}(x^k) \subseteq D^k$, and $x^k \triangleleft_{(B^k \cup D^k) \cap N_1} x^{k+1}$. From this point, the arguments follow Case 1, except in this case we have $i \in D^k$. This is due to the fact that $i' \in D^k$ by $i' \in D_{j'}^0(x^k)$. By $i' \notin N_i$, Lemma 3.A.3 gives $i \in D_{i'}^1(x^k)$, implying $i \in D^k$.

Case 2b: The condition of reaching Case 2a does not hold, however, there exists a player $i' \in (\operatorname{argmin}_{j \in I^k \setminus \bar{N}_i} x^k(\bar{N}_j)) \cap N_{\bar{N}_i}$ and a $j' \in I^k \cap (N_i \cap N_{i'})$ such that $x^k(\bar{N}_i) = x^k(\bar{N}_{j'}) = x^k(\bar{N}_{i'})$. Then, $i' \in \operatorname{argmin}_{j \in I^k} x^k(\bar{N}_j)$. Let player i move to his best response from x^k and call the resulting effort profile \hat{x}^k . Due to $x^k(\bar{N}_i) < 1$ this is an increase, hence $\hat{x}_i^k > x_i^k$. We get $\hat{x}^k(\bar{N}_i) = \hat{x}^k(\bar{N}_{j'}) = 1$, and $i' \in \operatorname{argmin}_{j \in I^k} \hat{x}^k(\bar{N}_j)$. By Lemma 3.A.7 there exists $x^{k+1} \in R_s(\hat{x}^k)$ and a partnership with a buffer B^k, C^k, D^k in x^{k+1} such that $C^k \subseteq \operatorname{argmin}_{j \in I^k} \hat{x}^k(\bar{N}_j)$, $D^k = N_{C^k} \cap I^k = N_{C^k} \cap I_{-1}^+(x^k)$, $D_{i'}(\hat{x}^k) \subseteq$

D^k , and $x^k \triangleleft_{(C^k \cup D^k) \cap N_1}$. Clearly, $i \notin C^k$, hence we either have $i \in D^k$ and $i \notin I^{k+1}$, or $x_i^{k+1} = \hat{x}_i^k > x_i^k$. Furthermore, notice that $j' \in D_{i'}(\hat{x}^k)$, hence $j' \in D^k$. Therefore, if $i \in I^{k+1}$, then we get $x^{k+1}(\bar{N}_i) < 1$ as $D^k \cap N_i \neq \emptyset$.

Case 2c: The conditions of reaching Cases 2a and 2b do not hold. If we had a player $i' \in \operatorname{argmin}_{j \in I^k \setminus \bar{N}_i} x^k(\bar{N}_j) \cap N_{\bar{N}_i}$ and a $j' \in I^k \cap (N_i \cap N_{i'})$ such that $x^k(\bar{N}_i) = x^k(\bar{N}_{j'}) > x^k(\bar{N}_{i'})$, it would contradict $i \in \operatorname{argmin}_{j \in I^k} x^k(\bar{N}_j)$. Hence, if we select an arbitrary $i' \in (\operatorname{argmin}_{j \in I^k \setminus \bar{N}_i} x^k(\bar{N}_j)) \cap N_{\bar{N}_i}$ and a $j' \in I^k \cap (N_i \cap N_{i'})$, we must have $x^k(\bar{N}_i) < x^k(\bar{N}_{j'})$. As in Case 2b, let \hat{x}^k denote the profile we get by letting player i move to his best response. Since $x^k(\bar{N}_i) < 1$, we have $\hat{x}_i^k > x_i^k$, and due to $i \in \operatorname{argmin}_{j \in I^k} x(\bar{N}_j)$, for every $j \in N_i \cap I^k$ we have $1 = \hat{x}^k(\bar{N}_i) \leq \hat{x}^k(\bar{N}_j)$. Since $x \in X_+^*$, and since no player in the neighborhood of i' has increased, we have $x^k(\bar{N}_{i'}) \leq 1$. Hence, $i' \in \operatorname{argmin}_{j \in I^k} \hat{x}^k(\bar{N}_j)$. Therefore, by Lemma 3.A.7 we have a profile $x^{k+1} \in R_s(x^k)$ and a partnership with a buffer B^k, C^k, D^k such that $C^k \subseteq \operatorname{argmin}_{j \in I^k} \hat{x}^k(\bar{N}_j)$, $D_{i'}(\hat{x}^k) \subseteq D^k$, and $x \triangleleft_{(C^k \cup D^k) \cap N_1} x^{k+1}$. We now show that $i \notin C^k$. Suppose the contrary. If we had $i \in C^k$, then $\min_{j \in I^k} \hat{x}^k(\bar{N}_j) = 1$. By the connectedness of C^k this would mean that there is a pair $i^0 \in I^k \setminus \bar{N}_i$ and $j^0 \in I^k \cap N_i$ with $g_{i^0 j^0} = 1$ and $\hat{x}^k(\bar{N}_{j^0}) = \hat{x}^k(\bar{N}_{i^0}) = 1$. However, this must mean $x^k(\bar{N}_i) = x^k(\bar{N}_{j^0}) < x^k(\bar{N}_{i^0})$, contradicting the fact that Case 2b's conditions are not satisfied. Hence, $i \notin C^k$. The fact that $x_i^{k+1} > x_i^k$ if $i \in I^{k+1}$ and $x_i^{k+1} = b_i(x^k) = 0$ otherwise follows similarly to Case 1.

We now show that each of the three properties are satisfied for y , starting with property 2. Notice that since C^k is never empty when I^k is non-empty, the process described above will terminate in x^κ with $\kappa < n$. Consider the sets $C^k \cup D^k$. Clearly, we have $I_{-1}^+(x) \cup \{i\} = \bigcup_{k=0}^\kappa (C^k \cup D^k)$. Since for every $k \geq 0$ we have $x \triangleleft_{(C^k \cup D^k) \cap N_1} x^{k+1}$, B^k, C^k, D^k are partnerships with buffers in y , covering the entirety of $I_{-1}^+(x) \cup \{i\}$, and since the players in $C^k \cup D^k$ do not change their efforts after reaching the profile x^{k+1} , the Addition property of pillar sets implies $x \triangleleft_{N_1} y$ for $y = x^\kappa$, satisfying property 2. Since every myopic contributor in y is at best response due to the properties of partnerships and buffers, we have $y \in X_+^*$.

Now we show $y_i = b_i(y)$. Clearly, $i \in C^k \cup D^k$ for some $k \leq \kappa$. If $i \in C^k$ for some $k \leq \kappa$, then $x_i^{k+1} = b_i(x^{k+1})$ by property 2 of Definition 3.A.6. Since $D^k = N_{C^k} \cap I^k$, and since $x^{k+1}((N_{C^k} \setminus D^k) \cap I_{-1}) = 0$ by construction, we have $x^{k+1}(N_{C^k} \setminus \{1\}) = 0$. The construction of y guarantees $y(N_{C^k} \setminus \{1\}) = 0$ as well, meaning that $x_i^{k+1} = b_i(x^{k+1})$ implies $y_i = b_i(y)$. If $i \in D^k$ for some $k \leq \kappa$, then $y_i = 0$. Furthermore, since $B^k \subseteq N_i$, and since $y(B^k) \geq 1 - x_1$, by $i \in N_1$, we must also have $y(\bar{N}_i) \geq 1$, meaning that $y_i = b_i(y)$.

Finally, we show $|\bar{L}(y)| \leq |\bar{L}(x)|$. If $x_i > 0$, then by the construction of y it holds that $I_{-1}^+(y) \subseteq I_{-1}^+(x)$. Therefore, no island that was not a clique could become larger, and no island that was a clique could become a non-clique, implying $|\bar{L}(y)| \leq |\bar{L}(x)|$. The same holds if $x_i = 0$ and $y_i = 0$, as, once again, no free rider in x is a contributor in y . If $x_i = 0$ but $y_i > 0$, then we must have had a $k \leq \kappa$ with $i \in B^k$. As shown in Case 2, this can only happen if i itself is a member of a clique, hence y satisfies property 3 as well. \blacksquare

Lemma 3.5.4. Fix $s \in \tilde{S}$ such that $s_{-1} = b_{-1}$. If there exists an $x \in X$ with $|[x]_s| > 1$, and $x_1 < 1$, then there exists an action profile $y \in X^* \cap R_s(x)$, that can be reached via updates only by myopic players, such that $x \triangleleft_{N_1} y$ and $|\bar{L}(x)| \geq |\bar{L}(y)|$.

Proof. First, let $I = \{1, 2\}$ with $g_{12} = 1$. For any initial profile x such that $x_1 < 1$, letting player 2 move to a best response will satisfy the statement for $y = (x_1, 1 - x_1)$, since $1 - x_1 > 0$, and $M = \{2\}$ is a pillar set of $I_{-1} = \{2\}$ in every profile. Clearly, $|\bar{L}(x)| = |\bar{L}(y)| = 0$.

We proceed by induction for the number of players, n . Assume that the statement is satisfied for every network of fewer than n players.

Let $x^1 = x$. For $k \geq 1$ we will construct the profile $x^{k+1} \in R_s(x^k)$ such that the following properties are satisfied:

1. Only myopic players update between x^k and x^{k+1} .
2. $x^{k+1} \in X_+^*$, and if $x^k \notin X^*$, then $x^{k+1} \neq x^k$.
3. $x^k \triangleleft_{N_1} x^{k+1}$.
4. $|\bar{L}(x^k)| \geq |\bar{L}(x^{k+1})|$.

Let $\kappa = \min\{k \geq 1: x^k = x^{k+1}\}$. Such a κ must exist due to the fact that $X_1 \times X'_{-1}$ is finite, we cannot leave it by Lemma 3.3.8 and Remark 3.3.9, and that for every $1 \leq k \leq \kappa$ it holds that $\phi(x^k) < \phi(x^{k+1})$. Clearly, by the transitivity of \triangleleft_{N_1} , the profile $y = x^\kappa$ satisfies all of the desired properties. The construction of x^{k+1} , given x^k proceeds as follows. We distinguish four cases.

Case 1: There exists $i \in I_{-1}$ with $x^k(\bar{N}_i) < 1$. In this case Lemma 3.5.3 guarantees the existence of x^{k+1} which satisfies properties 1, 3, and 4. Furthermore, x^{k+1} is such that $x^{k+1} \in X_+^*$ and $x_i^{k+1} = b_i(x^{k+1})$, meaning that $x^{k+1} \neq x^k$, satisfying property 2 as well.

Case 2: For every $i \in I_{-1}$ we have $x^k(\bar{N}_i) \geq 1$, and $I_{-1}^0(x^k) \neq \emptyset$. In this case we have an $i \in I_{-1}$ such that $x_i^k = 0$. Since myopic players behave identically in the presence of the free-rider i as they would in a game played by the set $I \setminus \{i\}$ on the network $G|_{I \setminus \{i\}}$, we can invoke the induction hypothesis for a network with $n - 1$ players. Therefore, there exists $x^{k+1} \in R_s(x^k)$ which satisfies property 1, as well as $x^{k+1} \in X_+^*$, because all myopic players other than i are at their best response. We have $x^k \triangleleft_{N_1 \setminus \{i\}} x^{k+1}$ by the induction hypothesis, which, due to $x_i^k = 0$ implies property 3. Since $x_i^k = x_i^{k+1}$, we have $|\bar{L}(x^k)| \geq |\bar{L}(x^{k+1})|$ as well, satisfying property 4. If for every $j \in I_{-1} \setminus \{i\}$ we have $x_j^{k+1} = x_j^k$, then we must have had $x^k \in X^*$ due to $x^k(\bar{N}_i) \geq 1$, satisfying property 2 as well.

Case 3: For every $i \in I_{-1}$ we have $x^k(\bar{N}_i) \geq 1$, $I_{-1}^0(x^k) = \emptyset$, and we have two players, $i, j \in I_{-1}$ with $g_{ij} = 1$ and $x^k(\bar{N}_i) \neq x^k(\bar{N}_j)$. In this case, property 4 is vacuous, as $|\bar{L}(x^k)|$ cannot increase. By Lemma 3.A.7, there exists a profile \hat{x}^k and a partnership with buffer B, C, D in \hat{x}^k with $C \subseteq \operatorname{argmin}_{i \in I_{-1}} x^k(\bar{N}_i)$, $D = N_C \setminus \{1\}$, and $x^k \triangleleft_{(C \cup D) \cap N_1} \hat{x}^k$.

Since every element of $I_{-1} \setminus (C \cup D)$ is separated from B by the buffer of free-riders D in the profile \hat{x}^k , by $s_{-1} = b_{-1}$, every myopic player will behave identically as they would in a network without $C \cup D$, $G|_{I \setminus (C \cup D)}$. Therefore, we can apply the induction hypothesis on the set of players $I \setminus (C \cup D)$. Hence, there exists a profile x^{k+1} such that for every $j \in C \cup D$ it holds that $x_j^{k+1} = \hat{x}_j^k$, every member of $I_{-1} \setminus (C \cup D)$ is at his best response in x^{k+1} , and $\hat{x}^k \triangleleft_{N_1 \setminus (C \cup D)} x^{k+1}$. Clearly, x^{k+1} satisfies properties 1, 2, and 4. Since $\hat{x}^k \triangleleft_{N_1 \setminus (C \cup D)} x^{k+1}$ implies $x^k \triangleleft_{N_1 \setminus (C \cup D)} x^{k+1}$, and since $x^k \triangleleft_{(C \cup D) \cap N_1} \hat{x}^k$ implies $x^k \triangleleft_{(C \cup D) \cap N_1} x^{k+1}$, the Addition property of pillar sets gives $x^k \triangleleft_{N_1} x^{k+1}$, satisfying property 3.

Case 4: For every $i \in I_{-1}$ we have $x^k(\bar{N}_i) \geq 1$, $I_{-1}^0(x^k) = \emptyset$, and we for every pair of players $i, j \in I_{-1}$ with $g_{ij} = 1$ it holds that $x^k(\bar{N}_i) = x^k(\bar{N}_j)$. If G is a disconnected network, we apply the induction hypothesis to each of its components. If G is connected, then for every $i \in I_{-1}$ it holds that $x^k(\bar{N}_i) = \alpha \geq 1$. If $\alpha = 1$, then $x^k \in X^*$ and we are done. If $\alpha > 1$ and I_{-1} is a clique, then by letting every myopic player move to his best response in any order we reach a suitable x^{k+1} .

If $\alpha > 1$ and I_{-1} is not a clique, then there must exist $i, j, j' \in I_{-1}$ such that $g_{ij} = g_{ij'} = 1$ but $g_{jj'} = 0$. Let the players j and j' move to their best responses from x^k and call the resulting effort profile x' . Since $\alpha > 1$, both players must decrease, meaning that we have $x'(\bar{N}_j) = x'(\bar{N}_{j'}) > x'(\bar{N}_i)$, and $i \in \operatorname{argmin}_{i' \in I_{-1}} x'(\bar{N}_{i'})$. By Lemma 3.A.7 there exists an effort profile \hat{x}^k and a partnership with buffer B, C, D in \hat{x}^k such that $C \subseteq \operatorname{argmin}_{i' \in I_{-1}} x'(\bar{N}_{i'})$, $D = N_C \setminus \{1\}$, $D_i(x') \subseteq D$, and $x^k \triangleleft_{(C \cup D) \cap N_1} \hat{x}^k$.

Notice that $j, j' \in D$, since $j, j' \in D_i^0(x')$. From here the analysis of this case follows that of Case 3. \blacksquare

Lemma 3.5.6. Fix $s \in \tilde{S}$ such that $s_{-1} = b_{-1}$. If there exists an $x \in X$ with $||x|_s| > 1$ and $x_1 < 1$, then there exists an action profile $y \in [x]_s \cap X^*$, such that $x \triangleleft_{N_1} y$ and for every $F \in \mathcal{F}(y)$ it holds that either $F \subseteq N_1$ or $F \cap N_1 = \emptyset$, and $|\overline{L}(x)| \geq |\overline{L}(y)|$.

Proof. By Lemma 3.5.4 and by the transitivity of \triangleleft_{N_1} , we can take a $x' \in R_s(x) \cap X^*$ and show that a $y \in R_s(x')$ exists that satisfies the desired properties. If the statement does not hold for x' , then we must have $F \in \mathcal{F}(x')$ with an $i \in F \cap N_1$ and a $j \in F \setminus N_1$ such that $g_{ij} = 1$.

Because of $||x|_s| > 1$, we must have either $s_1(x') > x'_1$ or $s_1(x') < x'_1$. If $s_1(x') < x'_1$, then by Lemma 3.5.4 there exists x^1 such that $x' \triangleleft_{N_1} x^1$, $|\overline{L}(x')| \geq |\overline{L}(x^1)|$, and $x^1 \in X^*$. If no faction in x^1 has members in both N_1 and $I_{-1} \setminus N_1$, we are done. If there are such factions, then again we must have $s_1(x^1) \neq x^1_1$. If $s_1(x^1) < x^1_1$, we repeat the same argument until we find a profile of similar properties for which either the statement holds, or where player 1 increases. Since negative contributions are not possible, player 1 cannot decrease forever.

If $s_1(x') > x'_1$, then let player 1 revise, and let the resulting profile be denoted by x'' . For every $i' \in F \cap N_1$ we have $x''(\overline{N}_{i'}) > x''(\overline{N}_j) = 1$, and we also have $j \in \operatorname{argmin}_{j' \in F} x''(\overline{N}_{j'})$. Hence, by Lemma 3.A.7 there exists a profile $\hat{x} \in R_s(x'')$ and a partnership with a buffer B, C, D in \hat{x} such that $C \subseteq \operatorname{argmin}_{j' \in F} x''(\overline{N}_{j'})$, and $D = N_C \cap I_{-1}^+(x'') = N_C \cap F$. Hence, $B \in \mathcal{B}(\hat{x})$ with $B \cap N_1 = \emptyset$, a contradiction by Proposition 3.4.7. Therefore, every faction of x' is either completely inside or completely outside N_1 . \blacksquare

Lemma 3.5.8. Fix $s \in \tilde{S}$ such that $s_{-1} = b_{-1}$ and for every $x \in X$ it holds that $s_1(x) \leq 1$. If there exists an $x \in X$ with $||x|_s| > 1$ and $x_1 < 1$, then there exists an action profile $y \in [x]_s \cap X^*$, such that $x \triangleleft_{N_1} y$, for every $F \in \mathcal{F}(y)$ we either have $F \subseteq N_1$ or $F \cap N_1 = \emptyset$, and for every $F \in \mathcal{F}_1(y)$ it holds that F is a clique.

Proof. By Lemmas 3.5.4 and 3.5.6, and by the transitivity of \triangleleft_{N_1} we can take $x \in X^*$, such that for every $F \in \mathcal{F}(x)$ we either have $F \subseteq N_1$ or $F \cap N_1 = \emptyset$. We prove that for every such x , if there is a $F \in \mathcal{F}_1(x)$ that is not a clique, then we have a $y \in R_s(x) \cap X^*$ with $x \triangleleft_{N_1} y$ and $|\overline{L}(y)| < |\overline{L}(x)|$.

Since there exists a $F \in \mathcal{F}_1(x)$ that is not a clique, there exist players $i, j, j' \in F$ with $g_{ij} = g_{ij'} = 1$ but $g_{jj'} = 0$. Consider the profile $x' = (s_1(x), x_{-1})$. Clearly, $x' \neq x$, otherwise we have $||x|_s| = 1$.

If $s_1(x) < x_1$, then by Lemma 3.5.4 there exists a $x'' \in R_s(x') \cap X^*$ such that $x' \triangleleft_{N_1} x''$ and $\bar{L}(x'') \leq \bar{L}(x')$. Since $x'(\bar{N}_i) < 1$ by player 1's decrease, we must have $x'' \neq x'$ and $x''_1 = x'_1 < x_1$. We then consider the profile $(s_1(x''), x''_{-1})$, and if player 1 decreases again, we repeat the above argument. Since player 1 cannot decrease forever, as in Lemma 3.5.6, we must have an effort profile $\tilde{x} \in R_s(x) \cap X^*$ with $x \triangleleft_{N_1} \tilde{x}$ and $\bar{L}(\tilde{x}) \leq \bar{L}(x)$ such that $s_1(\tilde{x}) > \tilde{x}_1$. We therefore assume $s_1(x) > x_1$.

We discuss two cases.

Case 1: $s_1(x) \in (x_1, 1)$. Then, it holds that $x'(\bar{N}_i) = x'(\bar{N}_j) = x'(\bar{N}_{j'}) > 1$. Let x'' denote the profile we get by letting j and j' move to a best response, both of whom decrease. Now we have $x''(\bar{N}_i) < x''(\bar{N}_j) = x''(\bar{N}_{j'})$, and $i \in \operatorname{argmin}_{i' \in F} x''(\bar{N}_{i'})$. Take $J = F \cap I_{-1}^+(x'')$. By Lemma 3.A.7 there exists an effort profile $\hat{x} \in R_s(x'')$ and a partnership with a buffer B, C, D in \hat{x} such that $C \subseteq \operatorname{argmin}_{i' \in F} x''(\bar{N}_{i'})$, $D = N_C \cap F$, $D_i(\hat{x}) \subseteq D$, and $x'' \triangleleft_{(C \cup D) \cap N_1} \hat{x}$.

Notice that $x''_j > 0$ and $x''_{j'} > 0$ imply $j \in D$ and $j' \in D$ respectively, as in this case they are members of J , and of $D_i^0(x'')$, hence $\hat{x}_j = \hat{x}_{j'} = 0$. Since the players in $I_{-1} \setminus (C \cup D)$ did not change their efforts between x and \hat{x} we clearly have $x \triangleleft_{N_1 \setminus (C \cup D)} \hat{x}$, meaning that, by the Addition property of pillar sets, we have $x \triangleleft_{N_1} \hat{x}$ as well. Since no player who was a free rider in x has positive efforts in \hat{x} , but we had $j, j' \in \bar{L}(x)$ who are both free riders in \hat{x} , we must have $\bar{L}(\hat{x}) < \bar{L}(x)$.

By Lemma 3.5.6, and by the transitivity of \triangleleft_{N_1} , there exists a $y \in R_s(\hat{x})$ which returns all myopic players to playing best responses, satisfying all desired properties.

Case 2: $s_1(x) = 1$. In this case, let x'' denote the profile we get by letting all of N_1 move to best response, leading to $x''(N_1) = 0$. Then, for every $i \in I \setminus \bar{N}_1$ we have $x''_i = b_i(x'')$, as every contributor is a member of a partnership, and should a free rider $j \in I \setminus \bar{N}_1$ have a consumption lower than 1, by Lemma 3.4.4 we would have a $\tilde{x} \in R_s(x'')$ where $\{j\} \in \mathcal{F}(\tilde{x})$, a contradiction by Proposition 3.4.7. Since every myopic player is at best response in x'' and we have $s_1(x'') \leq 1$ by assumption, we must have $s_1(x'') < 1$, otherwise $|[x'']_s| = 1$, a contradiction.

Let player 1 move from x'' and call the resulting effort profile \hat{x} . Since for every $j \in I \setminus \bar{N}_1$ we had $x_j = b_j(x)$ while $\hat{x}_j = b_j(\hat{x})$ holds as well, for every $i \in \bigcup_{F \in \mathcal{F}_1(x)} F = N_1 \cap I_{-1}^+(x)$ it must hold that $\hat{x}(\bar{N}_i) = \hat{x}_1$. Let M be a maximum independent set of $\bigcup_{F \in \mathcal{F}_1(x)} F$. Notice that $M \in \mathcal{M}_x(N_1)$. Let the members of M update in any order and call the resulting effort profile \hat{y} . Then, for every $i \in M$ we will have $\hat{y}_i = 1 - \hat{x}_1$ and for every $j \in I \setminus M$ we have $\hat{y}_j = \hat{x}_j$. Since the only pillar set of N_1 in \hat{y} is M , we have $x \triangleleft_{N_1} \hat{y}$. Furthermore, as there was a faction inside N_1 in profile x that was not a clique, whereas there are no such factions in \hat{y} , we must have $\bar{L}(\hat{y}) < \bar{L}(x)$.

If $\hat{y} \in X^*$, then we are done. Otherwise, Lemma 3.5.6 guarantees the existence of $y \in R_s(\hat{y})$ which satisfies all desired properties. \blacksquare

3.C Proofs of Propositions 3.6.2 and 3.6.8

Proposition 3.6.2. There exists $\bar{\delta}$ such that for every $G \in \mathcal{G}_1$ and every $\delta > \bar{\delta}$ it holds that

$$\underline{u}(G) = \bar{u}(G) = f_1\left(\max_{M \in \mathcal{M}_1^{nd}(G)} |M|\right).$$

Proof. For $x \in X$ let $I_{-1}^1(x) = \{i \in I_{-1} : x_i = 1\}$ denote the set of myopic players with effort level 1. Let M be a maximum independent set of $N_1(G)$ and consider the following strategy, $s_1 \in \tilde{S}_1$:

$$s_1(x) = \begin{cases} 0 & \text{if } I_{-1}^1(x) \subseteq M, \\ 1 & \text{otherwise.} \end{cases}$$

To show the statement, we first argue that by the strategy profile $s = (s_1, b_{-1})$ the game reaches an effort profile with probability 1 where the farsighted player's instantaneous payoff equals $f_1(|M|)$. Let $x^M \in X$ be the effort profile for which $x_i = 1$ if $i \in M$ and $x_i = 0$ otherwise. Notice that $A(s) = \{x^M\}$, as in every other effort profile the farsighted player or a member of M would change his effort. Furthermore, for every $x \in X$ it holds that $x^M \in R_s(x)$, and it is reachable by fewer than $n + |M|$ revisions as follows: If $I_{-1}^1(x) \not\subseteq M$, then player 1 revises to play 1, then each of the myopic players revise in any order to play 0. This takes n steps. Once all myopic players are at 0, player 1 decreases his efforts to 0, and then the members of M revise in any order all of them increasing to 1, reaching the profile x^M in $|M|$ more steps. If $I_{-1}^1(x) \subseteq M$, then all members of $M \setminus I_{-1}^1(x)$ revise any order, all of whom increasing to 1, once again, reaching the profile x^M in at most $|M|$ steps.

With x^M being the unique absorbing profile, and it being reachable in a bounded number of steps for every $\omega \in \Omega$ we have that

$$\lim_{t \rightarrow \infty} u_1^t(\omega, s) = f_1(|M|).$$

Since there exists a strategy that guarantees a long-run payoff of $f_1(|M|)$ there must be a sufficiently high $\bar{\delta}$ such that for every $\delta > \bar{\delta}$ and every SSPE s^* we have

$$\lim_{t \rightarrow \infty} u_1^t(\omega, s^*) \geq \lim_{t \rightarrow \infty} u_1^t(\omega, s) = f_1(|M|).$$

On the other hand, $f_1(|M|) = \max_{x \in X'} \pi_1(x)$, so by Lemma 3.3.8 and Remark 3.3.8 we also have

$$\lim_{t \rightarrow \infty} u_1^t(\omega, s^*) \leq f_1(|M|).$$

Hence, $\underline{u}(G) = \bar{u}(G) = f_1(|M|)$. ■

Proposition 3.6.8. For every $G \in \mathcal{G}$ it holds that

$$\underline{u}(G) = \begin{cases} f_1(\underline{e}(G)) & \text{if } \underline{e}(G) > 0, \\ f_1(1) - c_1 & \text{if } \underline{e}(G) = 0. \end{cases}$$

Proof. We first consider the case that $\underline{e}(G) = 0$. Clearly, in any network the farsighted player can guarantee the long-run payoff $f_1(1) - c_1$ by playing the strategy $s_1(x) = 1$ for every $x \in X$. Hence we must have $\underline{u}(G) \geq f_1(1) - c_1$. If $\underline{e}(G) = 0$, then there must exist an $\tilde{x} \in \tilde{X}$ such that $\max_{M \in \mathcal{M}_1^{\text{ind}}(G|_{I(\tilde{x})})} |M| = 0$, which happens if and only if $N_1 \cap I(\tilde{x}) = \emptyset$. Take $\omega = (\tilde{x}, 1)$. In this game, player 1 plays as if he was the only player as he gets 0 from his neighbors and no player will change his effort level no matter what player 1 does. Hence, in every SSPE we must have $s_1^*(\tilde{x}) = 1$, meaning that $\underline{u}(G) \leq f_1(1) - c_1$ also holds.

We now consider the case $\underline{e}(G) > 0$. We first show that $\underline{u}(G) \leq f_1(\underline{e}(G))$. Let $\underline{x} \in \operatorname{argmin}_{\tilde{x} \in \tilde{X}} \max_{M \in \mathcal{M}_1^{\text{ind}}(G|_{I(\tilde{x})})} |M|$ and let $\underline{\Omega} = \{(\underline{x}, i) | i \in I\}$.

Since in every SSPE, the players outside $I(\underline{x})$ never change their efforts, and since no player in $I(\underline{x})$ is connected to a player outside of it with positive efforts, they do not influence the other players' payoffs in any way. We can therefore remove them from the network and take the network $G|_{I(\underline{x})}$. Since $I(\underline{x}) \subseteq \bar{N}_1$ we can apply Proposition 3.6.2. Hence for every $\underline{\omega} \in \underline{\Omega}$ it holds that

$$\underline{u}(G|_{I(\underline{x})}) = f_1(\underline{e}(G|_{I(\underline{x})})).$$

By the choice of \underline{x} , we have $f_1(\underline{e}(G|_{I(\underline{x})})) = f_1(\underline{e}(G))$, and we have $\underline{u}(G|_{I(\underline{x})}) \geq \underline{u}(G)$ due to

$$\underline{u}(G) = \inf_{\omega \in \Omega, s^* \in S^*} \lim_{t \rightarrow \infty} u_1^t(\omega, s^*) \leq \inf_{\underline{\omega} \in \underline{\Omega}, s^* \in S^*} u_1^t(\underline{\omega}, s^*).$$

Finally, we show that if $\underline{e}(G) > 0$ then $\underline{u}(G) \geq f_1(\underline{e}(G))$ also holds. We prove this by induction. For $n = 2$ the statement holds by Proposition 3.6.2. Now fix n and assume that it holds for every network with fewer than n players. For networks that belong to \mathcal{G}_1 the statement holds by Proposition 3.6.2, so we can assume that $G \notin \mathcal{G}_1$.

Suppose that $\underline{u}(G) < f_1(\underline{e}(G))$. Let $\hat{X} = \{\hat{x} \in X : \tilde{\mathcal{B}}(\hat{x}) \neq \emptyset\}$. Since $G \notin \mathcal{G}_1$, \hat{X} must be non-empty, as any effort profile \hat{x} for which $\hat{x}_j = 1$ for some $j \notin \overline{N}_1$ and $\hat{x}(N_j) = 0$ will satisfy this property. For every $\hat{x} \in \hat{X}$ and every $B \in \tilde{\mathcal{B}}(\hat{x})$ we have that $\underline{e}(G) \leq \underline{e}(G|_{I \setminus \overline{N}_B})$, since combining any closed partnership in the reduced network $G|_{I \setminus \overline{N}_B}$ with B gives a closed partnership in G . We now construct a strategy $s^{**} \in \tilde{S}_1 \times \{b_{-1}\}$ such that for every $\omega \in \Omega$ it holds that $\lim_{t \rightarrow \infty} u_1^t(\omega, s^{**}) \geq f_1(\underline{e}(G))$ to achieve a contradiction.

Consider the following strategy:

$$s_1^{**}(x) = \begin{cases} 0 & \text{if } x \in X^* \text{ and either } x(N_1) = 0 \text{ or } x(N_1) \geq \underline{e}(G), \\ 1 & \text{otherwise.} \end{cases}$$

If player 1 adopts s_1^{**} , then by Lemma 3.4.1 the game eventually reaches a state $y \in X^*$ for which either $y(N_1) \geq \underline{e}(G)$ or $y(N_1) = 0$. If it is the former, we are done. If it is the latter, then by the definition of s^{**} , there exists $y^0 \in R_{s^{**}}(y) \cap X_+^*$ for which $y^0(N_1) = 0$ and $y_1^0 = 0$. There are two cases. If $\tilde{\mathcal{B}}(y^0) \neq \emptyset$, then take a $B \in \tilde{\mathcal{B}}(y^0)$ and remove \overline{N}_B from the game altogether. By the induction hypothesis it is true that $\underline{u}(G|_{I \setminus \overline{N}_B}) \geq f_1(G|_{I \setminus \overline{N}_B})$, hence player 1 has an SSPE in the reduced game with player set $I \setminus \overline{N}_B$ which guarantees at least a payoff of $f_1(G|_{I \setminus \overline{N}_B}) \geq f_1(\underline{e}(G))$.

If $\tilde{\mathcal{B}}(y^0) = \emptyset$, then there must exist $i^0 \in N_1$ with $y^0(\overline{N}_{i^0}) < 1$. By Lemma 3.4.4 there exists $\hat{y}^0 \in R_{s^{**}}(y^0)$ in which $\hat{y}_{i^0}^0 = 1$, $\hat{y}^0(N_{i^0}) = 0$, and for every $j \notin \overline{N}_{i^0}$ we have $\hat{y}_j^0 = y_j^0$.

Once i^0 usurped his neighborhood's efforts we again let every other myopic player reach their best responses again without player 1 moving. By Lemma 3.4.1, there exists a $y^1 \in R_{s^{**}}(\hat{y}^0) \cap X^*$ such that $y_1^1 = y_1^0 = 0$. Notice that $y_{i^0}^1 = \hat{y}_{i^0}^0 = 1$, while $y^1(N_{i^0}) = \hat{y}^0(N_{i^0}) = 0$. We now repeat the argument we had for y^0 . If $\tilde{\mathcal{B}}(y^1) \neq \emptyset$, then take a $B \in \tilde{\mathcal{B}}(y^1)$, remove \overline{N}_B from the game, and apply the induction hypothesis to show the statement. If not, and if $\underline{e}(G) = 1$, then we are done. If not, and if $\underline{e}(G) > 1$, then there must exist an $i^1 \in N_1$ for which $y^1(\overline{N}_{i^1}) < 1$. Then, by Lemma 3.4.4 there exists $\hat{y}^1 \in R_{s^{**}}(y^1)$ for which $\hat{y}_{i^1}^1 = 1$, $\hat{y}^1(N_{i^1}) = 0$, and for every $i \in I \setminus \overline{N}_{i^1}$ it holds that $\hat{y}_i^1 = y_i^1$.

We continue this way until, for some $k \leq |N_1|$, we reach a \hat{y}^k such that for every $i \in N_1$ we have $\hat{y}_i^k = b_i(\hat{y}^k)$. Notice that in \hat{y}^k a maximal independent set of N_1 is playing 1 while the rest of N_1 is playing 0. Let this independent set be denoted by M^* and notice that $\underline{e}(G) \leq |M^*|$.

Once we reach \hat{y}^k we let all remaining myopic players reach their best replies as per Lemma 3.4.1. Let this effort profile be denoted by y^* . It is easy to see that

$y^* \in X^*$, and that $\pi_1(y^*) = f_1(|M^*|) \geq f_1(\underline{e}(G))$. Note that the definition of s^{**} ensures that we either find a closed partnership that we can remove, or the game reaches y^* with probability 1. Hence, a minimum long-run instantaneous payoff of $f_1(\underline{e}(G))$ is attainable from any effort profile via strategy s^{**} and the same must hold for every SSPE profile s^* . Hence $\underline{u}(G) \geq f_1(\underline{e}(G))$. ■

Chapter 4

Best-response dynamics in directed network games

“I do detest everything which is not perfectly mutual.”

Lord Byron

This chapter¹ studies the cycling behavior of a simple learning process in network games of non-reciprocal relationships. Unlike in the well-studied case of reciprocal relations, in directed network games one-sided best-response dynamics may produce cycles. This chapter presents two subclasses of such games in which best-response cycles cannot occur, leading to convergence to the set of Nash equilibria. The first such class features a hierarchical relationship between the players, and is characterized by a directed acyclic network. The second features weak interaction of the players, more specifically, the total external effects of each player’s decision are lower than the effect on own payoffs.

4.1 Introduction

The Nash equilibrium is one of the central concepts of game theory. One of its most attractive properties is that it applies to all games that can be represented in a normal form. This great generality, however, does not come with a flawless record in

¹This chapter is based upon Bayer and Kozics (2019). I would like to thank Botond Kőszegi, Miklós Pintér, and Ádám Szeidl for comments on previous versions of this chapter, as well as Yannick Viossat for feedback and suggestions.

applied situations. The concept’s success as an accurate predictor of outcomes of real-life strategic situations is mixed, giving excellent predictions in some (e.g. congestion games), but poor ones in others (e.g. one-shot public goods games, beauty contests). Nash equilibrium’s excellent theoretical properties and its shortcomings in applied settings together imply that there are fundamental differences in behavior in games where it provides an accurate prediction and in those where it fails to.

The theory of learning provides the toolkit to capture these differences. By hypothesizing that equilibrium behavior arises through the adaptation rather than the introspection of the players, learning serves as the principal toolkit to provide motivation to be interested in the equilibrium behavior in games: If general learning processes converge to a Nash equilibrium, then it may be used as a prediction. Additionally, as some Nash equilibria may be more likely to be limit points of learning processes than others, learning may help resolve problems arising from equilibrium selection.

In this chapter, we analyze a generalized version of Chapter 2’s model of weighted network games, called directed network games. As explained previously in Chapters 2 and 3, network games constitute an extremely rich framework, with a wide range of applications in many fields of science. Unlike in the models of Chapter 2 and Chapter 3, directed network games allow for players to have non-reciprocal relations to each other. Non-reciprocal relationships may represent (1) one-sided links, in which case a player influences another player’s payoffs but not vice versa, (2) two-sided, complementary, but unequal links, in which case one player has more of an influence over another player’s payoffs than the other way around, or (3) two-sided, parasitic links which provide benefits for one player but provide harm to the other.

A brief overview of relevant theoretical literature is as follows: Public goods games played on networks have been introduced by Bramoullé and Kranton (2007). For these games, the stability of Nash equilibria with respect to a simple learning process, the continuous best-response dynamic has been established by Bramoullé et al. (2014). Uniform convergence of the continuous best-response dynamic to a single Nash equilibrium in all such games has been shown by Bervoets and Faure (2019). Bervoets et al. (2016) constructs an unsophisticated, convergent learning processes. All of these papers assume reciprocal relations between the players of the network, relying on the theory of potential games (Monderer and Shapley, 1996) to achieve their results. The same assumption is made in Ballester et al. (2006)’s closely related model on peer effects.

To our knowledge, this is the first model focusing on non-reciprocal relationships in the class of network games. As shown by Monderer and Shapley (1996), the existence

of a potential function is closely related to the reciprocity of interactions between every pair of players, a property that cannot be exploited in our more general setting. We focus on the cycling properties of one-sided best-response dynamics. It is easy to see that even in a simple game of non-reciprocal interactions, a game played on a directed cycle network, best-response cycles may emerge, leading to non-convergence, and a lack of motivation for the Nash equilibrium. Nevertheless, classes of networks exist that give rise to good cycling behavior.

In this chapter we identify two such subclasses, called hierarchical games, and games of weak interaction. Hierarchical games are characterized by directed acyclic networks. They are applicable in situations when pairwise interactions are one-sided, as described by point (1). Populations that adhere to this criterion have some underlying hierarchy. Players on the highest level are not affected by any other player, players on intermediate levels are affected by those on higher levels but not by those on lower levels, while players on the lowest level have no effect on any other player. A specific application of hierarchical games in the realm of environmental economics is waste-water management along a river. In this case the players of the game are municipalities located along the river, while the hierarchy follows the flow of the river, with upstream municipalities at the top of the hierarchy. Other examples of networks satisfying this criterion include trophic networks and military command structures.

The second subclass of games with good cycling properties are games of weak interaction. Such games are characterized by networks with weights that are relatively small in absolute value. Specifically, for a game to be of weak interaction, every player's influence on his own payoffs has to be higher than his influence on every other player's payoff combined. The strength of the interactions represents the level of the interconnectedness of the network. In social networks, low interconnection can be interpreted as a form of individualism. In economics, weak interaction is a characteristic of efficient markets, while strong interaction indicates the possibility of arbitrage. In biology, weak, but non-negligible interaction is a characteristic of the foraging behavior of most solitary animals, as well as of some social animals.

Our results imply that the Nash equilibrium retains some of its predictive power in games on networks even if the relations are non-reciprocal. Specifically, the acyclicity of the best-response dynamic means that many intuitive one-sided learning processes converge for these two subclasses of games, if the strategy spaces are finite – which they are in most applications. The predictive power of the Nash equilibrium is therefore shown if (1) there is a clear hierarchy on the players, or (2) interactions are not too strong.

This chapter is organized as follows: Section 4.2 presents our two main concepts, directed network games and best-response dynamics. Section 4.3 contains the main results, the acyclicity of best-response dynamics in hierarchical games and games of weak interaction. Section 4.4 concludes.

4.2 The model

Let $I = \{1, \dots, n\}$ be the set of players. For $i \in I$ and bounds $\underline{x}_i < \bar{x}_i$ the set $X_i = [\underline{x}_i, \bar{x}_i]$ is called the action set of player i , $X = \prod_{i \in I} X_i$ is called the set of action profiles, and $X_{-i} = \prod_{j \neq i} X_j$ the set of action profile for all players except i . Let $x_i \in X_i$ denote the action taken by player i , $x_{-i} \in X_{-i}$ the truncated action profile of all players except player i , and $x = (x_i)_{i \in I}$ the action profile of all players.

Definition 4.2.1. The tuple $G = (I, X, (\pi_i)_{i \in I})$ is called a *directed network game* with payoff functions $\pi_i: X \rightarrow \mathbb{R}$ given by

$$\pi_i(x) = f_i \left(\sum_{j \in I} w_{ij} x_j \right) - c_i x_i, \quad (4.1)$$

where for every $i \in I$, we have that $f_i: \mathbb{R} \rightarrow \mathbb{R}$ is twice differentiable, $f'_i > 0$, $f''_i < 0$, $w_{ij} \in \mathbb{R}$, $c_i \geq 0$.

The interpretation is the following. Each player produces a specialized good with linear production technology, incurring costs c_i for every unit of the good produced. Players derive benefits from the consumption of their own goods and they are affected by their opponents' production decisions. Player i 's enjoyment of player j 's good is represented by the weight $w_{ij} \in \mathbb{R}$. Crucially, unlike in Chapters 2 and 3, we do not impose reciprocal relations, i.e. $w_{ij} \neq w_{ji}$ may hold. As before, we normalize the interaction parameter of each player i with himself, w_{ii} , to 1. The overall benefits of player i are given by the benefit function f_i over the weighted sum of his and his opponents' goods.

Note that if $w_{ij} > 0$ and $w_{ji} > 0$, then the goods of players i and j are strategic substitutes. If $w_{ij} < 0$ and $w_{ji} < 0$ then their goods are strategic complements. If $w_{ij} > 0$ and $w_{ji} < 0$, then players i and j share a parasitic link.

Since the benefit functions f_i are increasing and concave, and the cost parameters c_i are positive, for a given x_{-i} , there is a unique value of x_i that maximizes $\pi_i(x)$. For player $i \in I$, let b_i denote the best-response functions, i.e. for every $i \in I$ and $x \in X$, $b_i(x) = \operatorname{argmax}_{x_i} \pi_i(x)$. Let the target values t_i be implicitly defined by $f'_i(t_i) = c_i$ as in Chapter 2. Unlike in Chapter 3 we do not assume that players have symmetric

target consumption levels. If f'_i is always larger (smaller) than c_i then t_i is defined as positive (negative) infinity.

Lemma 4.2.2. *Let $G = (I, X, (\pi_i)_{i \in I})$ be a weighted network game. Then, for every $i \in I$ and $x \in X$ the best response functions are the following:*

$$b_i(x) = \begin{cases} \underline{x}_i & \text{if } t_i - \sum_{j \neq i} w_{ij} x_j < \underline{x}_i, \\ t_i - \sum_{j \neq i} w_{ij} x_j & \text{if } t_i - \sum_{j \neq i} w_{ij} x_j \in [\underline{x}_i, \bar{x}_i], \\ \bar{x}_i & \text{if } t_i - \sum_{j \neq i} w_{ij} x_j > \bar{x}_i. \end{cases} \quad (4.2)$$

Proof. First, we calculate the unconstrained best response for player i . The first-order condition is:

$$\frac{\partial \pi_i(x)}{\partial x} = f'_i \left(\sum_{j \in I} w_{ij} x_j \right) - c_i = 0, \quad (4.3)$$

as $w_{ii} = 1$. Combining with $f'_i(t_i) = c_i$, we get that the unconstrained best response, $\tilde{b}_i(x)$ is

$$\tilde{b}_i(x) = t_i - \sum_{j \neq i} w_{ij} x_j. \quad (4.4)$$

As the second-order condition is

$$\frac{\partial^2 \pi_i(x)}{\partial x^2} = f''_i \left(\sum_{j \in I} w_{ij} x_j \right) < 0, \quad (4.5)$$

$\tilde{b}_i(x)$ is indeed maximizing the payoff. This means, that for every $x_i > b_i(x)$, a marginal increase of x_i decreases $\pi_i(x)$, while for every $x_i < b_i(x)$, a marginal increase of x_i increases $\pi_i(x)$. Therefore, if $\tilde{b}_i(x) \in [\underline{x}_i, \bar{x}_i]$, then $b_i(x) = \tilde{b}_i(x)$. If $\tilde{b}_i(x) < \underline{x}_i$, then $b_i(x) = \underline{x}_i$, as choosing a larger x_i would decrease the payoff. Similarly, if $\tilde{b}_i(x) > \bar{x}_i$, then $b_i(x) = \bar{x}_i$. ■

Lemma 4.2.3. *Every directed network game has a Nash equilibrium.*

Bramoullé et al. (2014)'s analogue result using Brouwer's fixed-point theorem for a positive and symmetric weight matrix is directly applicable in the directed network case. Let the set of Nash equilibria be denoted by X^* .

Although Lemmas 4.2.2 and 4.2.3 have close analogues in existing models with symmetric interaction, most results, especially those concerning the cycling and convergence of best-response dynamics do not hold in the directed network case. We raise an example of a well-known case of non-reciprocal interaction.

Example 4.2.4. Let $I = \{1, 2, 3\}$, let $X_i = [0, 1]$ for $i \in I$, $t = (1, 1, 1)^\top$, and

$$W = \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix}.$$

By Lemma 4.2.2, the best-response functions are $b_1(x) = 1 - x_3$, $b_2(x) = 1 - x_1$, $b_3(x) = 1 - x_2$. The only Nash equilibrium is $x^* = (0.5, 0.5, 0.5)^\top$.

Example 4.2.4 is inspired by games with cyclic best-response dynamics, such as the rock-paper-scissors game. Such games are extensively used in applications involving evolutionary game theory, particularly in modeling ecological interactions.

We now introduce the class of learning processes this chapter is focused on, the best-response dynamic.

Definition 4.2.5. A sequence of action profiles $(x^k)_{k \in \mathbb{N}}$ is called a *one-sided best-response dynamic*, or *best-response dynamic* for short, if for every $k \in \mathbb{N}$ there exists an $i^k \in I$ such that $x_{-i^k}^k = x_{-i^k}^{k+1}$ and $x_{i^k}^{k+1} = b_{i^k}(x^k)$.

The set of sequences that satisfy Definition 4.2.5 are the same as the infinite best-response compatible paths of Chapter 2. In a one-sided learning process, each period sees exactly one player changing his action. As in Chapter 2 we make no explicit assumptions on how the revision opportunities are assigned, only that each period sees only one player revising. In a one-sided process governed by the best-response dynamic, every player revises to his best response given the previous period's action profile. The actions of every other player remain the same. It is possible that the actions of all players remain the same, i.e. $x^k = x^{k+1}$ for some k , but only if at least one player's action was already his best response.

Definition 4.2.6. A sequence of action profiles $(x^k)_{k \in \mathbb{N}}$ has a cycle if there exist three time periods, $k < k' < k''$ such that $x^k = x^{k''}$, but $x^k \neq x^{k'}$.

A process has a cycle, if it revisits an action profile in two different time periods, between which there is a time when it moves away from it. Cycling properties are extremely important for the convergence of the learning process. For instance, in case the absence of cycles can be shown for a process, then it will always converge in a finite (i.e. discretized) setting of the game.

Cycles in a best-response dynamic are called best-response cycles. As shown in Chapter 2's Section 4, in games with a symmetric weight matrix there are no best-response cycles. However, this is not true in general for directed network games. This is shown by the following example.

k	x_1^k	x_2^k	x_3^k	i^k	$\sum_{j \in I} w_{ikj} x_j$	$b_{ik}(x_{ik})$
0	1	0	0	3	0	1
1	1	0	1	1	2	0
2	0	0	1	2	0	1
3	0	1	1	3	2	0
4	0	1	0	1	0	1
5	1	1	0	2	2	0
6	1	0	0			

Table 4.1: The best-response cycle of Example 4.2.7.

Example 4.2.7. Consider the game outlined in Example 4.2.4, with $I = \{1, 2, 3\}$, $X_i \in [0, 1]$ for $i \in I$, $t = (1, 1, 1)^\top$ and

$$W = \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix}.$$

Notice that the game has a single Nash equilibrium, $x^* = (0.5, 0.5, 0.5)^\top$.

Suppose that the initial action profile of the players is $x^0 = (1, 0, 0)^\top$. Table 4.1 shows a best-response cycle of length 6.

We recall another definition from Chapter 2, that of a best-response potential game.

Definition 4.2.8. [Best-response potential game (Voorneveld, 2000)] A game $G = (I, X, (\pi_i)_{i \in I})$ is a *best-response potential game*, if there exists a *best-response potential function* $\phi : X \rightarrow \mathbb{R}$ such that for every $i \in I$, and every $x_{-i} \in X_{-i}$ it holds that

$$\operatorname{argmax}_{x_i \in X_i} \pi_i(x) = \operatorname{argmax}_{x_i \in X_i} \phi(x). \quad (4.6)$$

The symmetry of the weight matrix W and the game having a best-response potential are tightly related. For symmetric weight matrices, Bramoullé et al. (2014) has shown that the function $\phi(x) = x^\top t - \frac{1}{2} x^\top W x$ is an exact potential of Ballester et al. (2006)'s game of negative externalities, given by the following, quadratic payoff functions:

$$u_i(x) = t_i x_i - \frac{1}{2} x_i^2 - \sum_{j \neq i} w_{ij} x_i x_j.$$

Since u_i and π_i give rise to the same best-response functions, ϕ is, as discussed in Chapter 2, a best-response potential. However, by Theorem 4.5 of Monderer and Shap-

ley (1996), ϕ can only be an exact potential of the game characterized by $(I, X, (u_i)_{i \in I})$, if for every $i, j \in I$ we have

$$\frac{\partial^2 u_i}{\partial x_i \partial x_j} = \frac{\partial^2 u_j}{\partial x_i \partial x_j},$$

which amounts to

$$w_{ij} = w_{ji}.$$

This suggests that ϕ has no analogous forms that could serve as a best-response potential for the $w_{ij} \neq w_{ji}$ case. As shown in Example 4.2.7, non-reciprocal relationships allow for the emergence of best-response cycles, the absence of a best-response potential, and thus, non-convergence.

4.3 Classes of games with acyclical best-response dynamics

As shown in the previous section, allowing for non-reciprocal interactions in network games has grave consequences on the convergence properties of even the simplest learning processes, i.e. the one-sided best-response dynamic which may now produce best-response cycles. Nevertheless, there are useful classes of network games where the absence of cycles can still be shown without restrictive assumptions on the learning process. In this section we present two of them, hierarchical games and games with weak interactions.

Consider the following example.

Example 4.3.1. Let $I = \{1, 2, 3\}$ be a set of three cities with industrial zones located along a river. The river flows from city 1 to city 2 and then to city 3. Each city i decides on the amount of money spent on cleaning the industrial waste in the river out of a budget of 2, i.e. $X_i = [0, 2]$. The cities' target values describe the point at which the marginal benefits of an extra euro's worth of cleaner water are the same as the costs for that city, and are as follows: $t = (1, 2, 3)^\top$. The targets increase downstream, indicating that in the absence of cleaning expenses from other cities, a downstream city requires higher cleaning efforts due to the extra pollution from its upstream counterparts' industries. The interaction weights between the cities are given

as follows:

$$W = \begin{pmatrix} 1 & 0 & 0 \\ 0.5 & 1 & 0 \\ 0.25 & 0.5 & 1 \end{pmatrix}.$$

City 1 is unaffected by the decisions of cities 2 and 3 due to its upstream position. City 2 derives benefits from city 1's cleaning efforts, but less so than from that of its own. City 3 derives benefits from both cities' cleaning efforts, again less so than from that of its own, and less so for the more distant city 1.

As per Lemma 4.2.2 the best-response functions of the three players are as follows: $b_1(x) = 1$, $b_2(x) = \min\{2, \max\{0, 2 - 0.5x_1\}\}$, and $b_3(x) = \min\{2, \max\{0, 3 - 0.25x_1 - 0.5x_2\}\}$. The unique Nash equilibrium is $x^* = (1, 1.5, 2)^\top$.

Example 4.3.1 reflects a much-stylized situation, inspired by one-dimensional models of river pollution such as the one appearing in Pimpunchat et al. (2009). As suggested by the example, typical models of pollution cannot be accommodated by a symmetric interaction matrix, as the external effects typically go only one way. A further characteristic of this example is the hierarchical relationship of the best-responses, as the payoffs of cities are independent of the actions of downstream cities, and the underlying network has no cycles.

In addition to river pollution, such hierarchies are present in most production chains where goods – and therefore externalities – flow downwards in the production chain, or in some social networks such as the military where orders are traveling down the chain of command. Many trophic networks also have hierarchical features with apex predators on the highest level of the hierarchy and prey animals located on lower levels. A formal characterization of such networks is as follows.

Definition 4.3.2. A game is called *hierarchical* if for its underlying network W it holds that $w_{ij} = 0$ for every $i < j \in I$.

In other words, if the underlying network is acyclic, we have a hierarchical game. In Definition 4.3.2 it is assumed that players with lower indices are higher up in the hierarchy, i.e. player 1 is unaffected by any other player's action, player 2 is only affected by player 1, etc., which is without loss of generality.

Proposition 4.3.3. *For every hierarchical game and every best-response dynamic $(x^k)_{k \in \mathbb{N}}$ it holds that there are no cycles in $(x^k)_{k \in \mathbb{N}}$.*

Proof. Suppose that there is a cycle $(x^k)_{k \in \{\ell, \dots, \ell''\}}$ with $\ell < \ell''$ such that $x^\ell = x^{\ell''}$ but $x^\ell \neq x^{\ell'}$ for some $\ell' \in \{\ell, \dots, \ell''\}$. Let i denote the player with the lowest index such that there exists an $\ell_i \in \{\ell, \dots, \ell''\}$ with $x_i^\ell \neq x_i^{\ell_i}$. Such a player must exist, otherwise no player has changed his action during this time period, contradicting the definition of a cycle.

Since for every player $i' < i$ we have that $x_{i'}^\ell = x_{i'}^{\ell+1} = \dots = x_{i'}^{\ell''}$, by the definition of hierarchical games we have that $\sum_{j \neq i} w_{ij} x_j^\ell = \sum_{j \neq i} w_{ij} x_j^{\ell+1} = \sum_{j \neq i} w_{ij} x_j^{\ell''}$ and hence $b_i(x^\ell) = b_i(x^{\ell+1}) = \dots = b_i(x^{\ell''})$. However, since $(x^k)_{k \in \{\ell, \dots, \ell''\}}$ comes from a best-response dynamic, this contradicts the $x_i^\ell \neq x_i^{\ell_i}$. ■

The idea behind the proof of Proposition 4.3.3 is the observation that for a player to revisit an action played earlier, a player in a higher level of the hierarchy must have changed his action at some point. Because of this, the player at the highest level cannot revisit earlier actions, as once he reaches his best response, he stays there indefinitely, for the movements on the lower levels do not change his payoffs, and therefore, his best response. Hence, the highest level player must not change his action throughout a cycle, meaning that the second highest player also cannot revisit earlier actions, and so on.

To set up the next class of games, consider a parametric version of Examples 4.2.4 and 4.2.7.

Example 4.3.4. Consider a game with circular externalities, similar to the one presented in Example 4.2.4, with the strength of the interactions represented by the parameter $\delta \in [0, 1]$.

$$W = \begin{pmatrix} 1 & 0 & \delta \\ \delta & 1 & 0 \\ 0 & \delta & 1 \end{pmatrix}.$$

By Lemma 4.2.2, the best-response functions are $b_1(x) = 1 - \delta x_3$, $b_2(x) = 1 - \delta x_1$, $b_3(x) = 1 - \delta x_2$. The only Nash equilibrium is $x^* = (1/(1 + \delta), 1/(1 + \delta), 1/(1 + \delta))^\top$.

If the strength of the interaction is high, i.e. $\delta = 1$, then best-response cycles may exist as per Example 4.2.7. If the strength of the interaction is zero, i.e. $\delta = 0$, then the game is a hierarchical game, as well as a weighted network game, meaning that best-response cycles cannot exist.

In Table 4.2 we show the sequence of action profiles that follows if players receive revision opportunities in the same order as in the best-response cycle of Example 4.2.7, starting, again, in the action profile $(1, 0, 0)^\top$. It is clear from the table and from the

k	x_1^k	x_2^k	x_3^k	i^k	$\sum_{j \in I} w_{ikj} x_j$	$b_{ik}(x_{ik})$
0	1	0	0	3	0	1
1	1	0	1	1	δ	$1 - \delta$
2	$1 - \delta$	0	1	2	$\delta - \delta^2$	$1 - \delta + \delta^2$
3	$1 - \delta$	$1 - \delta + \delta^2$	1	3	$\delta - \delta^2 + \delta^3$	$1 - \delta + \delta^2 - \delta^3$

Table 4.2: The best-response dynamic of Example 4.3.4.

best-response functions, that in this order of revisions, the player holding the revision opportunity in period k will play $\sum_{\ell=0}^k (-\delta)^\ell = (1 - (-\delta)^{k+1})/(1 + \delta)$ in the next, for parameter values of $\delta < 1$. Playing on in this order will produce no cycles, and lead to convergence to the Nash equilibrium.

Intuition suggests that if the interaction parameters are low enough, then the game maintains the good cycling and convergence properties characteristic of games with reciprocal interaction. Example 4.3.4 suggests that the critical value of cross-player interaction which leads to cycling is the value by which players interact with themselves. In the next definition we introduce games of weak interaction along the above ideas, formally.

Definition 4.3.5. A game is of *weak interaction* if for every $i \in I$ the underlying network satisfies $\sum_{j \in I \setminus \{i\}} |w_{ji}| < 1$.

By Definition 4.3.5, a game is said to be of weak interaction if the sum of absolute value of every player's outgoing weights is lower than unity, i.e. the player's interaction parameter with himself. In such games, the total external effects of a player's actions on other players' best-responses are lower than the effect of his actions on his own best-response. Notice that the definition does not impose that the incoming weights, i.e. a player's sensitivity to the other players' actions, be lower than unity. Crucially, weak interaction is markedly different from no interaction, as the external effects are of the same order as the players' influence on themselves. However, as the number of players grow large, the per capita level of interaction approaches zero.

Examples of games of low interaction include games in efficient financial and economic networks. In this context, low interaction means that the players' own financial or production decisions affect own profits more than other players on the market. If the converse was true, there would be a possibility of arbitrage by increasing the profits of other agents (by increasing volume if the weights are positive, and by dropping it

if the weights are negative), and being compensated by them in return. A biological example of low interaction would be foraging games, where, typically, animals' foraging strategies affect the given animal more than its opponents.

The final result of this chapter states that games of weak interaction can never give rise to best-response cycles.

Proposition 4.3.6. *For every game of weak interaction and every best-response dynamic $(x^k)_{k \in \mathbb{N}}$ it holds that there are no cycles in $(x^k)_{k \in \mathbb{N}}$.*

Proof. Let the function $L: X \rightarrow \mathbb{R}$ be defined as follows:

$$L(x) = \sum_{i \in I} |x_i - b_i(x)|.$$

Clearly, the function's value is 0 if and only if x is a Nash equilibrium.

Fix $k \in \mathbb{N}$. We now calculate the change in the value of L between periods k and $k + 1$. If $x_{i^k}^k = b_{i^k}(x^k)$, then $x^{k+1} = x^k$, hence the change is zero. If $x_{i^k}^k = x_{i^k}^{k+1} + d = b_{i^k}(x^k) + d$ for some $d \neq 0$, then by the properties of best-response dynamics we have

$$L(x^{k+1}) - L(x^k) = \sum_{j \neq i^k} |x_j^k - b_j(x^{k+1})| - \sum_{j \neq i^k} |x_j^k - b_j(x^k)| - |d|.$$

For every $j \in I \setminus \{i^k\}$ we have $|x_j^k - b_j(x^{k+1})| - |x_j^k - b_j(x^k)| \leq |b_j(x^k) - b_j(x^{k+1})|$, hence we get

$$L(x^{k+1}) - L(x^k) \leq \sum_{j \neq i^k} |b_j(x^k) - b_j(x^{k+1})| - |d|.$$

Finally, by the properties of the best-response functions we have

$$L(x^{k+1}) - L(x^k) \leq \sum_{j \neq i^k} |w_{ji}| |d| - |d| < 0,$$

where the last inequality is due to the weak interaction property.

Hence, every time a player changes his action in a best-response dynamic, L 's value strictly decreases, meaning that we cannot get cycles. ■

Proposition 4.3.6 is shown directly by calculating the change in the aggregated distances from their current best responses, as players revise their actions. The weak interaction property guarantees that the total distance can never increase along a best-response dynamic, meaning that acyclicity is guaranteed. Since all inequalities used in the proof are tight, the result itself is tight. As demonstrated by Example 4.2.7, allowing the total magnitude of the interaction parameters to sum up to 1 leads to the emergence of best-response cycles.

Since the weakness of interaction is a good measure of the individualism present in a social network, this result implies that more individualistic societies allow for an increased predictive power of the concept of Nash equilibrium, as games played on these societies will have better convergence properties. In games played on financial networks our result hints that more efficient markets allow for more accurate predictions of the Nash equilibrium.

4.4 Conclusion

In this chapter we analyze the cycling behavior of a simple learning process, called the best-response dynamic in directed network games. Directed network games are an extension of the private provision of public goods game, or local public goods game, allowing for players to have non-reciprocal relationships with one another. To our knowledge, the paper on which this chapter is based is the first to consider such relations in this gameclass. Non-reciprocity of relations occurs in social networks between two agents if a link goes only one way, a link is more significant for one than for the other, or the relationship is beneficial for one party and harmful for the other. In economics, most externalities, such as pollution are non-reciprocal: they either only travel one way, or the magnitude by which the agents affect one another is different. In biology, trophic networks typically feature one-way interactions, while many other interspecies relations, such as parasitism may be grossly non-reciprocal.

There are several existing results on the acyclicity and convergence of several learning processes if relationships are reciprocal. These results rely on the theory of potential games, particularly on the fact that classic results by Monderer and Shapley (1996) imply that reciprocal relationships give way to the existence of a best-response potential function. Because of this, every one-sided best-response update, or any update taking the player closer to his current best response will increase the value of the potential, meaning that these processes are acyclical.

While best-response potentials do not exist in general for games featuring non-reciprocal relations, alternative methods may be used to analyze learning in subclasses of directed network games. In this chapter we present two such subclasses, hierarchical games and games of weak interaction. Hierarchical games are characterized by directed acyclical networks, featuring one-way links with no cycles. Games of weak interaction are networks in which every player's actions affect the other players' decisions less so than his own. Both games have a wide range of applications and together they cover

a significant part of the applications featuring non-reciprocal interaction.

In this chapter we show that both of the above mentioned subclasses lead to acyclical best-response dynamics with strong implications on the convergence of these learning processes. In discretized games, acyclicity is a sufficient condition for convergence. In case of a best-response dynamic, the convergence is to the set of Nash equilibria. As most applications are closely captured by discretized strategy spaces, our results show that the predictive power of the Nash equilibrium extends beyond games with reciprocal relationships.

Additional work is needed to establish convergence conditions in games with continuous strategy spaces. By introducing additional assumptions, such as on the order of players receiving revision opportunities, further classes of games may be identified with good cycling and convergence properties. These plans are left for future research.

Chapter 5

A two-phenotype model of immune evasion by cancer cells

“And now for something completely different.”

Monty Python

This chapter¹ proposes a model with two types of cancer cells differentiated by their defense mechanisms against the immune system. “Selfish” cancer cells develop defense mechanisms that benefit the individual cell, whereas “cooperative” cells deploy countermeasures that increase the chance of survival of every cell. Our phenotypes capture the two main features of the tumor’s efforts to avoid immune destruction, crypticity against immune cells for the selfish cells, and tumor-induced immunosuppression for the cooperative cells. We identify steady states of the system and show that only homogeneous tumors can be stable in both size and composition. We show that under generic parameter values, a tumor of selfish cells is more benign than a tumor of cooperative cells, and that a treatment against cancer crypticity may promote immunosuppression and increase cancer growth.

¹This chapter is based upon Bayer, Brown, and Staňková (2018). The majority of this project was conceived during my research visit at the University of Illinois at Chicago between July 23 and August 23, 2016. The visit was sponsored by the European Union’s Horizon 2020 research and innovation program. I would like to thank Jessica Cunningham, Abdel Halloway, and Christopher Whelan for feedback and suggestions.

5.1 Introduction

The immune system influences cancer initiation and progression. One of the hallmarks of cancer is evasion of the immune system (Hanahan and Weinberg, 2011). The immune system is how multi-cellular organisms mount an adaptive response to diseases and pathogens (including cancer). There are innate and adaptive mechanisms of immunity. Innate immunity provides a relatively undirected but permanent defense against pathogens using a variety of white blood cells that include phagocytes and natural killer cells. Dendritic cells (a subset of phagocytes) provide a link between the innate and adaptive immune system. Upon encountering foreign proteins or molecules (antigens) the dendritic cells can entrain T-lymphocytes (T-cells) by modifying their receptor proteins to recognize the antigen. As part of the adaptive immune system, the killer T-cells now possess receptors that will recognize the surface antigen on nucleated cells such as infectious protozoans, viral infected normal cells, and even cancerous cells. Upon making contact, the T-cell is able to breach the target cell's membrane and introduce lethal cytotoxins.

When treating cancer, a high immune system infiltration into a tumor often begets a positive prognosis (Parcesepe et al., 2016) as immune cells inhibit the growth and spread of the tumor. Immunotherapy tries to trigger an effective immune response to the cancer (Dimberu and Leonhardt, 2011). Such therapies may introduce retroviruses into the tumor. The retroviruses present antigens that induce an immune response. Or, the patient's own immune system may be boosted by entraining T-cells on cancer cells drawn from the patient (Morgan et al., 2006). The antigens and entrained killer T-cells are then injected into the patient as a form of 'vaccine' to create a more directed and effective immune response to the cancer cells within the tumor. Finally, checkpoint therapies directly target the immune evasion traits of cancer cells (Goswami et al., 2016) making it easier for T-cells to encounter and attack cancer cells.

The adaptive immune response to cancer cells via T-cell activation has been conceptualized and modeled as a competition between killer T-cells (as well as associated macrophages, helper-cells, neutrophils, etc.) and the cancer cells. Eftimie et al. (2011) provide an extensive review of common spatially homogeneous mathematical models describing the interactions between a malignant tumor and the immune system, starting from the single equation models for tumor growth and adding complexity. Nani and Freedman (2000) construct a model of competition between normal cell and cancer cells during the process of immunotherapy. Robertson-Tessi et al. (2012) and Nakada et al. (2016) focus on the components of the immune system. Our model complements

the literature by focusing on the various strategies available to the cancer cells in the game played by the cancer cells and the immune system.

Unfortunately for the patient, cancer cells evolve resistance to the immune system and to immunotherapies (Ribas, 2015; Sharma et al., 2017). A number of papers have used agent based models to examine the evolution of resistance to various types of treatment or drugs by imagining two populations of cancer cells – one susceptible to treatment, the other not or less so (Panetta, 1998; Tomasetti and Levy, 2010; Sun et al., 2016). Other models see resistance as the cancer cells outcompeting the immune cells (De Pillis and Radunskaya, 2001). Baar et al. (2016) build a stochastic agent-based model where resistant phenotypes and genotypes are selected via mutation as they interact with a heterogeneous population of T-cells.

We consider an evolutionary game between two types of cancers cells in response to the immune system (Dhodapkar, 2013), either early in cancer progression or early in the application of immunotherapy. Rather than susceptible versus non-susceptible, we are interested in modeling two ways for how cancer cells evade killer T-cells. The two strategies available to the cancer cells correspond to the two major subsets of immune-resistant cancer cells proposed by Gajewski et al. (2013). The first represents a non-cooperative, selfish strategy where the cancer cell using this strategy enjoys some resistance to the immune system, but in a manner that has no direct influence on other cancer cells. Biologically, this models resistance strategies that involve the cancer cells down-regulating or ceasing to present the antigen required for the T-cell to recognize the cancer cell. This amounts to the cancer cells evolving camouflage and crypticity (Maeurer et al., 1996; Seliger et al., 1997; Hicklin et al., 1999; Johnsen et al., 1999). The second strategy amounts to cooperation as the cancer cells shut down the immune response as a collective public good. Biologically, this models resistance strategies where the cancer cell co-opts signaling pathways by directly signaling the T-cells to cease or by signaling macrophages and/or helper T-cells to cease supporting the production and entraining of killer T-cells. The collective consequence of this second strategy is immunosuppression within the tumor (Stewart and Smyth, 2011).

In what follows we develop a predator-prey-like population model that includes the two types of cancer cells and the T-cells. Like Babbs (2012) and Kareva and Berezovskaya (2015), we see the interaction between cancer cells and the T-cells as a modified predator-prey system. A selfish cell enjoys the protection provided by its cooperative counterparts while providing no benefits to other cancer cells. Cooperative cells work together to suppress the immune system. Their rate of survival increases with the ratio of cooperative cells. They gain no protection from the selfish cells. The

interplay of selfishness and cooperation represents a public goods game, selfishness being the individually optimal strategy, while cooperation being the socially optimal one. In our model, any given selfish cell is more likely to survive an encounter with an immune cell than a cooperative one. Since cooperative cells gain strength in numbers and selfish cells do not, a tumor comprised of cooperative cancer cells may be worse for the patient.

As specific goals, we examine when one or the other cancer strategy will outcompete the other. We then analyze the dynamics and equilibria of tumors comprised of cancer cells using either the selfish or cooperative strategy. Of considerable interest are extinction thresholds where below a critical size the immune system can drive the cancer extinct whereas above these thresholds the cancer will grow to sizes largely unaffected by the immune system. Using comparative statics, we consider the effects of various parameters corresponding to different forms of therapies on the model. Of specific interest are therapies that may unwittingly switch the stable steady state from selfish to cooperative cancer cells thus worsening the patient's prognosis (Smyth et al., 2006). Section 5.2 introduces the model. In Section 5.3 we discuss the steady states and the evolutionary stable strategies (ESS) of the model. In Section 5.4 we discuss the effects of different therapies. Section 5.5 concludes this chapter.

5.2 The model

We imagine a patient with a clinically detectable primary tumor that may or may not be metastatic. We assume that it is a solid tissue cancer, while not specifying the exact kind of cancer. The model considers the interaction between the cancer cells and the immune system within the tumor. For the moment we are imagining a population of killer T-cells. Our focus will be on the dynamics and steady-state population levels of both cancer and immune cells (Section 5.3). We then consider how various forms of therapy acting through the model's parameters alter tumor growth and prolong the progression free survival of the patient (Section 5.4).

We use a series of ordinary differential equations to model two cancer cell phenotypes, selfish and cooperative. Let $x_s(t)$ and $x_c(t)$ denote the population sizes of the selfish and cooperative types at time $t \in [0, \infty)$, respectively. Let $x(t) = x_s(t) + x_c(t)$ denote the total tumor mass, and $g(t) = x_c(t)/x(t)$ be the proportion of cooperative cells. Let $y(t)$ denote the total amount of killer T-cells (referred to as T-cells, for simplicity) in the tumor. We assume that the cells are well mixed, and so the ratio of the

cooperative phenotype in any cell's interaction radius is $g(t)$. We refer to the values of $x_s(t)$, $x_c(t)$, $x(t)$, and $y(t)$ as *population* counts. The value $g(t)$ is referred to as the *composition* of the tumor. We henceforth omit the time variable t , whenever it does not cause confusion. For a time varying $z(t)$ we use the usual notation \dot{z} to denote its time derivative.

The change in the number of both cancer cell populations and the total population is given by:

$$\dot{x}_c = r_c(1 - \sqrt{\frac{x}{K}})x_c - g\mu_c(g)\sqrt{xy}, \quad (5.1)$$

$$\dot{x}_s = r_s(1 - \sqrt{\frac{x}{K}})x_s - (1 - g)\mu_s(g)\sqrt{xy}, \quad (5.2)$$

$$\dot{x} = ((1 - g)r_s + gr_c)(1 - \sqrt{\frac{x}{K}})x - ((1 - g)\mu_s(g) + g\mu_c(g))\sqrt{xy}, \quad (5.3)$$

with $r_s, r_c, K \in \mathbb{R}_+$, and functions $\mu_c, \mu_s: [0, 1] \rightarrow \mathbb{R}_+$. For ecological realism we assume that $\dot{x}_c = \dot{x}_s = \dot{x} = 0$, whenever $x = 0$, i.e. the patient is cancer-free. As will be made clear, this is the continuous extension of the model defined by 5.1, 5.2, and 5.3.

The non-negative parameters r_s and r_c denote the growth rates of the selfish and the cooperative phenotypes, respectively. We assume that in the absence of immune reaction (or other exogenous factors, e.g. treatment), the population of cancer cells grows until it reaches its carrying capacity, K , upon which, the population stabilizes.

Standard logistic growth takes the form $1 - x/K$, indicating a linearly decreasing per capita growth rate with population size. Our root-logistic formulation of $1 - \sqrt{x/K}$ permits more tractable analytic results. Furthermore, below carrying capacity, per capita growth rates are lower than logistic, and maximum population growth occurs at $4K/9$ rather than at $K/2$. A peaking of total population growth rate at a smaller population size is consistent with a necrotic core in tumors. Cells proliferate at a slower rate in the interior of the tumor than at the exterior (Adam and Maggelakis, 1989; Foryś and Mokwa-Borkowska, 2005), due to a decreased resource availability. Our model therefore captures the increased resource scarcity, and the decreased resource uptake of a growing tumor, resulting in a slower-than-logistic growth function. In all other ways we follow the typical assumptions of Lotka-Volterra competition models.

Note that x is the sum of the population of selfish and cooperative phenotypes, and thus the reproduction rate of the tumor is the linear combination of the reproduction rates of the phenotypes using their frequencies as weights.




Figure 5.1: Per capita growth rates of logistic (red), and root-logistic (blue) models. Both assume decreasing per capita growth rates due to diminishing per capita resources in the environment, but root-logistic growth rate is lower. In case of cancer, the existence of a necrotic core implies that only a subset of the population participates in resource uptake, thereby lowering the growth rate of tumors.

The negative terms on the right-hand sides of (5.1), (5.2), and (5.3) capture the rate at which tumor cells have lethal encounters with T-cells. The encounter rate is given as \sqrt{xy} to keep the analysis tractable. Note that \sqrt{x} is proportional to the perimeter of a 2-dimensional object of area x , and in many cases, \sqrt{x} serves as an approximation for the surface area of a 3-dimensional object of volume (or mass) x . Therefore, our formulation can be interpreted as all encounters happening on (or in an area proportional to) the surface region of the tumor (Robertson-Tessi et al., 2012).

The function μ_c denotes the rate at which a cooperative cell is killed by a T-cell upon encounter, and the function μ_s denotes the rate with which a selfish cell is killed upon encounter. They are given by:

$$\mu_c(g) = \gamma(1 - g\varepsilon_c), \quad (5.4)$$

$$\mu_s(g) = \gamma(1 - g\varepsilon_c)(1 - \varepsilon_s), \quad (5.5)$$

with $\varepsilon_c, \varepsilon_s \in [0, 1]$.

The parameter γ denotes the lethality of the T-cells, while parameters ε_s and ε_c describe the resistance of the selfish and the cooperative phenotypes, respectively. The value $\mu_c(g)$ is the probability that a cooperative cancer cell, encountering an active T-cell, is destroyed, given composition g . The value $\mu_s(g)$ has the same interpretation for selfish cells. As in a public goods game, the cooperative cancer cells contribute a shared resistance that benefits all cancer cells equally, including the selfish cells. In addition to the public resistance provided by the cooperative cells, the selfish cells benefit themselves through their own resistance strategy. Every selfish cell benefits

from the resistance provided by their cooperative counterparts but not vice versa. As a result selfish cancer cells always face a lower death rate than cooperative cells. In (5.3) the linear combination of the phenotypical death rates $\mu_x(g) = (1 - g)\mu_s(g) + g\mu_c(g) = (1 - g)\gamma(1 - g\varepsilon_c)(1 - \varepsilon_s) + g\gamma(1 - g\varepsilon_c)$ is determined using the frequencies as weights.

The factor $1/(1 - \varepsilon_s) = \mu_c/\mu_s$ represents the advantage of selfishness. A value of 1, coinciding with zero selfish resistance $\varepsilon_s = 0$, means selfish cells survive an encounter with a T-cell with the same rate as cooperative cells. The higher the ε_s value, the higher the difference between the survival rates, favoring the selfish phenotype.

Notice that $\mu_c(g)$ and $\mu_s(g)$ are given as linear functions of g , and $g \in [0, 1]$. Hence, as the number of cancer cells x approaches zero, it can be inferred from (5.3) by L'Hôpital's rule that \dot{x} also converges to zero. Similarly, as x_c, x_s approach zero, \dot{x}_c and \dot{x}_s both converge to zero. Therefore, the assumption $\dot{x}_c = \dot{x}_s = \dot{x} = 0$ whenever $x = 0$ is a continuous extension of (5.1), (5.2), and (5.3).

Let $\omega \in \mathbb{R}^+$ denote the maximum tumor mass a patient can survive. Once the total cell count of cancer cells reaches ω , the patient dies. We assume $x(0) < \omega$, meaning that the patient is alive at time 0. Let $T = \inf\{t \geq 0: x(t) \geq \omega\}$ ² denote the patient's survival time.

The change in the population of T-cells is given by:

$$\dot{y} = r_i(1 - \frac{y}{L + a\sqrt{x}})y - \delta g\sqrt{x}y, \quad (5.6)$$

with $a, \delta, L \in \mathbb{R}_+$. In the absence of tumor cells, the population of T-cells grows logistically with a growth rate (or replenishment rate) r_i , and a carrying capacity L . If the patient is cancer-free, i.e. $x = 0$, we define $\dot{y} = r_i(1 - \frac{y}{L})y$.

As in De Angelis and Mesin (2001) and De Pillis and Radunskaya (2001), an increased tumor mass increases the immune response. This is modeled by a rise in the carrying capacity of T-cells. We assume that this increase is proportional to the encounter rate per T-cell, \sqrt{x} . The non-negative parameter a indicates the strength of the increase. Setting $a = 0$ means no change in immune response.

As a deviation from standard predator-prey models we assume that cooperative cancer cells inhibit the growth of the T-cell population. This captures the immunosuppression effect of a tumor, which in our model is induced only by the cooperative phenotype. The effect on the growth rate of the T-cells is proportional to the number of cooperative cancer cells encountered, $g\sqrt{x}$, factored by a constant δ that measures the effectiveness of immunosuppression. This linear formalization is also in line with

² $\inf \emptyset = \infty$.

De Angelis and Mesin (2001). For a detailed meta-study on the molecular mechanics of tumor-induced immunosuppression the reader is referred to Wu et al. (2015).

By rearranging (5.1) and (5.2) we can derive the evolution of g .

$$\dot{g} = \left(\frac{\dot{x}_c}{x} \right) = \frac{\dot{x}_c x - \dot{x} x_c}{x^2} = g(1 - g) \left((r_c - r_s) \left(1 - \sqrt{\frac{x}{K}} \right) + (g\varepsilon_c \varepsilon_s - \varepsilon_s) \frac{y}{\sqrt{x}} \gamma \right). \quad (5.7)$$

In (5.7) it is stated that the frequency of cooperative cells evolves autonomously as a result of differences of birth and death rates of the two phenotypes. The factors $1 - \sqrt{\frac{x}{K}}$, and $\gamma \frac{y}{\sqrt{x}}$ express the relative importance of birth and death rates in determining which phenotype proliferates at a higher rate. For instance, in case of a low T-cell population, the difference between r_c and r_s determines whether the ratio of the cooperative phenotype increases or decreases. With a high T-cell population, the phenotype with the higher resistance prevails.

As a technical remark, note that for $x = 0$, the composition parameter g and its time-derivative are undefined. As a consequence, with x approaching zero, the time-derivative of g can be unbounded in magnitude. Nevertheless, the trajectories remain continuous, and hence g will remain between 0 and 1 as per Equation 5.7.

The system defined by (5.3), (5.6), and (5.7) is equivalent to the one defined by (5.1), (5.2), and (5.6).

5.3 Steady-state analysis

In this section we focus our attention on the steady states of the model. Steady states are combinations of tumor population size, tumor composition, and T-cell population size for which all three components of the model are not changing in time. A brief summary of our findings is as follows: No steady states exist with both cancer cells types (cooperative and selfish) at positive population sizes. The steady state with no cancer cells and no immune cells is always unstable. The steady state with no cancer cells but the presence of immune cells may or may not be stable. Finally, a variety of steady states exist with positive populations of one cancer cell type and immune cells. Their stability properties can be illustrated using zero-growth isocline diagrams. We start with the condition that determines which cancer cell type outcompetes the other. The sign of the following expression determines whether a given cancer cell type can both invade and resist invasion from the other type:

$$\frac{r_c}{r_s} - \frac{1}{1 - \varepsilon_s}.$$

If the ratio of reproduction rates is larger than the advantage of selfishness, then the tumor (unless eliminated by the T-cells) is cooperative at the ESS, otherwise, it is selfish. Since a purely cooperative tumor is more aggressive than a purely selfish one thanks to the former's ability to suppress immune responses, the composition of the tumor, and therefore the sign of the above expression is crucial for the patient's prognosis.

The remainder of this section contains the formal definitions, statements, and proofs.

Definition 5.3.1. The triple (x^*, y^*, g^*) is called a steady state of the dynamic system defined by (5.3), (5.6), and (5.7), if $\dot{x} = 0$, $\dot{y} = 0$, and $\dot{g} = 0$ are all satisfied.

A steady state tumor described by (x^*, y^*, g^*) is called mixed if $g^* \in (0, 1)$, cooperative if $g^* = 1$, and selfish if $g^* = 0$. Cooperative and selfish tumors comprise an ESS if they are linearly stable for deviations in x and y , and if they resist an invasion of the other phenotype. Since our model precludes mixed steady states, we only discuss stability for homogeneous tumors.

Definition 5.3.2. Fix g . The set of (x, y) pairs for which $\dot{x} = 0$ is called the zero-growth isocline curve of x . The set of (x, y) pairs for which $\dot{y} = 0$ is called the zero-growth isocline curve of y .

For $t \geq 0$ the triplet $(x(t), y(t), g(t))$ can be thought of as a snapshot of the tumor and the immune system at time t , specifying the tumor size, the immune cell count, and the composition of the tumor. As per Definition 5.3.1, in a steady state, both population counts and the tumor composition are constant.

Clearly, if (x^*, y^*, g^*) is a steady state, then (x^*, y^*) is located along the zero-growth isocline curves of both x and y , given g^* . The system yields three trivial steady states, as listed in the following remark. For ecological realism we include the steady-states with $x = 0$, even as the state parameter g is undefined in this case.

Remark 5.3.3. The following triples are all steady states.

1. $x = 0, y = L$.
2. $x = 0, y = 0$.
3. $x = K, y = 0, g \in [0, 1]$.

In Remark 5.3.3, Case 1 describes a cancer-free patient with T-cells matching the body's carrying capacity. This patient is healthy. Case 2 describes a patient who is

cancer-free, but has no T-cells. Cases 1 and 2 are undefined in the model, but their stability properties are well-defined. This may be interpreted as an unrelated immune deficiency. This patient is cancer-free but is unprotected against diseases due to a lack of T-cells. Since $\omega < K$, case 3 describes a patient who has already succumbed to cancer, and has no immune protection.

From (5.7) one can deduce that both $g = 0$ and $g = 1$ guarantee the compositional stability of the tumor, meaning that $\dot{g} = 0$ is assured. We subdivide this section into three subsections: we briefly show non-existence of steady states in mixed tumors before going on to discuss steady states in selfish and cooperative tumors, respectively.

5.3.1 Nonexistence of steady states in mixed tumors

In this subsection we show that unless the advantage of selfishness happens to be the same as the ratio of reproduction rates, then two steady state conditions, the size of the tumor being constant, and its composition being constant can only be satisfied for purely selfish or purely cooperative tumors.

Proposition 5.3.4. *Suppose that $0 < g < 1$, and $0 < x < K$.*

- *If we have $\frac{r_c}{r_s} > \frac{1}{1-\varepsilon_s}$ and $\dot{x} = 0$, then it holds that $\dot{g} > 0$.*
- *If we have $\frac{r_c}{r_s} < \frac{1}{1-\varepsilon_s}$ and $\dot{x} = 0$, then it holds that $\dot{g} < 0$.*

Proof. First we consider $r_c(1 - \varepsilon_s) > r_s$. Since $\dot{x} = 0$, we must have $\dot{x}_c = -\dot{x}_s$. If $\dot{x}_c > 0$ and $\dot{x}_s < 0$, then $\dot{g} > 0$ holds, since the cooperative population is rising and the selfish one dwindling. Suppose that $\dot{x}_s \geq 0$. Then we have

$$0 \leq (1 - g)r_s(1 - \sqrt{\frac{x}{K}}) - (1 - g)\gamma\frac{y}{\sqrt{x}}(1 - g\varepsilon_c)(1 - \varepsilon_s).$$

We multiply by $g/(1 - g)$, divide by $(1 - \varepsilon_s)$ and invoke $r_c(1 - \varepsilon_s) > r_s$ to get

$$0 \leq g\frac{r_s}{1 - \varepsilon_s}(1 - \sqrt{\frac{x}{K}}) - g\gamma\frac{y}{\sqrt{x}}(1 - g\varepsilon_c) < gr_c(1 - \sqrt{\frac{x}{K}}) - g\gamma\frac{y}{\sqrt{x}}(1 - g\varepsilon_c) = \dot{x}_c,$$

meaning that $\dot{x}_c > 0$, contradicting $\dot{x} = 0$, since both population counts are rising. Therefore, $\dot{x}_s < 0$ and $\dot{x}_c > 0$ must hold, implying $\dot{g} > 0$.

The case $r_c(1 - \varepsilon_s) < r_s$ follows via similar arguments. ■

Proposition 5.3.4 states that constant tumor size cannot coexist with constant tumor composition. Moreover, if the ratio of the reproduction rates r_c/r_s is higher than the advantage of selfishness, $1/(1 - \varepsilon_s)$, then constant tumor size implies higher relative

proliferation of the cooperative cells. If the ratio of reproduction rates is lower, the opposite holds, and the selfish cells will outcompete the cooperative cells. Therefore, no non-trivial steady states exist in which the tumor is a mix of the two phenotypes (the knife-edge case $r_c/r_s = 1/(1 - \varepsilon_s)$ is not considered).

5.3.2 Steady states in selfish tumors

We now consider selfish tumors. We show that steady states in this case correspond to the solution set of a second-degree polynomial in \sqrt{x} . If there are two solutions, the one with a larger tumor population is stable if and only if the tumor population exceeds a quarter of the carrying capacity. The steady state with the lower tumor cell population is unstable and functions as a strong Allee threshold (Taylor and Hastings, 2005) below which tumor growth is negative.

The zero-growth isoclines of the selfish case can be calculated by setting $g = 0$, $\dot{x} = 0$, and $\dot{y} = 0$, and are as follows: The zero-growth isocline of y is

$$y = L + a\sqrt{x}, \quad (5.8)$$

while the zero-growth isocline of x is

$$y = \sqrt{x}\left(1 - \sqrt{\frac{x}{K}}\right)\frac{r_s}{\gamma(1 - \varepsilon_s)}. \quad (5.9)$$

This allows us to calculate the steady states of the selfish case.

Proposition 5.3.5. *The triple $(x^*, y^*, 0)$ is a steady state, if x^* solves the following equation*

$$0 = -\frac{r_s}{\sqrt{K}}x^* + \sqrt{x^*}(r_s - a\gamma(1 - \varepsilon_s)) - \gamma(1 - \varepsilon_s)L$$

and $y^* = L + a\sqrt{x^*}$.

Proof. Setting the two right hand sides of (5.8) and (5.9) equal and rearranging gives

$$0 = -\frac{r_s}{\sqrt{K}}x^* + \sqrt{x^*}(r_s - a\gamma(1 - \varepsilon_s)) - \gamma(1 - \varepsilon_s)L.$$

■

Let

$$D_s = (r_s - a\gamma(1 - \varepsilon_s))^2 - 4\frac{L}{\sqrt{K}}r_s\gamma(1 - \varepsilon_s)$$

denote the value of the discriminant. Then, the exact values of the steady states are as follows:

- if $D_s > 0$, there are two steady states, $(x_{s1}^*, y_{s1}^*, 0)$, and $(x_{s2}^*, y_{s2}^*, 0)$, with

$$\sqrt{x_{s1,s2}^*} = \frac{(r_s - a\gamma(1 - \varepsilon_s)) \pm \sqrt{D_s}}{2\frac{r_s}{\sqrt{K}}}$$

and $y_{s1,s2}^* = L + a\sqrt{x_{s1,s2}^*}$, as long as $x_{s1,s2}^* > 0$.

- if $D_s = 0$, there is a single steady state, $(x_{s0}^*, y_{s0}^*, 0)$, where $x_{s0}^* = (r_s - a\gamma(1 - \varepsilon_s))\sqrt{K}/2r_s$, and $y_{s0}^* = L + a\sqrt{x_{s0}^*}$, as long as $x_{s0}^* > 0$.
- if $D_s < 0$ (or if $x_{s1,s2}^* < 0$), there are no non-trivial steady states in selfish tumors.

The zero-growth isoclines intersect only if $D_s \geq 0$ is satisfied, with two points of intersection if $D_s > 0$, and one point of intersection if $D_s = 0$.

Proposition 5.3.6. *Let $D_s > 0$. If $x_{s1}^* > \frac{K}{4}$, then $(x_{s1}^*, y_{s1}^*, 0)$ is stable in x and y .*

The proof is shown as an appendix. It relies on showing that the elements of the problem's Jacobian matrix take the following signs:

$$J(x_{s1}^*, y_{s1}^*) = \begin{pmatrix} - & - \\ + & - \end{pmatrix},$$

with the negativity of the top left element hinging on the condition $x_{s1}^* > K/4$. Graphically, this condition means that the larger equilibrium is to the right of the peak of the x zero-growth isocline. Thus, the trace of the Jacobian is negative, and the determinant positive, meaning that the steady state is stable. As long as the advantage of selfishness is larger than the ratio of phenotypical reproduction rates, this is an ESS.

Corollary 5.3.7. *Let $D_s > 0$. If $x_{s1}^* > K/4$, and if $r_c/r_s < 1/(1 - \varepsilon_s)$ then $(x_{s1}^*, y_{s1}^*, 0)$ is an ESS.*

Proof. By Proposition 5.3.6 we have linear stability in x and y , and since the advantage of selfishness is higher than the ratio of reproduction rates, Proposition 5.3.4 ensures that the selfish phenotype resists invasion by the cooperative one. ■

We illustrate the steady states in the selfish case, as well as some typical dynamic properties in the following example.

Example 5.3.8. Consider a tumor of selfish cells and the following parameter values: $K = 625$, $L = 20$, $r_s = 0.2$, $\gamma = 0.1$, $\varepsilon_s = 0.7$. The interpretation is the following. The carrying capacity of tumor cells is significantly higher than that of T-cells. The unconstrained rate of proliferation is 20%. In 10% of encounters, T-cells become activated,

and the cancer cells resist destruction with a rate of 70% once a T-cell is activated, for an effective kill rate of 3% per encounter. These parameters satisfy $D_s > 0$, hence we have a total of five steady states. Fix $a = 1$ and $\omega = 470$.

Figure 5.2 shows the zero-growth isoclines, phase diagram, and steady states of this system. Black dots denote the trivial steady states of Remark 5.3.3, $(0, 0)$, $(0, L)$, $(K, 0)$, and red ones denote the steady states where both the tumor size and the T-cell count is non-zero from Proposition 5.3.5. In Figure 5.3 we plot the growth of the tumor




Figure 5.2: Phase diagram of a selfish tumor with a positive discriminant. The steady state with a lower cancer cell population is unstable, and serves as an Allee threshold. The steady state with a higher cancer cell population is stable and is below the lethal tumor mass. Patient could survive with the cancer burden indefinitely.

population as a function of population size, moving along the y zero-growth isocline, a smooth function of x . Notice that the tumor's growth rate along the zero-growth isocline is maximized near a cancer cell population size of 150, or 24% of carrying capacity. The change of population, compared to distances between steady states, is small, a useful property for discretization of the model in conducting simulations. By Proposition 5.3.5, we have the two steady states, $x_{s1}^* = 281.6$, $y_{s1}^* = 36.78$ and $x_{s2}^* = 20$, $y_{s2}^* = 24.47$. The interpretation is the following. At the theoretical maximum population count, the cancer cells are at 45% carrying capacity, while the T-cells are at 184% of L , meaning that the immune reaction to the tumor burden is significant. At the Allee threshold, the tumor burden is 3.2% of carrying capacity, whereas T-cells are at 122% of L .

Below a tumor burden of 20, a strong immune system overcomes and eliminates the cancer despite its relatively large growth potential. Above the Allee threshold, the reproduction of cancer cells offsets even a strong immune response. Tumor growth is maximized for values of about 150, and stops at the theoretical maximum of 281.6.




Figure 5.3: The growth rate of a selfish tumor along the y zero-growth isocline. x_{s2}^* serves as a strong Allee threshold. Tumor population stabilizes at x_{s1}^* . Tumor growth is maximized at the quarter of the tumor's carrying capacity.

The model has two reasons for this absence of growth. First, a larger tumor induces a larger immune reaction that directly hinders the proliferation of cancer cells, and second, a larger tumor will face a lower abundance, and a lower uptake of resources that the cells use to reproduce.

Note that if for any reason (i.e. hidden tumors, or immune deficiency) a smaller tumor burden can manage to exceed the Allee threshold, reversing tumor growth becomes harder and harder as time passes. Another implication of this example is that if a successful treatment eradicates a large fraction of the tumor, but not enough to go below the Allee threshold, then in time the cancer will return.

Example 5.3.9. Consider the same parameter values as example 5.3.8, but with $\gamma = 0.2$ instead of 0.1, corresponding to a stronger immune system. Then $D_s < 0$, and the zero-growth isocline of y lies above the zero-growth isocline of x . Figure 5.4 shows the appropriate phase diagram. As per Proposition 5.3.5, there are no non-trivial steady states. Note that the steady state, $(0, L)$ behaves as an attractor for every initial condition with $y(0) > 0$, meaning that the cancer is always eradicated (unless the lethal tumor mass is reached before that happens).

5.3.3 Cooperative steady states

We now examine tumors composed of cooperative cancer cells. As before, the set of steady states corresponds to the solution set of a second degree polynomial in \sqrt{x} , and




Figure 5.4: Phase diagram of a selfish tumor with a negative discriminant. The zero-growth isoclines do not intersect. A healthy immune system drives the cancer cells to extinction. Patient is expected to become cancer free.

in case of two cooperative steady states the one with a higher cancer cell population count is stable if the number of cancer cells exceed $K/4$. Similarly to the selfish case, we calculate the zero-growth isocline curves by setting (5.3) and (5.6) to 0 and substituting $g = 1$. The zero-growth isocline of y in the cooperative case is

$$\dot{y} = 0 \Leftrightarrow y = (L + a\sqrt{x})(1 - \frac{\delta}{r_i}\sqrt{x}) = L + a\sqrt{x}(a - \frac{\delta L}{r_i}) - a\frac{\delta}{r_i}x, \quad (5.10)$$

while the zero-growth isocline of x is

$$\dot{x} = 0 \Leftrightarrow y = \sqrt{x}(1 - \sqrt{\frac{x}{K}})\frac{r_c}{\gamma(1 - \varepsilon_c)}. \quad (5.11)$$

Proposition 5.3.10. *The triple $(x^*, y^*, 1)$ is a steady state, if x^* solves the following equation:*

$$0 = (-\frac{r_c}{\sqrt{K}} + a\gamma(1 - \varepsilon_c)\frac{\delta}{r_i})x^* + (r_c - \gamma(1 - \varepsilon_c)(a - \frac{\delta}{r_i}L))\sqrt{x^*} - \gamma(1 - \varepsilon_c)L,$$

$$\text{and } y^* = \sqrt{x^*}(1 - \sqrt{\frac{x^*}{K}})\frac{r_c}{\gamma(1 - \varepsilon_c)}.$$

Proof. Setting the right hand sides equal and rearranging gives

$$\frac{\dot{x}}{\sqrt{x}} = (-\frac{r_c}{\sqrt{K}} + a\gamma(1 - \varepsilon_c)\frac{\delta}{r_i})x + (r_c - \gamma(1 - \varepsilon_c)(a - \frac{\delta}{r_i}L))\sqrt{x} - \gamma(1 - \varepsilon_c)L,$$

hence the roots of the above expression are the steady-state tumor masses. ■

Let

$$D_c = (r_c - \gamma(1 - \varepsilon_c)(a - \frac{\delta}{r_i}L))^2 - 4\frac{L}{\sqrt{K}}r_c\gamma(1 - \varepsilon_c) + 4a\gamma^2(1 - \varepsilon_c)^2\frac{\delta}{r_i}L$$

denote the value of the discriminant in the cooperative case. The exact values of the steady states are as follows:

- if $D_c > 0$, there are two steady states, $(x_{c1}^*, y_{c2}^*, 1)$, and $(x_{c2}^*, y_{c2}^*, 1)$, where

$$\sqrt{x_{c1,c2}^*} = \frac{(r_c - \gamma(1 - \varepsilon_c)(a - \frac{\delta}{r_i}L)) \pm \sqrt{D_c}}{2\frac{r_c}{\sqrt{K}} - 2\gamma(1 - \varepsilon_c)\frac{\delta}{r_i}},$$

and $y_{c1,c2}^* = (L + a\sqrt{x_{c1,c2}^*})(1 - \frac{\delta}{r_i}\sqrt{x_{c1,c2}^*})$, provided that y_{c1}^* and y_{c2}^* are positive.

- if $D_c = 0$, there is a single steady state, $(x_{c0}^*, y_{c0}^*, 1)$, for which we have $\sqrt{x_{c0}^*} = \frac{(r_c - \gamma(1 - \varepsilon_c)(a - \frac{\delta}{r_i}L))}{2\frac{r_c}{\sqrt{K}} - 2\gamma(1 - \varepsilon_c)\frac{\delta}{r_i}}$, and $y_{c0}^* = (L + a\sqrt{x_{c0}^*})(1 - \frac{\delta}{r_i}\sqrt{x_{c0}^*})$, provided that y_{c0}^* is positive.
- if $D_c < 0$, there are no steady states.

By analyzing the zero-growth isoclines in both cases one can see that the main difference between the purely selfish and the purely cooperative case comes from the suppression of T-cells by the cooperative phenotype. Notice that the T-cell zero-growth isocline includes a factor of $1 - (\delta/r_i)\sqrt{x}$. This factor is increasing in T-cell reproduction rate r_i and decreasing in the suppression parameter δ , as well as the encounter rate per T-cell \sqrt{x} . Hence a positive δ leads to a diminished T-cell population as a direct result of suppression.

Note however, that immunosuppression also indirectly raises the T-cell population. A lower T-cell population permits a higher steady-state tumor size, which in turn raises the T-cell population via an increased carrying capacity of T-cells. This induces a higher effective replenishment rate of the T-cell population.

We now state the conditions of a cooperative stable steady state.

Proposition 5.3.11. *Let $D_c > 0$. If $x_{c1}^* > K/4$, then $(x_{c1}^*, y_{c1}^*, 1)$ is linearly stable in x and y .*

The proof is shown as an appendix. The negativity of the Jacobian's trace is assured by the $x_{c1}^* > K/4$ condition, whereas the positivity of the determinant is due to the fact that the T-cell zero-growth isocline intersects the cancer's zero-growth isocline from below. If the advantage of selfishness is lower than the ratio of phenotypical reproduction rates, this steady state is also an ESS.

Corollary 5.3.12. *Let $D_c > 0$. If $x_{s1}^* > K/4$, and if $r_c/r_s > 1/(1 - \varepsilon_s)$ then $(x_{c1}^*, y_{c1}^*, 1)$ is an ESS.*

Proof. By Proposition 5.3.11 we have linear stability in x and y , and since the advantage of selfishness is lower than the ratio of reproduction rates, Proposition 5.3.4 ensures that the cooperative phenotype resists an invasion by the selfish phenotype. ■

Together, Corollaries 5.3.7 and 5.3.12 mean that there is always a unique ESS of the system, and the phenotype therein is determined by the advantage of selfishness.

Example 5.3.13. Consider a tumor with only cooperative cells and the following parameter values: $K = 625$, $L = 20$, $r_c = 0.2$, $r_i = 0.1$, $\gamma = 0.1$, $\varepsilon_c = 0.7$, $a = 1$. We retain the values for the reproduction rate and resistance parameter of example 5.3.8 so that we can better demonstrate the effect of immunosuppression. Figure 5.5 shows the isocline diagram, plotted for $\delta = 0.003$. The interpretation is that T-cells are killed with a rate of 0.3% per encounter. In this example, the cancer cell population and the T-cell population act as competitors, as both zero-growth isoclines have negative slopes at the stable non-trivial steady-state. Notice that the cancer's zero-growth isocline is




Figure 5.5: Phase diagram of a cooperative tumor with two non-trivial steady states. The steady state with lower cancer cell population is unstable, and serves as an Allee threshold. The steady state with higher cancer cell population is stable, and is above the lethal tumor mass. Patient prognosis is poor.

unaffected by the change in the suppression parameter and remains a concave curve, as in example 5.3.8. However, with $\delta > 0$, the zero-growth isocline for y has an upward sloping part and becomes downward sloping. Two non-trivial steady states exist.

In Figure 5.6 we plot the same diagram with $\delta = 0.005$, corresponding to a stronger immunosuppression ability of the cooperative phenotype (for this calibration, the zero-growth isocline of y reduces to a line). In this case, the steady state (x_{c1}^*, y_{c1}^*) is out of bounds, since $x_{c1}^* > K$ and $y_{c1}^* < 0$, meaning that we only get one non-trivial steady

state, (x_{c2}^*, y_{c2}^*) , a saddle point, leaving $(0, L)$ and $(K, 0)$ as stable steady states. A complete suppression of the immune system is therefore possible.




Figure 5.6: Phase diagram of a cooperative tumor with a single non-trivial steady state. The zero-growth isoclines intersect in the positive quadrant only once. A complete suppression of the immune system becomes possible. Patient prognosis is grave.

5.4 Treatment effects

The principal objectives behind cancer modeling are to improve and inform the research into therapy. Here we explore the implications of our model for cancer therapy strategies. We do this by investigating how a change in the tumor's micro-environment affects the tumor and the immune system and then draw conclusions on their influence on the patient's prognosis.

Patient outcome in our model can be measured as survival time and/or as the tumor's steady-state size. Our survival time measure relates to what is known as progression free survival time following the initiation of therapy. This can be actual time before patient death, or it may represent the time before another line of therapy becomes necessary. The steady-state size of the primary tumor may be of interest as this may correlate with the likelihood or rate of metastases. We can use our model to examine the nature and effectiveness of various forms of immunotherapy. Specifically, we are interested in how changes in model parameter values and outcomes can be interpreted from a therapeutic standpoint. Immunotherapy likely influences three key parameters of the model:

1. *Increasing γ , the rate at which cancer cells have lethal encounters with T-cells.*

In immunotherapy this is often accomplished by using retroviruses or vaccines that

tag the cancer cells with an antigen that can be detected by the patient's immune system, thus eliciting the production of antigen specific T-cells. Alternatively, immune cells from the patient may be entrained and cultured *in vitro* to recognize cancer cells collected through biopsy. These immune cells are then injected into the patient.

2. *Decreasing δ , the suppression parameter of the cooperative phenotype.* Targeted therapy can be used to lower the cooperative phenotype's ability to suppress the immune system. Such therapy by neutralizing the cancer cells' signaling or by killing cells that overexpress PDL-1, either reduce δ , or strongly select against cancer cells with high values of δ .

3. *Decreasing ε_c and ε_s , the resistance rates of the two phenotypes.* The T-cells kill cancer cells through direct contact, and the injection of proteins that initiate apoptosis (cell death). Cancer cells may exaggerate glycolysis to produce lactic acid. The secreted lactic acid by reducing Ph creates a protective moat that impairs contact with an encountered T-cell. This is often associated with the upregulation of carbonic dehydratase IX (CAIX) that assists the cancer cell in maintaining and surviving an acidic environment. Buffer therapies can be added to the immunotherapy as a means of raising the Ph of the tumor and rendering the cancer cells' acidic moat ineffective to lower ε_c (Ibrahim-Hashim et al., 2017). Checkpoint inhibitors may unmask the selfish phenotype, making the selfish phenotype more susceptible to destruction by the T-cells (Pardoll, 2012).

With such therapies in mind, in this section we consider the effects of these key parameters on the trajectory of tumor growth and the steady-state size of the tumor. We do this by using comparative statics and simulations on a discretized model, using finite difference approximation. Our conclusions are as follows.

1. Increased immune system efficiency and/or decreased cooperative resistance prolongs the patient's survival time and lowers tumor mass at the stable steady state.
2. Decreased immunosuppression prolongs the patient's survival time in tumors that feature the cooperative phenotype. It decreases tumor mass at the stable steady state, if it is cooperative, and has no effect on the stable steady state, if it is selfish.
3. Decreased crypticity has ambiguous effects. It may increase or decrease the patient's survival time and increase or decrease tumor mass at the stable steady state.

The first two findings are in line with expectations, while the third may seem counter-intuitive. The reason for this is that reduced crypticity lowers the advantage of selfishness, and thus pushes the cancer cells towards an ESS with cooperative cancer cells. If the tumor becomes cooperative instead of selfish at the stable steady state, the patient's prognosis may be worse due to the cooperative phenotype's immunosuppression. Treatments that increase the immune system's efficiency, decrease the cooperative resistance, or decrease immunosuppression have no effect on the advantage of selfishness. Hence, these treatment strategies do not influence the tumor's steady-state composition and there is no ambiguity in their benefits on the patient's prognosis.

The remainder of this section formalizes, extends, and illustrates the above results.

5.4.1 Effective treatment strategies

Any treatment that does not change the advantage of selfishness can be shown to be effective. We first formalize the effects of the parameters γ and ε_c on the steady-state tumor size.

Proposition 5.4.1.

1. If D_c is positive then, for every parametrization, the following comparative statics hold:

- a. $\frac{\partial x_{c1}^*}{\partial \gamma} < 0$, $\frac{\partial x_{c2}^*}{\partial \gamma} > 0$,
- b. $\frac{\partial x_{c1}^*}{\partial \varepsilon_c} > 0$, $\frac{\partial x_{c2}^*}{\partial \varepsilon_c} < 0$.

2. If D_s is positive then, for every parametrization, the following comparative statics hold:

- a. $\frac{\partial x_{s1}^*}{\partial \gamma} < 0$, $\frac{\partial x_{s2}^*}{\partial \gamma} > 0$,
- b. $\frac{\partial x_{s1}^*}{\partial \varepsilon_c} = 0$, $\frac{\partial x_{s2}^*}{\partial \varepsilon_c} = 0$.

Proof. 1.a. Consider the zero-growth isoclines of the cooperative case, (5.10) and (5.11). Since $D_c > 0$ we have two intersection points, and since the zero-growth isocline of x includes the origin, while that of y has a positive y -intercept, it must hold that at x_{c2}^* the zero-growth isocline of y intersects that of x from above, while at x_{c1}^* the zero-growth isocline of x intersects that of y from below. Raising γ affects only the zero-growth isocline of x , compressing it towards the x -axis. It follows that the points of intersection must move closer together, hence tumor mass decreases at the stable steady state, and the Allee threshold is achieved at a higher tumor mass.

b. Decreasing ε_c has the same effect as increasing γ .

2.a. This case is identical to 1.a, but instead of the isoclines of the cooperative case we need to use (5.8) and (5.9).

b. Changing ε_c has no effect on the zero-growth isoclines in the selfish case, hence the selfish steady states do not change. ■

Proposition 5.4.1 confirms the first point outlined at the beginning of this section. Treating for immune system efficiency and/or cooperative resistance decreases tumor mass at the stable steady state, making metastases less likely. In addition, this treatment increases the Allee threshold, thereby making the cancer population more vulnerable to other forms of treatment (Tobin et al., 2011). Figure 5.7 shows the qualitative effects of an increased immune efficiency parameter on the two cooperative steady states with original parameters taken from Example 5.3.13. The dashed red line shows the original cancer zero-growth isocline, the solid red line shows the zero-growth isocline after the increase. Note that in this example the treatment results in the cancer cell population at the stable steady state being lower than the critical tumor mass, hence a patient may survive indefinitely with the tumor burden due to this therapy.




Figure 5.7: The effect of an increased γ on cooperative steady states. The cancer zero-growth isocline is compressed towards the horizontal axis. Tumor mass becomes higher at the Allee threshold, and lower at the stable steady state. Steady-state tumor mass falls below the lethal value. Patient prognosis improves as a result of the treatment.

Next, we formulate the relationship between the suppression parameter and steady-state tumor sizes for cooperative tumors.

Proposition 5.4.2. *If D_c is positive then, for every parametrization, the following comparative statics hold: $\frac{\partial x_{c1}^*}{\partial \delta} > 0$, $\frac{\partial x_{c2}^*}{\partial \delta} < 0$.*

Proof. We again consider the two zero-growth isoclines of the cooperative case, (5.10) and (5.11). Raising δ has no effect on the cancer zero-growth isocline, while it decreases the slope of the T-cell isocline. Hence the points of intersection of the two curves must spread farther apart, implying the stated inequalities. ■

Proposition 5.4.2 confirms the second point highlighted in the beginning of this section. It implies that lowering the immunosuppression parameter via treatment has the same qualitative effect on steady states in cooperative tumors as strengthening the immune system. Clearly, such a treatment has no effect on a purely selfish tumor as the selfish phenotype does not display immunosuppression. Figure 5.8 shows the qualitative effects of a decreased immunosuppression parameter on the two cooperative steady states with original parameters retained from Example 5.3.13. The dashed blue line shows the original T-cell zero-growth isocline, the solid red line shows the zero-growth isocline after the treatment. Once again, this example shows a treatment that results in the cancer population at the stable steady state being lower than the critical tumor mass. We now consider the effects of therapy on the survival time of the




Figure 5.8: The effect of an decreased δ on cooperative steady states. The T-cell zero-growth isocline is tilted upward. Tumor mass becomes higher at the Allee threshold, and lower at the stable steady state. Steady-state tumor mass falls below the lethal value. Patient prognosis improves as the result of the treatment.

patient, T , a variable determining the viability of a treatment strategy. We develop a round of simulations in a discrete model to obtain approximations of the effects of treatments via γ and in δ on survival time. Let x_c^t, x_s^t, y^t denote cooperative, selfish, and T-cell populations at time $t = 0, 1, 2, \dots$ for a given initial state (x_c^0, x_s^0, y^0) . For $t \geq 0$ the dynamics of the state-variables were determined by finite difference approximations of equations (5.1), (5.2), (5.6) with time step 1. The differences between values of time

t and $t - 1$, $\Delta x_c^t = x_c^t - x_c^{t-1}$, $\Delta x_s^t = x_s^t - x_s^{t-1}$, and $\Delta y^t = y^t - y^{t-1}$ are given by the right hand sides of the respective equations, substituting the previous values x_c^{t-1} , x_s^{t-1} , and y^{t-1} . The time step 1 was chosen for intuitive simplicity, and because the differences are of a suitably small magnitude to ensure smooth dynamics (see e.g. Figure 5.3).

The configurations used are as follows: $K = 625$, $L = 20$, $r_s = 0.19$, $r_c = 0.2$, $\varepsilon_s = 0.7$, $\varepsilon_c = 1$, $a = 1$. We considered three possible critical tumor masses $\omega_1 = 150$, $\omega_2 = 200$, $\omega_3 = 250$. The initial conditions of the system are characterized by $x(0) = 40$, $y(0) = 20$, $g(0) = 0.7$. Note that recalibrating the model with different parameters and initial conditions does not change the qualitative effects but may lead to an early extinction of the tumor, making survival times infinite. For different combinations of γ and δ , Table 5.1 gives the times T_1 , T_2 , and T_3 at which the tumor mass exceeds the critical tumor masses ω_1 , ω_2 , and ω_3 , respectively, i.e. $T_i = \inf\{t > 0: x(t) > \omega_i\}$. The trends in each table are clear, for all three critical tumor masses, increasing γ or

	T_1	γ					T_2	γ			
		0.025	0.05	0.075	0.1			0.025	0.05	0.075	0.1
δ	0.001	13	15	20	41	δ	0.001	17	20	27	58
	0.003	13	15	19	31		0.003	16	19	25	48
	0.005	13	15	18	26		0.005	16	19	23	37
	0.007	13	15	17	23		0.007	16	18	22	31
	0.009	13	14	17	21		0.009	16	18	21	27

	T_3	γ			
		0.025	0.05	0.075	0.1
δ	0.001	20	24	36	90
	0.003	20	23	31	76
	0.005	20	23	28	54
	0.007	19	22	26	39
	0.009	19	22	25	33

Table 5.1: Survival times for different combinations of γ and δ . Patient prognosis improves with increase in γ and decrease in δ .

decreasing δ weakly increases survival time. Clearly, patients live the longest when the immune system is strongest and immunosuppression is weakest, corresponding to

the upper right corners of each table. Conversely, when the immune system is weak and immunosuppression is the strong, life expectancy is lowest, corresponding to the bottom left corners of the tables. Furthermore, notice that a stronger immune system means a stronger effect of immunosuppression. All of these observations predicted by our model fall in line with basic intuitions about immunotherapy.

5.4.2 Ambiguous effects of treating against cancer crypticity

Treating for the selfish resistance rate – the crypticity parameter – may have ambiguous effects on the patient. Specifically, we first show that reducing the resistance rate may both increase or decrease the tumor population at the stable steady state. On first glance this may seem counter-intuitive as it is easy to show that selfish steady states react the same way to a treatment of ε_s as cooperative ones to that of ε_c .

Proposition 5.4.3. *If D_s is positive then, for every parametrization, the following comparative statics hold: $\partial x_{c1}^*/\partial \varepsilon_s > 0$, $\partial x_{c2}^*/\partial \varepsilon_s < 0$.*

Proof. Very similar to Proposition 5.4.1 1.b. ■

Proposition 5.4.3 indicates that marginal changes of the selfish resistance parameter also constitute an effective treatment, as they have no effect on steady-state composition. However, if due to a major change of ε_s , the advantage of selfishness becomes lower than the ratio of the reproduction rates, r_c/r_s , then the system's ESS may become cooperative instead of selfish. This may worsen the patient's prognosis. We provide an example to illustrate this.

Example 5.4.4. Consider the following calibration: $K = 625$, $L = 20$, $r_c = 0.2$, $r_s = 0.19$, $\varepsilon_c = \varepsilon_s = 0.7$, $\delta = 0.001$, $\gamma = 0.1$. In this case, $r_c(1 - \varepsilon_s) = 0.06 < 0.19 = r_s$, so the advantage of selfishness is substantial. The system's only ESS is $(x_{s1}^*, y_{s1}^*, 0)$ where the tumor size is $x_{s1}^* = 261.5$. At the steady state, the cancer cells are at 42% of carrying capacity, while the T-cell population is 36.17, 181% of carrying capacity.

Suppose that the selfish phenotype's defenses are fully eliminated, i.e. $\varepsilon'_s = 0$ due to a successful treatment. The interpretation is that cancer cells are no longer able to block immune checkpoints. As a result, the advantage of selfishness is minimized, and $r_c(1 - \varepsilon'_s) > r_s$. The new ESS is $(x_{c1}^*, y_{c1}^*, 1)$. The tumor size is $x_{c1}^* = 349.2$, or 56% carrying capacity, which is a 34% increase in size. For the number of T-cells, we have 31.45, 157% of carrying capacity, a decrease of 13% in population. Additionally, we face a decrease in the Allee threshold, from $x_{s2}^* = 23.8$, or 3.8% carrying capacity to

$x_{c2}^* = 17.4$, 2.8% carrying capacity, a decrease of 27%. The tumor has switched from selfish to cooperative cancer cells.

Example 5.4.4 shows the importance of treatment effects on tumor composition: a treatment of the selfish phenotype's crypticity has led to an increase in the steady-state tumor size. However, this is not a general rule. As per Proposition 5.4.3, such a treatment always decreases tumor size if the tumor composition at the steady state is preserved.

The effects of treatment on the patient's prognosis are also ambiguous. Consider once again the configuration used for the previous simulation: $K = 25$, $L = 20$, $r_s = 0.19$, $r_c = 0.2$, $\varepsilon_c = 1$, $a = 1$, complete with $\gamma = 0.1$ and $\delta = 0.005$, and critical tumor masses $\omega_1 = 150$, $\omega_2 = 200$, $\omega_3 = 250$. The initial conditions for the system are $x(0) = 40$, $y(0) = 20$, $g(0) = 0.7$. Table 5.2 contains the T_1 , T_2 , and T_3 levels for different values of ε_s with the same initial conditions as before. Unsurprisingly,

ε_s	T_1	T_2	T_3
0	23	29	33
0.1	24	30	35
0.2	25	31	37
0.3	26	33	41
0.4	27	36	∞
0.5	27	45	∞
0.6	28	∞	∞
0.7	26	37	54
0.8	24	32	39
0.9	22	28	33
1	20	25	29

Table 5.2: Survival times with different values of ε_s for different critical tumor masses. Decreasing the selfish resistance rate below 0.6 decreases the patient's survival time.

survival times are minimized for the maximum resistance rate of 1, and go higher as ε_s is lowered. However, instead of a monotonic relationship, survival times are maximized for a resistance rate of 0.6 and go down again. This implies that treating a patient with resistance rate of 0.7 is dangerous, as exposing the selfish cancer cells to the immune system too much will lower life expectancy.

The explanation of this phenomenon comes from the phenotypes' competition for

resources. Treatment of the selfish phenotype results in increased proliferation by the cooperative phenotype, making the tumor more aggressive without slowing its growth. This is in line with Proposition 5.3.4, as lowering the advantage of selfishness results in a cooperative tumor. For $\varepsilon_s = 0$, the advantage of selfishness is nil, and lower than r_c/r_s , meaning that the tumor converges to the stable cooperative steady state, x_{c1}^* . For all other displayed values of ε_s , the advantage of selfishness is higher than the ratio of birth rates, meaning that the tumor either converges to the stable selfish steady state x_{s1}^* , or goes extinct. In the above cases, for $0.1 \leq \varepsilon_s \leq 0.6$, the tumor goes extinct (but not before potentially visiting the critical tumor mass), and for values $0.7 \leq \varepsilon_s \leq 1$, it converges to the stable selfish steady state.

Figure 5.9 shows the transient dynamics of the cancer and the immune system for four different values of selfish resistance rate, 0, 0.3, 0.6, and 0.9. For a value of 0, the tumor grows rapidly, then progresses to the cooperative steady state, with the selfish phenotype slowly going extinct. For a value of 0.3, the tumor shows rapid growth, but the cooperative cells are outcompeted by the selfish cells. As the cooperative phenotype declines, the T-cells experience a comeback, and eradicate all cancer cells. For a value of 0.6, the same thing happens, but in this case the selfish phenotype is more resistant, which results in a quicker extinction of the cooperative cells, a slower tumor growth, and a slower extinction of the tumor. For a value of 0.9, the cooperative cells are eradicated even faster, but the selfish cells are resistant enough to withstand the immune system, resulting in rapid tumor growth towards a stable equilibrium with large tumor size. Interestingly, optimal patient prognosis occurs when the selfish cells have moderately strong resistance. This produces a rapid eradication of the tumor cells and a smaller maximum tumor size prior to the immune system eventually eliminating the selfish cells as well.

Note that the ambiguous outcome of treatment effects in ε_s is maintained for different configurations, even if $r_c < r_s$. In this case, the advantage of selfishness is always larger than r_c/r_s , hence the only ESS is the stable selfish steady state. However, if the immunosuppression parameter is large enough, a treatment of the selfish resistance rate may still worsen the patient's prognosis, as without enough selfish cells to compete with the cooperative cells, the highly aggressive cooperative phenotype gains a larger growth rate. Thus, even though steady-state tumor composition is preserved and treatment decreases steady-state tumor size, the survival time of the patient may still decrease.

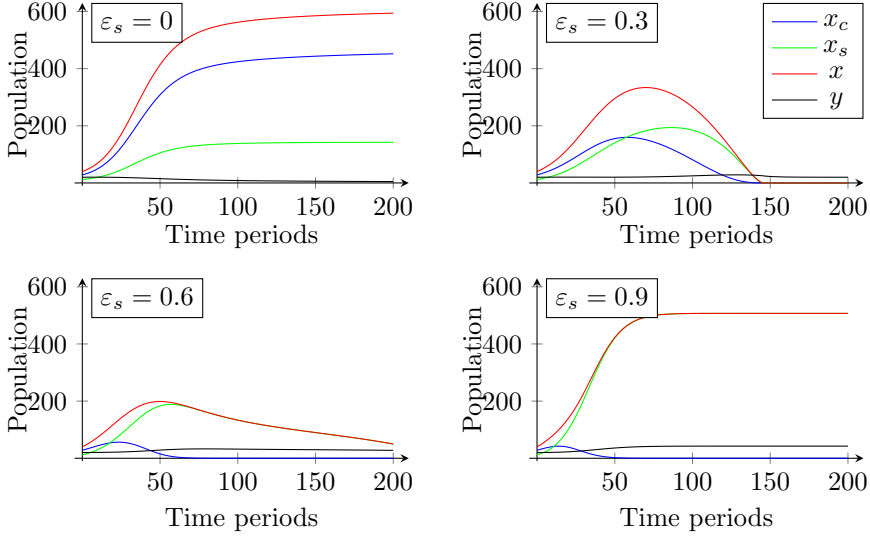


Figure 5.9: Transient dynamics of the system with different selfish resistance parameter values. For small and large values of ε_s the tumor progresses to its stable equilibrium, for intermediate values, it is eliminated. Intermediate values of selfish resistance result in better patient prognosis than extreme ones.

5.5 Conclusion

We use game theory to model the dynamics of cancer cells and the immune system as a predator-prey system. The cancer cells exhibit two strategies. One confers resistance (safety) to the focal individual with no effect on the survival of others, while the second confers safety to both the individual and its neighbors. We refer to the former strategy as “selfish” because it only helps itself, and the latter as “cooperative” because it provides a public good. The model always exhibits three general outcomes: extinction of the cancer cells with some residual population of T-cells, the coexistence of one or the other cancer strategy with a population of T-cells, or the extinction of the T-cells with cancer cells going to their carrying capacity (this can only happen when the cancer cells are cooperative).

The model always results in a single pure-strategy ESS. If the two strategies share the same growth parameters, save for those associated with providing the public good, then the selfish strategy outcompetes the cooperative one. In this way, the game is a kind of prisoner’s dilemma embedded within a population model based on logistic growth and “predation” from the immune system. Others who have addressed the

more general question of the evolution of cooperation, such as Axelrod and Axelrod (1984) have found that cooperation can prevail against selfishness if there are non-random interactions such as the clumping of individual by type in space (Nowak and May, 1992; Ale et al., 2013), or if cooperators can either withhold benefits from selfish individuals through conditional strategies such as Tit-for-Tat (Nowak and Sigmund, 1992), or otherwise punish defectors (Fudenberg and Maskin, 1986). Neither of these features for evolving cooperation occur within our model. However, cooperation can be the ESS if the intrinsic growth rate of the cooperators is sufficiently higher than that of the selfish individuals. If the advantage of selfishness is higher than the ratio of reproduction rates, the tumor is selfish, if it is lower, the tumor is cooperative.

As a predator-prey model, our immune-cancer model has several properties and consequences. As victims of the T-cells, the cancer cells experience mortality that increases with the population size of T-cells. Furthermore, the T-cells exert mortality as a Type II-like functional response. This means that the cancer cells experience safety in numbers, regardless their type, and so their isocline is humped shaped as seen in Rosenzweig and MacArthur (1963). The upward sloping region of a cancer cell's isocline (in the state space of cancer cell and T-cell population sizes) represents an Allee effect (Taylor and Hastings, 2005) as safety in numbers more than compensates for intra-cancer cell competition.

The T-cell population is not a predator in the traditional sense. In our model they have a carrying capacity set by the innate immune system of the patient. This attribute is in line with predator-prey models where the predators receive a subsidy from outside of the system (Nevai and Van Gorder, 2012). But, unlike these models, the T-cells gain no resources nor survival advantages from killing cancer cell. The growth rate of the T-cells does not increase with the mortality that they induce, rather, it is the overall population size of cancer cells that stimulates their growth rate. To incorporate this reality of T-cells, we let their carrying capacity increase with cancer cell population size. When the cancer cells are selfish, this produces a T-cell zero-growth isocline with a positive slope and a positive T-cell intercept. Combining the selfish cancer cell's and T-cell's isoclines produces up to two interior solutions. The first is an unstable extinction threshold, typical of many ecological models with an Allee effect. If the cancer population size drops too low they will go extinct and the cancer is cured. If the population size grows to a level above this threshold then a second, stable equilibrium is achieved (for some parameterizations this second interior equilibrium may yield a stable limit cycle) with a substantially higher cancer cell population and a higher T-cell population as well.

The interaction between the cooperative cancer cells and the immune system causes the model to deviate further from a standard predator-prey model. The immunosuppressive effect of the cooperative cancer cells means their presence exerts mortality on the T-cells — the cancer cells now have a predator-like effect on the T-cells even as they still enhance the T-cells’ carrying capacity. The combined effects of the cooperative cancer cells on the T-cells means that the T-cell’s zero-growth isocline can now either exhibit a hump (allowing for either two or even three interior equilibria), or it may be everywhere negatively sloped (allowing for either two interior equilibria or a single interior equilibrium). In this case, if the isoclines only permit a single interior equilibrium then it is a double extinction threshold, meaning that it separates the region where the T-cells drive the cancer cells extinct from the region where the cancer cells drive the T-cells extinct, meaning complete immunosuppression.

By considering the two immune evasion strategies our model reveals a more diverse array of transient dynamics and equilibria than the models of Robertson-Tessi et al. (2012) and Nakada et al. (2016). Furthermore, we have included the two ways that cancer cells might promote or inhibit T-cell persistence and recruitment in a more sophisticated and realistic fashion. Our two immune evasion strategies of selfish and cooperative mimic two broad classes of resistance strategies by cancer cells. In response to the immune system some cancer cells down-regulate the antigen that elicits T-cell attacks. This amounts to a form of crypticity that selfishly protects the focal cell while likely having little impact on the other cancer cells around it. On the other hand, cancer cells can also evolve the means to actually suppress T-cell recruitment and proliferation. This can be done by either directly signaling T-cells to cease functioning or by downregulating the cues that signal macrophages to stimulate T-cell proliferation (Gajewski et al., 2013).

Immunotherapy in our model acts through model parameters that enhance the lethal encounter rate of T-cells with cancer cells, reduce the advantage of selfishness (render crypticity ineffective), or decrease the immunosuppressive effect of cooperators (Ghirelli and Hagemann, 2013). These actions generally have positive therapeutic outcomes in terms of slowing the rate of tumor growth, increasing the extinction threshold below which the cancer will be cured, and decreasing the ultimate size of the tumor. If this size is below the lethal threshold for the patient then this may be a durable outcome, akin to adaptive therapy (Gatenby et al., 2009). However, by rendering crypticity ineffective, the tumor may shift from one dominated by selfish cells to one dominated by cooperative ones. As an unintended consequence, tumor size at the evolutionarily ESS may increase, and the patient’s prognosis may be substantially worsened.

Our model can also consider other therapies (singly or in combination) where a cytotoxic drug or radiation therapy may influence the efficacy of the patient’s natural immune system or additional immunotherapy. For instance, a cytotoxic drug may reduce the cancer cell’s intrinsic growth rate by targeting cells with high proliferation rates. Or a targeted therapy may reduce the population size of cancer cells to below their extinction threshold. This becomes akin to the abscopal effect (Mole, 1953). A number of clinical examples exist where radiation therapy can induce a stronger immune response. If the cancer cell population is sufficiently depressed by radiation therapy, ours and other models suggest an extinction threshold where the immune system can now eradicate the residual cancer cell population.

A main takeaway of this chapter is that in order to shrink, or fix tumor size through immunotherapy, the effects on tumor composition cannot be ignored. This is a well-known thesis in other aspects of cancer therapy, e.g. in the treatment of prostate cancer (You et al., 2017), but is yet to be adopted in immunotherapy. More specifically, our model identifies increasing immune efficiency and decreasing immunosuppression as viable treatment options. Both are shown to lower tumor burden and decrease the likelihood of cancer progression and metastasis via decreasing tumor size at the ESS, as well as to increase the patient’s life expectancy. In most dynamic models of heterogeneous tumors – this chapter’s included – the quantitative effects of treatment with respect to tumor composition are analytically intractable. The same holds for life expectancy or most other proxies of the patient’s prognosis. However, based on our work, a general qualitative statement can be formalized: treatment strategies that influence tumor composition are risky, whereas those that preserve composition are viable. Our findings merit more research aimed out identifying the risks and potential benefits of the former kind of treatment, and to validate our conclusions for the latter, whereas our proposed framework opens the possibility of integrating game theoretic concepts and ideas into the modeling of the interactions between heterogeneous tumors and the immune system.

5.A Appendix

The appendix of this chapter contains the more technical proofs.

Proposition 5.3.6. *Let $D_s > 0$. If $x_{s1}^* > \frac{K}{4}$, then $(x_{s1}^*, y_{s1}^*, 0)$ is linearly stable in x and y .*

Proof. The system is as follows.

$$\begin{aligned}\dot{x} &= r_s(1 - \sqrt{\frac{x}{K}})x - \gamma(1 - \varepsilon_s)\sqrt{x}y = F(x, y), \\ \dot{y} &= r_i(1 - \frac{y}{L + a\sqrt{x}})y = G(x, y).\end{aligned}$$

Consider the Jacobian of the problem.

$$J(x_{s1}^*, y_{s1}^*) = \begin{pmatrix} \frac{\partial F(x, y)}{\partial x} & \frac{\partial F(x, y)}{\partial y} \\ \frac{\partial G(x, y)}{\partial x} & \frac{\partial G(x, y)}{\partial y} \end{pmatrix} (x_{s1}^*, y_{s1}^*).$$

The elements are as follows.

1, $\frac{\partial F(x, y)}{\partial x}(x_{s1}^*, y_{s1}^*) < 0$. We can show this by taking

$$\frac{\partial F(x, y)}{\partial x} = r_s(1 - \frac{3\sqrt{x}}{2\sqrt{K}}) - \frac{y}{2\sqrt{x}}\gamma(1 - \varepsilon_s).$$

Note that since $\dot{x} = 0$ at the steady state, we have $r_s(1 - \frac{\sqrt{x_{s1}^*}}{\sqrt{K}}) = \gamma(1 - \varepsilon_s) \frac{y_{s1}^*}{\sqrt{x_{s1}^*}}$.

Thus, substituting gives

$$\frac{\partial F(x, y)}{\partial x}(x_{s1}^*, y_{s1}^*) = r_s(\frac{1}{2} - \frac{\sqrt{x_{s1}^*}}{\sqrt{K}}) < 0,$$

by using $\sqrt{x_{s1}^*} > \sqrt{K}/2$.

2, $\frac{\partial F(x, y)}{\partial y}(x_{s1}^*, y_{s1}^*) < 0$. This is shown by simply taking

$$\frac{\partial F(x, y)}{\partial y} = -\gamma(1 - \varepsilon_s)\sqrt{x} < 0,$$

for all $x > 0$.

3, $\frac{\partial G(x, y)}{\partial x}(x_{s1}^*, y_{s1}^*) > 0$. This is shown by taking

$$\frac{\partial G(x, y)}{\partial x} = r_i \frac{a}{2\sqrt{x}} \frac{y^2}{(L + a\sqrt{x})^2} > 0,$$

for all $x, y > 0$.

4, $\frac{\partial G(x, y)}{\partial y}(x_{s1}^*, y_{s1}^*) < 0$. This is shown by taking

$$\frac{\partial G(x, y)}{\partial y} = r_i - \frac{2y}{L + a\sqrt{x}}.$$

At the steady state we have $y_{s1}^* = L + a\sqrt{x_{s1}^*}$, thus $\frac{\partial G(x, y)}{\partial y}(x_{s1}^*, y_{s1}^*) = r_i - 2r_i = -r_i < 0$.

So, qualitatively, we have

$$J(x_{s1}^*, y_{s1}^*) = \begin{pmatrix} - & - \\ + & - \end{pmatrix},$$

meaning that $\text{Tr}(J(x_{s1}^*, y_{s1}^*)) < 0$, and $\text{Det}(J(x_{s1}^*, y_{s1}^*)) > 0$, thus the steady state (x_{s1}^*, y_{s1}^*) is stable for deviations in x and y . ■

Lemma 5.A.1. *For the steady state, $(x_{c1}^*, y_{c1}^*, 1)$, it holds that*

$$\frac{x_{c1}^*}{\sqrt{K}(L + \frac{ad}{r_i}x_{c1}^*)} > \frac{\gamma(1 - \varepsilon_c)}{r_c}.$$

Proof. At $(x_{c1}^*, y_{c1}^*, 1)$, the zero isocline of y intersects that of x from below. Hence, the derivative of the former, evaluated at the steady state, is greater:

$$\sqrt{x_{c1}^*} \left(\frac{1}{2} - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}} \right) \frac{r_c}{\gamma(1 - \varepsilon_c)} < \frac{a}{2} \sqrt{x_{c1}^*} - \frac{\delta}{2r_i} L \sqrt{x_{c1}^*} - a \frac{\delta}{r_i} x_{c1}^*. \quad (5.12)$$

At the same time, since the point (x_{c1}^*, y_{c1}^*) lies upon both isoclines we have

$$y_{c1}^* = \sqrt{x_{c1}^*} \left(1 - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}} \right) \frac{r_c}{\gamma(1 - \varepsilon_c)} = L + a \sqrt{x_{c1}^*} - \frac{\delta}{r_i} L \sqrt{x_{c1}^*} - a \frac{\delta}{r_i} x_{c1}^*.$$

Hence, we get

$$\sqrt{x_{c1}^*} \left(1 - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}} \right) \frac{r_c}{\gamma(1 - \varepsilon_c)} - L + a \frac{\delta}{r_i} x_{c1}^* = a \sqrt{x_{c1}^*} - \frac{\delta}{r_i} L \sqrt{x_{c1}^*}.$$

Applying 5.12 gives

$$\sqrt{x_{c1}^*} \left(1 - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}} \right) \frac{r_c}{\gamma(1 - \varepsilon_c)} - L + a \frac{\delta}{r_i} x_{c1}^* > 2 \sqrt{x_{c1}^*} \left(\frac{1}{2} - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}} \right) \frac{r_c}{\gamma(1 - \varepsilon_c)} + 2a \frac{\delta}{r_i} x_{c1}^*.$$

After rearranging we have

$$\frac{x_{c1}^*}{\sqrt{K}} \frac{r_c}{\gamma(1 - \varepsilon_c)} - L - a \frac{\delta}{r_i} x_{c1}^* > 0.$$

A final rearrangement gives

$$\frac{x_{c1}^*}{\sqrt{K}(L + \frac{ad}{r_i}x_{c1}^*)} > \frac{\gamma(1 - \varepsilon_c)}{r_c},$$

as stated. ■

Proposition 5.3.11 *Let $D_c > 0$. If $x_{c1}^* > \frac{K}{4}$, then $(x_{c1}^*, y_{c1}^*, 1)$ is linearly stable in x and y .*

Proof. The system is as follows.

$$\begin{aligned}\dot{x} &= r_c(1 - \sqrt{\frac{x}{K}})x - \gamma(1 - \varepsilon_c)\sqrt{x}y = F(x, y), \\ \dot{y} &= r_i(1 - \frac{y}{L + a\sqrt{x}})y - \delta\sqrt{x}y = G(x, y).\end{aligned}$$

Again, consider the Jacobian of the problem.

$$J(x_{c1}^*, y_{c1}^*) = \begin{pmatrix} \frac{\partial F(x, y)}{\partial x} & \frac{\partial F(x, y)}{\partial y} \\ \frac{\partial G(x, y)}{\partial x} & \frac{\partial G(x, y)}{\partial y} \end{pmatrix} (x_{c1}^*, y_{c1}^*).$$

As before, we have 1, $\frac{\partial F(x, y)}{\partial x}(x_{c1}^*, y_{c1}^*) = r_c(1 - \frac{3\sqrt{x_{c1}^*}}{2\sqrt{K}}) - \frac{y_{c1}^*}{2\sqrt{x_{c1}^*}}\gamma(1 - \varepsilon_c) = r_c(\frac{1}{2} - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}})$.

2, $\frac{\partial F(x, y)}{\partial y}(x_{c1}^*, y_{c1}^*) = -\gamma(1 - \varepsilon_c)\sqrt{x_{c1}^*}$. Since (x_{c1}^*, y_{c1}^*) is on the zero-growth isocline of x we have $\gamma(1 - \varepsilon_c)\sqrt{x_{c1}^*}y_{c1}^* = r_c(1 - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}})x_{c1}^*$, hence

$$\frac{\partial F(x, y)}{\partial y}(x_{c1}^*, y_{c1}^*) = -r_c(1 - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}})\frac{x_{c1}^*}{y_{c1}^*}.$$

3, $\frac{\partial G(x, y)}{\partial x}(x_{c1}^*, y_{c1}^*) = r_i\frac{a}{2\sqrt{x_{c1}^*}}\frac{y_{c1}^{*2}}{(L + a\sqrt{x_{c1}^*})^2} - \frac{\delta}{2\sqrt{x_{c1}^*}}y_{c1}^*$. Since the steady state is on the zero-growth isocline of y we have $r_i(1 - \frac{y_{c1}^*}{L + a\sqrt{x_{c1}^*}}) = \delta\sqrt{x_{c1}^*}$, hence

$$\frac{\partial G(x, y)}{\partial x}(x_{c1}^*, y_{c1}^*) = \frac{1}{2}r_i(a\sqrt{x_{c1}^*}\frac{y_{c1}^*}{(L + a\sqrt{x_{c1}^*})^2} + \frac{y_{c1}^*}{L + a\sqrt{x_{c1}^*}} - 1)\frac{y_{c1}^*}{x_{c1}^*}.$$

4, $\frac{\partial G(x, y)}{\partial y} = r_i(1 - \frac{2y}{L + a\sqrt{x}}) - \delta\sqrt{x}$. Using the steady-state condition, i.e. $r_i(1 - \frac{y_{c1}^*}{L + a\sqrt{x_{c1}^*}}) = \delta\sqrt{x_{c1}^*}$ gives

$$\frac{\partial G(x, y)}{\partial y}(x_{c1}^*, y_{c1}^*) = -r_i\frac{y_{c1}^*}{L + a\sqrt{x_{c1}^*}}.$$

As before, $\text{Tr}(J(x_{c1}^*, y_{c1}^*)) < 0$, for $x_{c1}^* > K/4$. For the determinant, we calculate

$$\begin{aligned}\frac{\partial F(x, y)}{\partial x}(x_{c1}^*, y_{c1}^*)\frac{\partial G(x, y)}{\partial y}(x_{c1}^*, y_{c1}^*) &= -r_cr_i(\frac{1}{2} - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}})\frac{y_{c1}^*}{L + a\sqrt{x_{c1}^*}}, \\ &= \\ &= -\frac{1}{2}r_cr_i(1 - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}})(a\sqrt{x_{c1}^*}\frac{y_{c1}^*}{(L + a\sqrt{x_{c1}^*})^2} + \frac{y_{c1}^*}{L + a\sqrt{x_{c1}^*}} - 1).\end{aligned}$$

Therefore, the determinant is as follows:

$$\begin{aligned}
& \text{Det}(J(x_{c1}^*, y_{c1}^*)) \\
&= \frac{1}{2} r_c r_i \left(\frac{\sqrt{x_{c1}^*}}{\sqrt{K}} - 1 - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}} \frac{a\sqrt{x_{c1}^*} y_{c1}^*}{(L + a\sqrt{x_{c1}^*})^2} + \frac{a\sqrt{x_{c1}^*} y_{c1}^*}{(L + a\sqrt{x_{c1}^*})^2} + \frac{\sqrt{x_{c1}^*}}{\sqrt{K}} \frac{y_{c1}^*}{L + a\sqrt{x_{c1}^*}} \right), \\
&= \frac{1}{2} r_c r_i \left(\left(1 - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}}\right) \left(\frac{a\sqrt{x_{c1}^*} y_{c1}^*}{(L + a\sqrt{x_{c1}^*})^2} - 1\right) + \frac{\sqrt{x_{c1}^*}}{\sqrt{K}} \frac{y_{c1}^*}{L + a\sqrt{x_{c1}^*}} \right).
\end{aligned}$$

We use the fact that the steady state lies along the zero-growth isocline of x and substitute $y_{c1}^* = \sqrt{x_{c1}^*} \left(1 - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}}\right) \frac{r_c}{\gamma(1-\varepsilon_c)}$ into the last bracketed term to get

$$\begin{aligned}
& \text{Det}(J(x_{c1}^*, y_{c1}^*)) = \\
& \frac{1}{2} r_c r_i \left(1 - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}}\right) \left(\frac{x_{c1}^*}{K} \frac{r_c}{\gamma(1-\varepsilon_c)} \frac{1}{L + a\sqrt{x_{c1}^*}} + \frac{y_{c1}^* a\sqrt{x_{c1}^*}}{(L + a\sqrt{x_{c1}^*})^2} - 1 \right).
\end{aligned}$$

Now we use the zero-growth isocline of y and substitute $y_{c1}^* = (L + a\sqrt{x_{c1}^*}) \left(1 - \frac{\delta}{r_i} \sqrt{x_{c1}^*}\right)$, which gives

$$\begin{aligned}
& \text{Det}(J(x_{c1}^*, y_{c1}^*)) = \\
& \frac{1}{2} r_c r_i \left(1 - \frac{\sqrt{x_{c1}^*}}{\sqrt{K}}\right) \left(\frac{x_{c1}^*}{\sqrt{K}} \frac{r_c}{\gamma(1-\varepsilon_c)} \frac{1}{L + a\sqrt{x_{c1}^*}} + a\sqrt{x_{c1}^*} \left(1 - \frac{\delta}{r_i} \sqrt{x_{c1}^*}\right) \frac{1}{L + a\sqrt{x_{c1}^*}} - 1 \right)
\end{aligned}$$

therefore, $\text{Det}(J(x_{c1}^*, y_{c1}^*)) > 0$ if and only if we have

$$\frac{x_{c1}^*}{\sqrt{K}} \frac{r_c}{\gamma(1-\varepsilon_c)} + a\sqrt{x_{c1}^*} \left(1 - \frac{\delta}{r_i} \sqrt{x_{c1}^*}\right) > L + a\sqrt{x_{c1}^*}.$$

Rearranging gives the condition as

$$\frac{x_{c1}^*}{\sqrt{K} \left(L + \frac{ad}{r_i} x_{c1}^*\right)} > \frac{\gamma(1-\varepsilon_c)}{r_c}.$$

By Lemma 5.A.1, this holds, hence the proof is complete. ■

Bibliography

- Adam, J.A., and Maggelakis, S.A., 1989. Mathematical models of tumor growth. IV. Effects of a necrotic core. *Mathematical Biosciences*, 97, 121-136.
- Ale, S.B., Brown, J.S., and Sullivan, A.T., 2013. Evolution of cooperation: combining kin selection and reciprocal altruism into matrix games with social dilemmas. *PloS one*, 8, p.e63761.
- Allouch, N., 2015. On the private provision of public goods on networks. *Journal of Economic Theory*, 157, 527-552.
- Axelrod, R., and Axelrod, R.M., 1984. The evolution of cooperation (Vol. 5145). Basic Books (AZ).
- Baar, M., Coquille, L., Mayer, H., Hölzel, M., Rogava, M., Tüting, T., and Bovier, A., 2016. A stochastic model for immunotherapy of cancer. *Scientific Reports*, 6, article nr. 24169.
- Babbs, C.F., 2012. Predicting success or failure of immunotherapy for cancer: insights from a clinically applicable mathematical model. *American Journal of Cancer Research*, 2, 204-213.
- Ballester, C., Calvó-Armengol, A., and Zenou, Y., 2006. Who's who in networks. Wanted: The key player. *Econometrica*, 74, 1403-1417.
- Ballester, C. and Calvó-Armengol, A., 2010. Interactions with hidden complementarities. *Regional Science and Urban Economics*, 40, 397-406.
- Bayer, P., Brown, J.S., and Staňková, K., 2018. A two-phenotype model of immune evasion by cancer cells. *Journal of Theoretical Biology*, 455, 191-204.
- Bayer, P., Herings, P., Peeters, R., and Thuijsman, F., 2017. Adaptive learning in weighted network games. *GSBE Research Memorandum No. 25*.

- Bayer, P., Herings, P., and Peeters, R., 2019. Myopic and farsighted players in the local public goods game. Working paper.
- Bayer, P., and Kozics, Gy., 2018. Best-response dynamics in directed network games. Working paper.
- Belhaj, M., Bramoullé, Y., and Deroïan, F., 2014. Network games under strategic complementarities. *Games and Economic Behavior*, 88, 310-319.
- Bervoets, S., and Faure, M., 2019. Stability in games with continua of equilibria. *Journal of Economic Theory*, 179, 131-162.
- Bervoets, S., Bravo, M., and Faure, M., 2016. Learning and convergence to Nash in games with continuous action sets. Working paper.
- Blume, L.E., Brock, W.A., Durlauf, S.N., and Ioannides, Y.M., 2010. Identification of social interactions. In *Handbook of Social Economics*, edited by Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, 853-964. Amsterdam: North Holland.
- Bourlès, R., Bramoullé, Y., and Perez-Richet, E., 2017. Altruism in networks. *Econometrica*, 85, 675-689.
- Bramoullé, Y., and Kranton, R., 2007. Public goods in networks. *Journal of Economic Theory* 135, 478-494.
- Bramoullé, Y., Kranton, R., and D'Amours, M., 2014. Strategic interaction and networks. *American Economic Review* 104, 898-930.
- Calvo-Armengol, A., and Zenou, Y., 2004. Social networks and crime decisions: The role of social structure in facilitating delinquent behavior. *International Economic Review* 45, 939-958.
- Camerer, C., Ho, T.H., Chong, J.K., 2002. Sophisticated experience-weighted attraction learning and strategic teaching in repeated games. *Journal of Economic Theory* 104, 137-188.
- De Pillis, L.G., and Radunskaya, A., 2001. A mathematical tumor model with immune resistance and drug therapy: an optimal control approach. *Computational and Mathematical Methods in Medicine*, 3, 79-100.
- De Angelis, E., and Mesin, L., 2001. Modelling of the immune response: Conceptual frameworks and applications. *Mathematical Models and Methods in Applied Sciences*, 11(09), 1609-1630.

- Dimberu, P.M., and Leonhardt, R.M., 2011. Cancer immunotherapy takes a multifaceted approach to kick the immune system into gear. *The Yale Journal of Biology and Medicine*, 84, 371-380.
- Dhodapkar, M.V., 2013. Personalized immune-interception of cancer and the battle of two adaptive systems — when is the time right? *Cancer Prevention Research*, 6, 173-176.
- Dubey, P., Haimanko, O., and Zapechelnyuk, A., 2006. Strategic complements and substitutes, and potential games. *Games and Economic Behavior* 54, 77-94.
- Eftimie, R., Bramson, J.L., and Earn, D.J.D., 2011. Interactions between the immune system and cancer: a brief review of non-spatial mathematical models. *Bulletin of Mathematical Biology*, 73, 2-32.
- Ellison, G., 1997. Learning from personal experience: One rational guy and the justification of myopia. *Games and Economic Behavior* 19, 180-210.
- Eksin, C., Molavi, P., Ribeiro, A., and Jadbabaie, A., 2012. Learning in linear games over networks. In *Communication, Control, and Computing (Allerton)*, 50th Annual Allerton Conference (434-440). IEEE.
- Elliott, M., and Golub, B., 2019. A network approach to public goods. *Journal of Political Economy* (forthcoming).
- Eshel, I., Samuelson, L., Shaked, A. 1998. Altruists, egoists, and hooligans in a local interaction model. *American Economic Review* 88, 157-179.
- Ewerhart, C., 2017. The lottery contest is a best-response potential game. *Economics Letters* 155, 168-171.
- Foryś U., and Mokwa-Borkowska, A., 2005. Solid tumour growth analysis of necrotic core formation. *Mathematical and Computer Modelling*, 42, 593-600.
- Fudenberg, D., and Levine, D.K., 1998. *The Theory of Learning in Games*. MIT press.
- Fudenberg, D., and Maskin, E., 1986. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica: Journal of the Econometric Society*, 533-554.
- Gajewski, T.F., Schreiber, H., and Fu, Y.X., 2013. Innate and adaptive immune cells in the tumor microenvironment. *Nature Immunology*, 14, 1014-1022.

- Gatenby, R.A., Silva, A.S., Gillies, R.J., and Frieden, B.R., 2009. Adaptive therapy. *Cancer research*, 69, 4894-4903.
- Ghirelli, C., and Hagemann, T., 2013. Targeting immunosuppression for cancer therapy. *The Journal of Clinical Investigation*, 123, p.2355-2357.
- Goswami, S., Aparicio, A., and Subudhi, S.K., 2016. Immune checkpoint therapies in prostate cancer. *Cancer Journal*, 22, 117-120.
- Goyal, S., 2012. *Connections: An Introduction to the Economics of Networks*. Princeton University Press.
- Hanahan, D., and Weinberg, R.A. 2011., Hallmarks of cancer: the next generation. *Cell*, 144, 646-674
- Helsley, R.W., and Zenou, Y., 2014. Social networks and interactions in cities. *Journal of Economic Theory*, 150, 426-466.
- Hicklin, D.J., Marincola, F.M., and Ferrone, S., 1999. HLA class I antigen downregulation in human cancers: T-cell immunotherapy revives an old story. *Molecular Medicine Today*, 5, 178-186.
- Ibrahim-Hashim, A., Roberson-Tessi, M., Xu, L., Balagurunathan, Y., Wojtkowiak, J.W., El-kenawi, A.E., Enriquez-Navas, P.M., Russell, S., Kam, Y., Lloyd, M.C., Bui, M., Brown, J.S., Anderson, A.R.A., Gillies, R.J., and Gatenby, R.A., 2017. Defining cancer subpopulations by adaptive strategies rather than molecular properties provides novel insights into intratumoral evolution. *Cancer Research* 77, 2242-2254.
- Jackson, M.O., and Zenou, Y., 2014. Games on networks. *Handbook of Game Theory*, Vol. 4, Peyton Young and Shmuel Zamir, eds., Elsevier Science.
- Jensen, M.K., 2010. Aggregative games and best-reply potentials. *Economic Theory* 43, 45-66.
- Johnsen, A.K., Templeton, D.J., Sy, M.S., and Harding, C.V., 1999. Deficiency of transporter for antigen presentation (TAP) in tumor cells allows evasion of immune surveillance and increases tumorigenesis. *The Journal of Immunology*, 163, 4224-4231.
- Kareva, I., and Berezovskaya, F., 2015. Cancer immunoediting: a process driven by metabolic competition as a predator-prey-shared resource type model. *Journal of Theoretical Biology*, 380, 463-472.

- König, M., Liu, X. and Zenou, Y., 2014. R&D networks: Theory, empirics and policy implications. *Review of Economics and Statistics* (forthcoming).
- Kukushkin, N.S., 2004. Best response dynamics in finite games with additive aggregation. *Games and Economic Behavior*, 48, 94-110.
- Kukushkin, N.S., 2015. Cournot tatonnement and potentials. *Journal of Mathematical Economics*, 59, 117-127.
- Leontief, W., 1970. Environmental repercussions and the economic structure: An input-output approach. *The Review of Economics and Statistics* 52, 262-271.
- Maeurer, M.J., Gollin, S.M., Martin, D., Swaney, W., Bryant, J., Castelli, C., Robbins, P., Parmiani, G., Storkus, W.J., and Lotze, M.T., 1996. Tumor escape from immune recognition: lethal recurrent melanoma in a patient associated with down-regulation of the peptide transporter protein TAP-1 and loss of expression of the immunodominant MART-1/Melan-A antigen. *Journal of Clinical Investigation*, 98, 1633-41.
- Mole, R.H., 1953. Whole body irradiation – radiobiology or medicine? *The British Journal of Radiology*, 26, 234-241.
- Monderer, D., and Shapley, L., 1996. Potential games. *Games and Economic Behavior* 14, 124-143.
- Morgan, R.A., Dudley, M.E., Wunderlich, J.R., Hughes, M.S., Yang, J.C., Sherry, R.M., Royal, R.E., Topalian, S.L., Kammula, U.S., Restifo, N.P., Zheng, Z., Nahvi, A., de Vries, C.R., Rogers-Freezer, L.J., Mavroukakis, S.A., and Rosenberg, S.A., 2006. Cancer regression in patients after transfer of genetically engineered lymphocytes. *Science*, 314, 126-129.
- Nakada, N., Nagata, M., Takeuchi, Y., and Nakaoka, S., 2016. Mathematical modeling and analysis of combinational immune boost for tumor elimination. In *AIP Conference Proceedings* (Vol. 1723, No. 1, p. 020002). AIP Publishing.
- Nani, F., and Freedman, H.I., 2000. A mathematical model of cancer treatment by immunotherapy. *Mathematical Biosciences*, 163, 159-199.
- Nevai, A.L., and Van Gorder, R.A., 2012. Effect of resource subsidies on predator–prey population dynamics: a mathematical model. *Journal of Biological Dynamics*, 6, 891-922.

- Nowak, M.A., and May, R.M., 1992. Evolutionary games and spatial chaos. *Nature*, 359, 826-829.
- Nowak, M.A., and Sigmund, K., 1992. Tit for tat in heterogeneous populations. *Nature*, 355, 250-253.
- Panetta, J.C., 1998. A mathematical model of drug resistance: heterogeneous tumors. *Mathematical Biosciences*, 147, 41-61.
- Parcesepe, P., Giordano, G., Laudanna, C., Febbraro, A., and Pancione, M., 2016. Cancer-associated immune resistance and evasion of immune surveillance in colorectal cancer. *Gastroenterology research and practice*, vol. 2016, Article ID 6261721, 8 pages.
- Pardoll, D.M., 2012. The blockade of immune checkpoints in cancer immunotherapy. *Nature Reviews Cancer*, 12, 252-264.
- Pimpunchat, B., Sweatman, W.L., Wake, G.C., Triampo, W. and Parshotam, A., 2009. A mathematical model for pollution in a river and its remediation by aeration. *Applied Mathematics Letters* 22, 304-308.
- Ribas, A., 2015. Adaptive immune resistance: How cancer protects from immune attack. *Cancer Discovery*, 5, 915-919.
- Robertson-Tessi, M., El-Kareh, A., and Goriely, A., 2012. A mathematical model of tumor-immune interactions. *Journal of Theoretical Biology*, 294, 56-73.
- Rosenzweig, M.L., and MacArthur, R.H., 1963. Graphical representation and stability conditions of predator-prey interactions. *The American Naturalist*, 97, 209-223.
- Sandler, T., and Hartley, K., 1995. *The Economics of Defense*. Cambridge University Press.
- Sandler, T., and Hartley, K., 2001. Economics of alliances: The lessons for collective action. *Journal of Economic Literature* 39, 869-896.
- Sandler, T., and Hartley, K., 2007. *Handbook of Defense Economics: Defense in a globalized world*. Elsevier.
- Seliger, B., Maeurer, M. J., and Ferrone, S. 1997. TAP off – tumors on. *Immunology today*, 18, 292-299.

- Sharma, P., Hu-Lieskovan, S., Wargo, J. A., and Ribas, A., 2017. Primary, adaptive, and acquired resistance to cancer immunotherapy. *Cell*, 168, 707-723.
- Smyth, M.J., Dunn, G.P., and Schreiber, R.D., 2006. Cancer immunosurveillance and immunoediting: the roles of immunity in suppressing tumor development and shaping tumor immunogenicity. *Advances in Immunology*, 90, 1-50.
- Stewart, T.J., and Smyth, M.J., 2011. Improving cancer immunotherapy by targeting tumor-induced immunosuppression. *Cancer Metastasis Rev*, 30, 125-140.
- Sun, X., Bao, J., and Shao, Y., 2016. Mathematical modeling of therapy-induced cancer drug resistance: connecting cancer mechanisms to population survival rates. *Scientific Reports*, 6.
- Taylor, C.M., and Hastings, A., 2005. Allee effects in biological invasions. *Ecology Letters*, 8, 895-908.
- Tobin, P.C., Berec, L., and Liebhold, A.M., 2011. Exploiting Allee effects for managing biological invasions. *Ecology Letters*, 14, 615-624.
- Tomasetti, C., and Levy, D., 2010. An elementary approach to modeling drug resistance in cancer. *Mathematical Biosciences and Engineering: MBE*, 7, 905.
- Voorneveld, M., 2000. Best-response potential games. *Economic Letters* 66, 289-295.
- Wu, A.A., Drake, V., Huang, H.S., Chiu, S., and Zheng, L., 2015. Reprogramming the tumor microenvironment: tumor-induced immunosuppressive factors paralyze T cells. *Oncoimmunology*, 4, p. e1016700.
- You, L., Brown, J.S., Thuijsman, F., Cunningham, J.J., Gatenby, R.A., Zhang, J., and Staňková, K., 2017. Spatial vs. non-spatial eco-evolutionary dynamics in a tumor growth model. *Journal of Theoretical Biology*, 435, 78-97.

Valorization

“We must know. We will know.”

David Hilbert

This addendum discusses the value-creation opportunities of the research presented in this dissertation. According to the Regulation Governing the Attainment of Doctoral Degrees of Maastricht University, examples of value-creation include “patents, licences, open source tools, software, making models and systems available, co-publications with social and/or economic stakeholders, publications in journals and newspapers, non-academic publications, appearances in the media, contributions to public debates, advice for social organisations or companies, projects for the SME portal, policy recommendations for governments, training programmes for professionals, public-sector workers, organisations or companies, participation in Top Institutes, in United Brains Limburg, collaborations in BioPartner, the Chemelot Campus, the Maastricht Health Campus, Campus Greenport Venlo, the Smart Services Campus, spin-offs and start-up companies, special collections, putting together exhibitions and/or catalogues, participation in administrative bodies.”

Of the twenty-one items listed (counting the collaborations with the various local institutes as one), one is directly applicable to the body of work presented: publications in journals and newspapers. The paper based upon Chapter 5 was published in the *Journal of Theoretical Biology* in 2018, while the paper based upon Chapter 2 is, at the moment of writing, in submission. The papers based upon Chapters 3 and 4 are planned for submission in early 2019.

Given the theoretical nature of the research in this dissertation, other opportunities of value-creation – in the sense that the term was used in the regulations – are indirect and speculative. The remainder of this addendum hypothesizes on the value of each chapter, possibly falling in the contributions to public debates, advice for social organizations or companies, and policy recommendations for governments categories.

Fulfilling some other categories, such as software, making models and systems available, co-publications with social and/or economics stakeholders, non-academic publications, and appearances in the media, might follow from this body of research, but are impossible to speculate.

Chapter 2 considers the properties of the Nash equilibrium in network games. The relevance and value of the Nash equilibrium is, as I was led to believe, beyond dispute. With applications in many distinct fields of science, such as economics, computer science, and biology, the same seems true for networks. The chapter's main result states that, in a setting featuring sequential updates, if each player's update takes him closer to his current best choice, then for any initial condition and for any weighted network the play of the game approaches equilibrium. This result ties in, somewhat indirectly, with one of the most important central policy questions in economics, whether markets should be nationalized, regulated, or left alone. The first step in answering this question for a given market is figuring out the implications of a *laissez-faire* approach. This chapter's result states that weighted network games do not require regulations to reach an equilibrium state.

The core issue in **Chapter 3** is the value of foresight in network games. The level of the agents' foresight is a central characteristic of the agents of financial and economic models. The existence and convergence results of the chapter guarantee that the value of foresight can be calculated for any simple graph. Additionally, the convergence results ensure that the main results of Chapter 2 remain true in case of a single farsighted player whenever the network is characterized by a simple graph. Finally, the chapter's results allow for a qualitative description of the change in the value that the farsighted player receives upon the addition or the removal of a link. Part of the final result, the fact that the farsighted player is better off if two of his neighbors sever contact, is reminiscent of a type of jealousy, implying an indirect but nonetheless interesting connection between network games and psychology.

The topic of **Chapter 4** is very similar to that of Chapter 2. It focuses on the effect of introducing non-reciprocal relationships into the game, represented by a more general class of graphs, called directed networks. The chapter's main result identifies two classes of networks in which best-response cycles cannot exist, and hence, with some permissive assumptions, the convergence results of Chapter 2 remain true in these two classes as well. This result also extends the implications made on the *laissez-faire* approach in Chapter 2 to situations where the players' relationship is hierarchical, or when the own effects of the players are larger than their external effects. It is noteworthy, however, that convergence is not generally established for the whole class

of directed networks, hence the results, and therefore the implications, of Chapter 2 may not be robust for non-reciprocal relationships.

Chapter 5 is written with a specific application, cancer therapy, in mind. While it remains a theoretical work, the potential for value-creation is more direct and less general. The chapter's main result is that immunotherapy of cancer may worsen the patient's prognosis if it causes the tumor to switch from a selfish immune evasion strategy – such as hiding from the killer T-cells – to a cooperative one – such as immunosuppression. It is left for future research to assess the practical extent of this result and address it in a change in therapy as necessary. The model itself is part of a growing strand of literature aiming to understand cancer by the tools and paradigms of game theory, adding to the value of both research fields.

I conclude this addendum by a reflection on the value-creation process of science. Given the many global problems that humanity will need to solve for a sustainable, prosperous, and healthy existence, I fully sympathize with the goal of making science useful. There are two issues to note, however. The first is that no discovery exists in a vacuum; the value of individual pieces of research may not be immediately visible, may be uncertain at the time of their writing, or may come in a way that no one, including the researcher, expects. By having researchers focus on creating value in the short run, society deprives itself from broader discoveries and a deeper understanding, maybe even from value, in the long run. The second is that, as David Hilbert's quote also seems to suggest, human curiosity is tricky; the drive for discovery is often separate from the drive to create value. I find it therefore a non-trivial question whether any focus on value-creation is needed, be it in the short- or the long-run, in order to maximize the societal value of research.

I do not make these points to criticize any particular standing policy. However, I do find it troublesome when researchers self-select on the basis of the perceived usefulness or uselessness of their fields because it interferes with the communication between disciplines. If I am advocating for anything, it is for theorists to try to be interested in some of the more practical questions. Their more abstract thinking may make a connection that an applied person is too focused to make. On the other hand, I wish to reiterate the concerns that a friend of mine had against, as he/she called it, the “marginalization of fundamental research”. Even if the amount of theoretical knowledge that will ever prove itself useful for humans is finite, I am quite positive we have not reached it yet. We should keep going.

Resumé

Péter Bayer was born in Budapest on June 7, 1990. He received his BA diploma in the Applied Economics program of the Corvinus University of Budapest in 2011. He worked at the university's Mathematics Department as a Junior Teaching Assistant from 2009 to 2013. After spending two years in educational doldrums, which saw some fledgling research projects, teaching, and a foray into actuarial mathematics, he entered the Economics MSc program at the Ludwig Maximilian University of Munich, graduating in 2015. In the second year of his Master's



he worked as a Research Assistant at the Max Planck Institute for Innovation and Competition and as a Teaching Assistant at the university's Chair of Organizational Economics. Between 2015 and 2019 he was a PhD candidate at the Department of Economics of Maastricht University under the supervision of Jean-Jacques Herings, Ronald Peeters, and Frank Thuijsman. In 2016 he spent a 1-month research visit at the University of Illinois at Chicago and, in 2018, 3 months at the Aix-Marseille School of Economics in Marseille. The results of his research are collected in this dissertation. He presented his work at various international conferences. Péter will further his academic career as a post-doctoral researcher at Université Grenoble Alpes.