

**Participation and Interaction in Projects
A Game-Theoretic Analysis**

Polevoy, Gleb

DOI

[10.4233/uuid:b40d40dd-f155-4b60-a245-04bdd0760861](https://doi.org/10.4233/uuid:b40d40dd-f155-4b60-a245-04bdd0760861)

Publication date

2016

Document Version

Final published version

Citation (APA)

Polevoy, G. (2016). *Participation and Interaction in Projects: A Game-Theoretic Analysis*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:b40d40dd-f155-4b60-a245-04bdd0760861>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

PARTICIPATION AND INTERACTION IN PROJECTS

A GAME-THEORETIC ANALYSIS

PARTICIPATION AND INTERACTION IN PROJECTS

A GAME-THEORETIC ANALYSIS

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K.C.A.M. Luyben,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op dinsdag 6 december 2016 om 12:30 uur

door

Gleb POLEVOY

Master of Science in Computer Science,
Technion — Israel Institute of Technology, Haifa, Israël,
geboren te Donetsk, Oekraïne.

This dissertation has been approved by the

promoters: prof. dr. C. Witteveen and prof. dr. C.M. Jonker
copromotor: dr. M.M. de Weerd

Composition of the doctoral committee:

Rector Magnificus,	chairman
Prof. dr. C. Witteveen,	Technische Universiteit Delft
Prof. dr. C.M. Jonker,	Technische Universiteit Delft
Dr. M.M. de Weerd,	Technische Universiteit Delft

Independent members:

Prof. dr. F.M. Brazier,	Technische Universiteit Delft
Prof. dr. habil. G. Schäfer	Centrum Wiskunde & Informatica (CWI) and Vrije Universiteit Amsterdam
Prof. dr. M. Uetz,	University of Twente
Prof. dr. E. Elkind,	Oxford University
Prof. dr. ir. D.H.J. Epema,	Technische Universiteit Delft, reserve member



Keywords: projects, agents, budgets, threshold, competition, interaction, recipro-
cation, convergence, habit, Nash equilibrium, subgame perfect equi-
librium, game theory

Author's email: geruskin@yahoo.com

Printed by: www.ridderprint.nl

Front & Back: The author's idea, honed and implemented by Remco Wetzels | On-
twerp & Illustraties | remcowetzels.nl

Copyright © 2016 by G. Polevoy

SIKS Dissertation Series No. 2016-49

The research reported in this thesis has been carried out under the auspices of SIKS,
the Dutch Research School for Information and Knowledge Systems.

This work has been supported by the project SHINE, the flagship project of DIRECT
(Delft Institute for Research on ICT at Delft University of Technology).

ISBN 978-94-6186-766-7

An electronic version of this dissertation is available at

<http://repository.tudelft.nl/>.

Typeset by the author with the \LaTeX Documentation System.

The saddest aspect of life right now is that science gathers knowledge faster than society gathers wisdom.

Isaac Asimov, 1988

This work is dedicated to my family

CONTENTS

Summary	xi
Samenvatting	xv
Preface	xix
1 Introduction	1
1.1 Related Work	3
1.1.1 Shared Effort in Common Projects	3
1.1.2 Reciprocal Interactions	6
1.2 SHINE Project.	7
1.3 Research Questions	7
1.4 Thesis Structure and Contributions	8
1.5 Reading Guide	10
References	11
A A Primer on Game Theory.	15
A.1 Noncooperative Games and Nash Equilibrium.	15
A.2 Efficiency of Nash Equilibria	15
A.3 Mixed Nash Equilibrium	16
2 Shared Effort Games	17
1 Introduction	18
2 Model	20
3 Pure Nash Equilibrium	21
3.1 Linear Project Functions	22
4 Mixed Nash Equilibrium	31
5 Simulations	31
5.1 Infinite-Strategy Fictitious Play for Shared Effort Games	32
5.2 Best Response in a 2-project Game.	33
5.3 The Simulation Method	33
5.4 Results and Conclusions	34
6 Related Work	40
6.1 Fictitious Play	41
7 Conclusions and Future Work.	42
References	43
A Omitted Proofs	45
B Simulation Results	50

3	Competition between Cooperative Projects	61
1	Introduction	62
2	Model	64
3	The Quota Model	66
4	The Success Threshold Model	70
5	Conclusions and Further Research	73
	References	74
4	Towards Decision Support in Reciprocation	77
1	Introduction	78
2	Modeling Reciprocation	80
2.1	Basics	80
2.2	Reciprocation	81
2.3	Context and Examples	81
3	Pairwise Interaction	83
3.1	<i>Fixed</i> Reciprocation	83
3.2	<i>Floating</i> Reciprocation	87
3.3	<i>Fixed</i> and <i>Floating</i> Reciprocation	90
4	Pairwise Interaction: Alternating Case	94
5	Multi-Agent Interaction	95
6	Simulations	102
7	Opinions	103
8	Related Work	105
9	Conclusions and Future Work	106
	References	109
5	The Game of Reciprocation Habits	113
1	Introduction	114
2	Background	117
3	Two Agents: Utility Maximization	118
3.1	Utility Definition for n Agents	118
3.2	Choosing Reciprocation Behavior for Two Agents	119
4	Two Agents: Maximizing Social Welfare	124
5	Two Agents: Reciprocation Attitude Game	127
5.1	PoA and PoS	129
5.2	Converging to NE	130
6	Two Agents: Reciprocation Attitude and Coefficient Game	132
6.1	PoA and PoS	133
6.2	Converging to NE	135
7	n Agents	135
8	Utility Maximization	136
9	Maximizing Social Welfare	138
10	Reciprocation Coefficient Game	139
10.1	PoA and PoS	140
10.2	Converging to NE	141

11	Discussion of the Model	142
11.1	Utility Aggregation	142
11.2	Strategy Space	144
12	Conclusions and Future Work	144
	References	146
6	Reciprocation Effort Games	149
1	Introduction	150
2	Model	152
3	General Convergence	157
4	Reciprocation Effort Game	159
5	Thresholded Reciprocation Effort Game	161
6	Exclusive Thresholded Reciprocation Effort Game	166
7	Extensive Reciprocation Effort Game	167
8	Conclusions and Further Research	168
	References	169
7	Conclusion	173
1	Main Results	174
1.1	The Connection to the SHINE Project	177
2	Future Research	178
2.1	Extensions	178
2.2	Other Directions	180
2.3	Other Approaches to Participation in Projects	180
	References	181
	Acknowledgements	183
	Curriculum Vitæ	185
	List of Publications	187
	References	188
	SIKS Dissertation Series	197

SUMMARY

PARTICIPATION AND INTERACTION IN PROJECTS A GAME-THEORETIC ANALYSIS

Much of what people, governments and companies do is dividing their time or efforts between several activities, and we need to help deciding how to divide the efforts most efficiently. Nice examples of such activities are writing for Wikipedia, authoring papers and books, sharing files on the Internet or just communicating with colleagues. In the example of dividing time between writing Wikipedia articles, authoring scientific papers and a book, the decision may be highly non-trivial. Assume, for instance, that in case of success, the book is the most cost-efficient enterprise. It can still happen that my efforts will not be high enough to be included in the list of the co-authors, or that the other authors will contribute little effort there, and the value of the book will be low. In these complex activities that depend on what the others do, we look for individually rational behavior which is also profitable for the whole group.

The ubiquity of such activities urges us to study them, in order to facilitate efficient effort division. Formally, we want to recommend stable division strategies that result in high social welfare. *Social welfare* means the total utility derived by the whole group of the involved agents (people, governments, etc.) Stability means that no agent can improve her own utility by doing something new, as long as the others keep doing what they are doing. This is the famous *Nash equilibrium*. Since a stable set of strategies (everyone's behavior) is a reasonable description of what will happen in the real life, we look for a Nash equilibrium with as high social welfare as possible.

We concentrate on games modeling dividing efforts between two common sorts of activities: *value-creating* activities like writing an article, and activities of *interaction*, such as communicating with colleagues. For each game we study the possible Nash equilibria and their social welfare. If all the Nash equilibria (which can be multiple) turn out to be *socially efficient* (which stands for equilibria which have social welfare which ratio to the maximum possible social welfare is close to 1), then no regulation is required. If there exist efficient as well as inefficient equilibria, then we may want to convince the participants to opt for the socially efficient equilibria. In the unlucky case when only inefficient equilibria exist, we may want to alter the whole situation by, e.g., subsidizing certain strategies of the agents.

First, we study the value-creating activities where the value of an activity is linear in the total effort contribution that the activity receives. We assume that this value is equally divided between all the contributors who have contributed at least the threshold, which is a fraction from the maximum contribution to the activity. This allows us to model activities like authoring papers, books or participating in a start-up. We find that for two participants all the Nash equilibria are socially efficient, while for more partici-

pants, some Nash equilibria are inefficient and regulation is often required to motivate the agents to divide their efforts in a socially efficient manner.

Consider competitive activities like publishing papers. A paper usually needs to be of a certain level to get published at a given venue, because a venue may have a quota on the number of published papers (resulting from the acceptance rate and the total number of submissions) or a minimum required level. The eternal question is: How can the venue guarantee the existence of socially efficient equilibria in the game of dividing effort between papers? We find that having a predefined minimum publication value is preferable to imposing a quota on the number of papers that are published, since in order to guarantee the existence of socially optimal Nash equilibria, the latter requires certain constraints on the effort budgets of the authors. Generally, we model and study activities where high enough a value has to be achieved in order to survive and actually attain their value.

The second kind of activities we model is *reciprocal interactions*. This means interactions where an agent (a person, government, etc.) acts on the other agent reacting on what the other agent has done to the acting person. This is, for example, the dynamics in an arms race and interpersonal quarrels. We first prove that such interactions stabilize around some limit value exponentially fast. In many cases, we also provide closed formulas for the limit of the actions. The limit mostly depends on the agents who act on more agents and react on how the other agents act on her less, i.e. which are stable.

We then ask how a smart agent can maximize her utility from such an interaction, defined as what she receives minus the effort incurred by her own actions in the limit. Since people often act on habits, as Kahneman describes in “Thinking, Fast and Slow”, we model that the agents always reciprocate, but they can choose their habits of reciprocation. A *habit* is represented by a parameter that defines how an agent reacts on the others’ actions. We prove that letting the kinder (inherently more positive) agents persist with acting kindly while letting the less kind agents react on the actions of the kind ones is beneficial for the acting agents and for the whole society. This is case if acting is easy; if it is hard, the kind agents should follow the less kind ones. Therefore, when acting is easy, the personal interests coincide with the social one and therefore, every Nash equilibrium is socially efficient. However, if we allow defining several habits simultaneously (by setting several parameters), then also the less efficient equilibria become possible, requiring wise choice and exemplifying that “where there is great power there is great responsibility”, like Churchill said.¹

Finally, we model dividing effort between several reciprocal interactions. We consider the cases: *a*) when no contribution threshold exists, *b*) when achieving a contribution threshold is required to enjoy the interaction, but everyone may interact, and *c*) when even interacting is possible only if the contribution threshold is achieved. In each case, we study the possible Nash equilibria and their social welfare.

To predict more and provide better advice, several directions seem promising. First, real agents often participate in value-creating, interactive and perhaps other activities. Our model would become more realistic from modeling these activities in one game. Additionally, we would like to model the influence of agents’ participation in one activity on how the same agents participate in another activity. Many real situations require further

¹This quote dates back to the French National Convention, 08/05/1793.

extensions to the models, such as the fact that colleagues come and go requires having a dynamically changing set of agents. Second, we may consider other participation models, such as the current participants in an activity voting on which other agents also may participate in this activity. The lack of omniscience in life motivates modeling the dynamics of knowledge about activities. Finally, real life requires considering many other aspects of participating in activities, such as the psychological appeal of advice and legal and social constraints on what people do.

We lay the foundation of realistic mathematical modeling and analysis of effort division between activities. The above mentioned future research directions can further facilitate decisions on effort division.

SAMENVATTING

DEELNAME EN INTERACTIE IN PROJECTEN

VANUIT HET OOGPUNT VAN SPELTHEORIE

Mensen, overheden en bedrijven moeten vaak hun tijd of inspanningen verdelen tussen meerdere activiteiten. Om te beslissen hoe we de inspanning het beste kunnen verdelen, hebben we hulp nodig. Interessante voorbeelden van dergelijke activiteiten zijn: Wikipedia beschrijvingen maken, artikelen en boeken schrijven, bestanden op internet delen of met collega's discussiëren. In het voorbeeld van de tijdverdeling tussen Wikipedia uitbreiden, wetenschappelijke artikelen schrijven en een boek schrijven, kan de beslissing vrij ingewikkeld zijn. Bijvoorbeeld, neem aan dat een succesvol boek het meestal kostenefficiënte project zou zijn. Mijn inspanning zou niettemin te laag kunnen zijn om bij de lijst van de coauteurs terecht te komen. De mogelijkheid bestaat ook dat de andere auteurs zo weinig zouden bijdragen zodat de waarde van het boek te laag zou zijn. Voor deze ingewikkelde activiteiten die afhangen van wat de anderen doen, zoeken we gedrag dat persoonlijk rationeel is en tegelijkertijd ook winst oplevert voor de gehele groep.

We verdelen onze inspanningen vaak tussen verschillende activiteiten. Daarom is het belangrijk om dit te bestuderen, zodat we een efficiënte inspanningsverdeling kunnen vinden. Formeel, willen wij stabiele verdelingstrategieën aanraden aan de samenleving die leiden naar een hoge sociale welvaart. *Sociale welvaart* staat voor het totale nut dat de hele groep van de betrokken agenten verkrijgt (mensen, overheden, enzovoort). Stabiliteit betekent dat geen agent zijn eigen nut kan verhogen door iets nieuws te doen, terwijl de andere agenten blijven doen wat ze nu doen. Dit is het beroemde Nash-evenwicht. Aangezien een stabiele verzameling van strategieën (het gedrag van iedereen) een redelijke beschrijving is van het echte leven, zoeken wij een Nash-evenwicht met de hoogste sociale welvaart mogelijk.

We richten ons op spellen die inspanningsverdeling modelleren tussen twee vaak voorkomende activiteitsoorten: *waardecreërende* activiteiten zoals een artikel schrijven, en *interactie*activiteiten zoals het communiceren met collega's. Voor elk spel bestuderen wij de mogelijke Nash-evenwichten en hun sociale welvaart. In het geval dat alle Nash-evenwichten (er kunnen meerdere van zijn) *sociaal efficiënt* blijken (dat staat voor evenwichten die sociale welvaart hebben die ongeveer 1 geven als we door de maximale mogelijke sociale welvaart delen), is er geen regeling nodig. Als er zowel efficiënte en niet-efficiënte evenwichten bestaan, kan het zijn dat we de deelnemers ervan willen overtuigen om voor de sociaal efficiënte evenwichten te kiezen. In het onfortuinlijke geval dat er alleen niet-efficiënte evenwichten bestaan, is het mogelijk wenselijk de gehele situatie te beïnvloeden door, bijvoorbeeld, bepaalde strategieën te subsidiëren.

Ten eerste bestuderen wij de waardecreërende activiteiten waar de activiteitswaarde lineair is in de totale inspanningsbijdrage die de activiteit ontvangt. We nemen aan dat

deze waarde gelijk verdeeld wordt tussen alle bijdragers die ten minste de drempel hebben bijgedragen. De drempel is een fractie van de maximale bijdrage aan de activiteit. Dit maakt het mogelijk om activiteiten als artikelen en boeken schrijven, of aan een start-up deelnemen te modelleren. We vinden dat in het geval van twee deelnemers alle Nash-evenwichten sociaal efficiënt zijn. In het geval van meer deelnemers, daarentegen, zijn sommige Nash-evenwichten niet-efficiënt en is regeling vaak nodig om de agenten te motiveren om hun inspanningen te verdelen op een sociaal efficiënte manier.

Neem bijvoorbeeld competitieve activiteiten, zoals het publiceren van artikelen. Een artikel moet gewoonlijk op een bepaald niveau zijn om gepubliceerd te worden door een gegeven conferentie, omdat een conferentie een quotum mag hebben op het aantal publicaties (een resultaat van de acceptatiegraad en het totale aantal commissies) of een minimaal vereist niveau. De eeuwige vraag is: Hoe kan een conferentie zorgen dat er sociaal efficiënte evenwichten bestaan in het spel van het inspanningsverdelen tussen artikelen? We vinden dat een minimale publicatiewaarde hebben voordeliger is dan het opleggen van een quotum met betrekking tot het aantal publicaties. De reden is: om het bestaan van een sociaal optimaal Nash-evenwicht te garanderen, vereist de tweede optie bepaalde beperkingen aan de inspanningsbegrotingen van de auteurs. Samenvattend: wij modelleren en bestuderen activiteiten die een voldoende hoge waarde moeten hebben om te overleven en werkelijk hun waarde te realiseren.

Het tweede soort van activiteiten dat we modelleren is *wederzijdse interacties*. Dit zijn activiteiten waar een agent (een persoon, een overheid, enzovoort) reageert op wat een andere agent heeft gedaan. Dit is, bijvoorbeeld, de dynamiek binnen een wapenwedloop of persoonlijke ruzies. Eerst bewijzen wij dat deze interacties zich rond een bepaalde limiet stabiliseren, en ze doen dit exponentieel snel. Vaak geven we een gesloten formule voor de limieten van de acties. Deze limiet hangt in grotere mate af van de agenten die reageren op meerdere anderen en van agenten die minder heftig reageren op de acties van de anderen (dus, die zich stabiel gedragen).

Daarna vragen we hoe een slimme agent haar nut van zo'n interactie kan maximaliseren. Het nut is gedefinieerd als de gekregen actie minus de kost van haar eigen acties, alles in de limiet. Aangezien mensen vaak hun gewoontes volgen, zoals Kahneman in "Thinking, Fast and Slow" beschrijft, modelleren wij dat de agenten altijd zo'n wederzijds gedrag vertonen (reciprocal behavior), maar ze kunnen hun gewoontes qua reacties kiezen. Een *gewoonte* is vertegenwoordigd door een parameter die definieert hoe een agent op de acties van de anderen reageert. Wij bewijzen dat als de aardigere (inherent meer positieve) agenten consequent positieve acties ondernemen terwijl de minder aardige agenten op de acties van de aardigere reageren, is dit gedrag nuttig zowel voor de individuele agenten als voor de hele groep. Dat geldt als actie ondernemen makkelijk is; anders zouden de aardigere agenten de minder aardige volgen. Daarom valt, als actie ondernemen makkelijk is, de persoonlijke interesse samen met die van de hele groep en daardoor is elk Nash-evenwicht sociaal efficiënt. Echter, als we meerdere gewoontes tegelijkertijd definiëren (door middel van meerdere parameters zetten) worden dan ook de minder efficiënte evenwichten mogelijk. Dit vereist wijs kiezen, want zoals Churchill zei "Waar grote kracht is, is er tevens grote verantwoordelijkheid".²

Ten slotte modelleren we de inspanningsverdeling tussen een aantal wederzijdse

²Dit citaat komt oorspronkelijk van de Franse Nationale Conventie, 08/05/1793.

interacties. We bestuderen de volgende opties: *a*) er bestaat geen bijdragedrempel, *b*) drempel bereiken is noodzakelijk om van de interactie te profiteren, maar allemaal mogen interacteren, en *c*) zelfs interacteren is alleen mogelijk als de bijdragedrempel bereikt is. Voor elke optie, bestuderen wij de mogelijke Nash-evenwichten en hun sociale welvaart.

Om meer te kunnen voorspellen en een beter advies te kunnen geven, lijken een aantal onderzoeksmogelijkheden veelbelovend. Ten eerste nemen echte agenten vaak deel aan waardecreërende, interactieve en misschien nog andere activiteiten. Ons model zou realistischer kunnen worden door deze activiteiten in hetzelfde spel te modelleren. Daarnaast zouden we de invloed willen modelleren van de afhankelijkheid van deelname van een agent van deelname in andere activiteiten. Veel situaties in de praktijk vereisen verdere uitbreidingen van het model; bijvoorbeeld, het komen en gaan van collega's vereist een dynamische verzameling van agenten. Ten tweede zouden wij andere deelnamemodellen kunnen bestuderen, zoals de bestaande deelnemers aan een activiteit die stemmen op welke andere agenten ook mogen deelnemen aan deze activiteit. Het gebrek aan alwetendheid in de werkelijkheid motiveert het modelleren van de dynamiek van de kennis over activiteiten. Ten slotte vereist het reële leven dat er rekening gehouden wordt met veel andere aspecten van deelname aan activiteiten, zoals de psychologische aantrekkelijkheid van advies en de wettelijke en de sociale beperkingen aan wat mensen doen.

We leggen de basis van realistisch wiskundig modelleren en analyse van inspanningsverdeling tussen activiteiten. De bovengenoemde mogelijkheden voor toekomstig onderzoek zouden beslissingen over inspanningsverdeling verder kunnen ondersteunen.

PREFACE

This work tackles some of the game theoretic aspects of agents dividing their efforts between activities we call projects and enjoy the fruit these projects yield. We concentrate on two major kinds of activities: *value-creating* activities like writing books and activities of *reciprocal interaction*, such as interaction between colleagues of even nations. For each sort of activities, we study the stable effort division strategies (Nash equilibria) and their efficiency. The overarching goal is to facilitate decision support about how to divide effort.

This research was performed at the Algorithmics group in the department of Software and Computer Technology, which resides in the faculty of Engineering, Mathematics and Computer Science (EEMCS) of Delft university of technology. This work constitutes the game-theoretic pillar of the SHINE project.³ SHINE means Sensing Heterogeneous Information Network Environment and it supports self-organizing agents in acquiring on-demand information from heterogeneous sources, like sensors or reports made by people, and reporting the gathered information in the appropriate form. Our work supports the project in the following aspects:

1. We facilitate understanding under what conditions people will invest their free time in the SHINE project by studying value-creating projects.
2. Since the participants in SHINE can both request and provide information, this forms an interaction network, which we study in detail. Namely, we predict the reciprocal interaction in such a network and suggest which habits are beneficial for the individuals and the society. We also consider agents participating in several such interactive projects.

Gleb POLEVOY
Delft, November 2016

³See <http://www.participatorysystems.nl/2013/02/shine/>, <http://mmi.tudelft.nl/~birna/projects/projects.html>, and <http://shine.tudelft.nl/home/>.

1

INTRODUCTION

People are meant to help one another, like a hand assists a hand, a leg assists a leg and one jaw assists the other one.

Marcus Aurelius

This chapter presents the problem of efficient behavior in sharing effort in projects and obtaining utility from the projects. The background is presented, together with the main practical motivation for the research and the gaps it fills. We also describe the relevance to the SHINE project for self-organizing information acquiring, which is an interesting concrete case where the theory applies. Then, we pose our research questions, and present our main contributions, together with the structure of the thesis. A reading guide closes the chapter.

1 Everything people, nations, companies, robots, computer programs and even you, the reader, do constitutes dividing own time or effort between several projects. A project is an abstraction for activity, which involves several agents, benefiting them all. For example, a person may work with her colleagues on a study, socialize with friends for leisure, etc. Other examples include programming projects like Linux [1], writing for Wikipedia [2], crowdsensing projects [3], co-authoring articles [4], manufacturing cars, playing sports together, or even driving on the same roads. Every common project constitutes an interaction between the contributors, because their utilities depend on each other's contributions. It is crucial to make these ubiquitous interactions more efficient. We aspire to do this by advising the agents themselves or their manager how to contribute own resources like time and effort to the projects so as to maximize own utility and the total utility.

Just imagine the difference between a smart person who manages his time well and a person who does not, or between a well organized distributed system and a messy one. Predicting such interactions is important for deciding how much to contribute. In addition, making such advice automatic is important both for devising decision support systems and for implementing artificial agents that interact with people.

Among the multitude of shared effort projects, we concentrate on projects that *yield a revenue to be shared*, and on projects that *are reciprocal interactions*, i.e. interactions where agents react to others' actions. The first kind represents a simple case of creating a common product and dividing its value between the creators. Such a project can be, for instance, a project at work, a common homework, an article [4], a book or Wikipedia [2] (the utility of Wikipedia is the community's recognition). We concentrate on project with a minimum contribution threshold, such that only the agents who contribute at least this threshold receive a share. In practice, these projects also face requirements: enterprises need to achieve a minimum profit to survive, papers need to receive a minimum grade from the reviewers to get published, etc. We study projects with and without such requirements. There is no analysis of the most efficient ways to contribute to such thresholded projects, so we advise how much to contribute and how to organize the whole process, to improve the total well-being. The second kind of projects stands for many sorts of communication where people are involved, such as politics [5–7] or relationships with friends and family [8]. We concentrate on the ubiquitous reciprocal interactions, meaning reactive interactions [9, 10]. Since there is no analysis of how such a process will unfold and what reciprocation habits are most expedient, we model and analyze the reciprocation process and strategic choice of habits for this process. We also advise on dividing one's efforts between several such reciprocal interactions.

In any kind of projects, assuming that the agents decide rationally on how to divide their time among the projects and, in the case of reciprocation, also how to reciprocate, we employ game theory to model and analyze these interactions. Game theory is a mathematical approach to study interaction [11], which analyzes rational agents at its core (though there exist other branches of game theory as well; see [12] for a nice primer, which is a broad, though not recent, overview). We need game theory, since game theory allows for rigorous analysis and crisp conclusions, and we need crisp results to implement decision support systems. The central model of interaction in game theory is a *game* [11, Section 2.1.1], where each of the participants, called *players* or *agents*, chooses

a way of behavior, called *strategy*. The simultaneously chosen strategies of all the players together are more or less preferable to a player, and this is modeled by a real *utility* function that a player obtains. A central concept in game theory is a *Nash equilibrium* (NE) [13], which is a strategy profile that is stable in the sense that no player is able to strictly improve her utility by deviating, if the others keep their strategies. In order to predict the situations that will be played and advise on which strategies are individually and socially efficient, we model the interactions as games, and look at their Nash Equilibria and at how large the utilities of the group in these equilibria are relatively to the maximum utility the group can achieve. Analyzing this allows either stating that no regulation is required, or, on the other hand, suggesting to the manager (e.g., the boss of several interacting colleagues) or the government which equilibrium the agents should choose, to maximize the utility of the group.

Sometimes, game theory recommends to contribute to projects in some way. There is much research on methods to influence contributions, such as online community participation [14] or the motives for contributing to Wikipedia [2]. This research gives hope that sociologists will suggest the ways to steer agents to act in the recommended way.

In this chapter, we first present the existing work necessary for grasping the contributions of the thesis. This includes the area we concentrate on, which is shared effort in public projects, and an important sort of common activities we elaborate on, namely reciprocal interactions, where agents repeatedly react to others' past actions. We describe an important application of our theory to the SHINE¹ project for acquiring information by self-organizing agents. Then, we pose the research questions this thesis aims to answer. We then present and briefly discuss our contributions and their meaning. Finally, we provide an efficient reading guide to the reader. We provide a brief introduction to game theory in the appendix of the chapter.

1.1. RELATED WORK

We are interested analyzing and predicting the behavior of agents who invest effort in projects and benefit from them. We concentrate on the following projects:

Shared effort games. The agents' contributions to these projects create a value, which is subsequently shared between the contributors.

Reciprocal interactions. Instead of being economically rational, people tend to adopt other ways of behavior [12, 15], not necessarily maximizing some utility function. People tend to reciprocate, i.e., react on the past actions of others [9, 10, 16, 17]. Since reciprocation is ubiquitous, we study reciprocation as a common project.

1.1.1. SHARED EFFORT IN COMMON PROJECTS

We first present an overview of motives to contribute to projects and of several aspects of contribution and dividing the revenue. Then, we present several related models and their analysis, concluding that no analysis of the general setting has taken place, a gap which we partially fill.

¹See <http://www.participatorysystems.nl/2013/02/shine/>.

1 Motivation to contribute to public projects, such as cooperative software development [18] or quality control in crowdsourcing [19], has been studied. These studies, especially those concentrating on concrete motivational techniques, are necessary to implement our recommendations. Wang et al. [20] model motivation to contribute to online traveling communities and conclude that both practical motives, such as supporting others, building relationships, or hoping for a future repay, as well as internal drives to participate are important. Forte and Bruckman [2] discuss why people contribute to Wikipedia. By questioning contributors, Forte and Bruckman conclude that the reasons are similar to those of scientists and include the desire to publish facts about the world. Bagnoli and Mckee [21] empirically check when people contribute to a public good, like building a playground. They find that if people know the threshold for the project's success and benefit from collective contributing, then they will contribute, in agreement with the theory of [22]. Bagnoli and Mckee argue that knowing such information is realistic, giving real cases of hiring a lobbyist and paying to a ski club as evidence. This conclusion supports the rationality assumption. The rewards people obtain from such contributions can be both extrinsic, like a payment and a record for the CV, and intrinsic, such as exercising one's favorite skills [23]. The concrete ways to motivate contribution are discussed as well. For instance, Harper et al. [24] find that explicitly comparing a person's contribution to the contribution that others provide helps focusing on the desired features of the system, but does not change the interest in the system per se. The influence of revealing how much people contribute to a movie rating community is experimentally studied in [14]. Initiating participation in online communities is experimentally studied in [25] on the example of the influence of similarity and uniqueness of ratings on participation. Such studies and more are necessary to implement recommendations about contribution.

We now discuss various ways of dividing a project's revenue. Sometimes, the agent who contributes the most obtains all the revenue (like in political campaigns [26]), sometimes, every agent obtains a revenue, roughly proportional to her contribution (this may take place in writing columns to a newspaper, every columnist receiving the part of the newspaper's fame, proportional to her contribution), while it can be that everyone obtains an equal share (for example, when constructing a public facility or co-authoring papers). Division of a value is thoroughly researched in the surplus sharing literature, such as the classical *Shapley value* [27], or [28]. Unlike this field that devises the division rules, we take division rules as given and analyze the agents' strategies in the resultant game.

We now present existing models of project contribution and dividing the revenue. While modeling effort sharing as a game, we were inspired by the effort market game model of Bachrach, Syrgkanis and Vojnović [29], where each agent divides her budget between the projects available to her, and subsequently all the contributors obtain certain shares of the project's revenue. This model, though close to ours, does not allow for a minimum contribution threshold. A more constrained model, called all-pay auction, consists of a shared effort game where only the contributor with the highest contribution obtains the project's value, while everyone pays. The equilibria of these games have been studied, for instance, by Baye, Kovenock and de Vries [30]. This work shows cases where each player obtains the expected payoff of zero, and where the winner obtains

the difference between the two highest valuations, while the rest obtains zero. All-pay auctions model lobbying, single-winner contests, political campaigns, striving for a job promotion (see e.g. [26]) and Colonel Blotto games with two players [31]. In the Colonel Blotto game, two generals divide their armies between battlefields, and at every battlefield, the larger force wins. The number of the won battlefields defines the utility of a general. In this model, there exists the maximum threshold at every battlefield, as the winner takes all. Roberson [31] analyzes the equilibria of this game and their expected payoffs. Any outcome is socially optimum, since this is a constant-sum game.

We now present what has been done for models that resemble ours. For a very specific case (N -approximate Vickrey conditions, which mean that every agent obtains at least a constant share of her marginal contribution), Bachrach et al. [29] bound the price of anarchy of shared effort games by the number of players. This work also shows upper bounds on the PoA for the case of convex project functions, where each player receives at least a constant share of its marginal contribution to the project's value. However, this condition does not hold when a threshold is introduced. Anshelevich and Hoefer [32] considered an undirected graph model, where the nodes are the players and each player divides its budget between its adjacent edges in minimum effort games (where the edges are the 2-player projects), each of which equally rewards both sides by measure of the project's success (i.e., duplication instead of division). Anshelevich and Hoefer prove the existence of equilibria, find the complexity of finding an NE, and find that the PoA is at most 2. A related setting of *multi-party computation games* appeared in [33]. There, the players are computing a common function that requires them to compute a costly private value, motivating free-riding. The work suggests a mechanism, where honest computation is an NE. This differs from our work, since Smorodinsky and Tennenholtz consider cost minimization, and the choice of the players is either honestly computing or free riding, no choice of projects.

To conclude, no equilibrium efficiency research has been done for sharing with a general threshold, and therefore we consider this important domain.

In some situations, such as economical investments, the projects obtain their modeled value only if they stand up to a competition. We consider two models for project competition: a quota or a minimum level-based success, which take place, for example, in the process of deciding whether to accept or reject a paper [34, 35]. A quota can be expressed in other ways, such as an acceptance rate. Since our thesis considers efficiency of equilibria, we naturally look into the influence of a quota or a minimum level on the prices of anarchy and stability. The price of anarchy is the ratio of the total utility of an NE with the least total utility and of the largest possible total utility. The price of stability is the ratio of the total utility of the socially best NE and of the largest possible total utility. Since the influence of competition on the efficiency of stable situations (NE) has not been studied in the context of projects with thresholds, we consider sharing effort with competition between such projects. This allows to better model investing effort in firms or investing time in a paper, since both the contributors to a project compete and a project has to receive enough, to be profitable at all.

1.1.2. RECIPROCAL INTERACTIONS

To predict reciprocation, we need a simple and yet powerful model for reciprocation. We now describe the two main streams of existent models of reciprocation and afterwards explain what we contribute.

Existent models of (sometimes repeated) reciprocation can be classified as either explaining existence or analyzing consequences. The following models consider the reasons for the existence of reciprocal tendencies, grouped by the nature of the reasons.

Direct evolution. The classical works of Axelrod [5, 36] consider discrete reciprocity and shows that it is rational for egoists, so that species evolve to reciprocate. Evolutionary explanation appears also in other places, such as [37, 38]. Axelrod and Hamilton [39] and Fletcher and Zwick [40] consider engendering reciprocation by both the genetical kinship theory (helping relatives) and by the utility from cooperating when the same pair of agents interact multiple times. Berg et al. [41] proves that people tend to reciprocate and considers possible motivations, such as evolutionary stability.

Indirectly evolved. Bicchieri [42, Chapter 6], explicitly considers the psychological and game theoretic aspects of norm emergence and the eventual game theoretic utility of behaving according to the norm. Van Segbroeck et al. [43] consider the evolution of fairness, and pursuing fairness as a motivation for reciprocation. The famous work of Trivers [44] shows, in much biological detail, that sometimes reciprocity is rational, and thus, people can evolve to reciprocate. He shows how various emotions related to altruism have evolved. For instance, moralistic aggression and guilt are considered as threats to cheaters. Suspicion has evolved to detect subtle cheating. He argues that people can find the balance between cheating and cooperating.

Strong reciprocity. Gintis [45, Chapter 11] considers discrete actions, discussing not only the rationally evolved tit-for-tat, but also reciprocity with no future interaction in sight, what he calls *strong reciprocity*. He models the development of strong reciprocity analytically, using societal evolutionary dynamics. Several possible reasons for strong reciprocity, such as a social part in the utility of the agents or expressing itself in emotions, are considered in [46].

Axiomatic. Reciprocal behavior is axiomatically motivated in [47], assuming agents care not only for the outcomes, but also for strategies, thereby pushed to reciprocate. Under their axioms, Segal and Sobel prove a representation theorem, saying, when the preferences can be captured by a unique linear combination of the outcome dependent utilities of the agents.

Another research direction assumes that reciprocal tendencies exist and analyzes what ways it makes interactions develop, i.e. the consequences of reciprocation. These models analyze reciprocal interactions by defining and analyzing a game, where the utility function of rational agents directly depends on showing reciprocation [9, 48–50]. The importance of reward/punishment or of incomplete contracts for the flourishing of reciprocal individuals in the society is shown in [10].

Since no analysis of uncurling of inborn reciprocity with time considers non-discrete interactions (unlike, say, the discrete one from Axelrod [5, 36]), we model how interactions evolve with time, given that people reciprocate, and analyze this process.

Having analyzed a given reciprocity process, we next consider fine-tuning reciprocity. Since people tend to act on habits [51], we concentrate on maximizing own utility by setting own habits of reciprocity. Finally, no analysis of participating in several reciprocity projects has been done, from the perspective of stable states and their efficiency.

1.2. SHINE PROJECT

This research aims to study the strategic aspects of SHINE.² SHINE, the flagship project of DIRECT (Delft Institute for Research on ICT at Delft University of Technology), builds a framework for receiving demands on heterogeneous information (environmental, social and urban), obtaining the required data and presenting the information to the requester. In order to be scalable, flexible, and safe from single failures, the system needs to be self-organizing. Self-organization requires the need to take into account the strategic aspect of the participants.

This work supports SHINE by concentrating on game theoretic aspects of crowdsensing, from motivation to participate in a crowdsensing project to the interpersonal dynamics between the participants. In order to analyze how to improve participation in crowdsensing projects, we model several projects where people can contribute to, making the projects obtain a value, which is subsequently divided between the contributors. Once people participate in our project, we want them to interact for everyone's benefit. First, we analyze a given reciprocal (reactive) interaction, and then, we model strategic choice of own habits, aimed to interact more efficiently. Finally, we model a person splitting her time and effort between several reciprocal interactions. In each modeled interaction, we look for Nash equilibria, which are situations where no person can strictly benefit by changing only her behavior, if the others keep behaving as before. These situations can be expected to sustain themselves, if they happen to occur. We look which equilibria are more efficient to the society, to facilitate decision support for choosing to which projects to contribute and which reciprocity habits to adopt.

The following research questions are relevant to any value-creating and reciprocal projects and in particular, for the goals of SHINE.

1.3. RESEARCH QUESTIONS

As mentioned at the beginning of the chapter, we model participating in projects and analyze the stable situations (NE) of this process and their efficiency. This allows for predicting the situation and advising on the more efficient ways to participate in the projects. Thus, the highest-level research question is:

What are the Nash equilibria in shared effort games and how efficient are they?

²See <http://www.participatorysystems.nl/2013/02/shine/>, <http://mmi.tudelft.nl/~birna/projects/projects.html>, and <http://shine.tudelft.nl/home/>.

1 Since we concentrate on two kinds of shared effort games, this question decomposes into the two following groups:

The first group related to projects that create a linear value. As we describe in Section 1.1.1, there is no analysis of the thresholded case, so we analyze it.

1. What are the Nash equilibria in shared effort games with equal sharing of a linear project's value to everyone who contributes above a threshold? How efficient are these equilibria?
2. What changes in the answer to question 1, if a project obtains its value only if it survives a competition between the projects?

Both questions allow for SHINE as a project. The second question models the minimum level SHINE has to attain in order to survive.

The second group is about reciprocal interaction. The first two questions direct one interaction but are required to eventually analyze dividing effort between several such projects.

3. As we explain in the end of Section 1.1.2, there is no simple model of inborn non-discrete reciprocal actions, so we model it and ask: In reciprocal interaction, what will the actions become in the long run?
4. Given the above model of reciprocation, we ask: Which habits³ of reciprocation prove to be most efficient in the long run?
5. Getting back to analyzing projects, the summarizing question is: What are the Nash equilibria in shared effort games where every project is reciprocal interaction? How efficient are these equilibria?

These questions allow to model the interaction between the participants in SHINE, as a particular case.

We now describe how we answer these questions.

1.4. THESIS STRUCTURE AND CONTRIBUTIONS

Much of this work is based on published papers. We have modified the original papers slightly to create a coherent story, and have moved the basic background to this chapter, while still keeping every chapter from Chapter 2 till including Chapter 6 self-contained.

We concentrate on two prominent classes of projects. First, we consider dividing effort between projects. A project's value, linear in the total received contribution, is divided between the contributors. Some projects, like paper co-authorship, possess a contribution threshold, necessary to receive a share. The necessary effort of mastering the interface and the basic rules of Wikipedia is an example of an absolute threshold [52], while assigning bonus points to students from homework exercises, where one needs to achieve at least some percentage of the best grade, to obtain the homework's credits is an example of a threshold, proportional to the investments in the project.

³For the sake of the presentation, we use the colloquial word "habit" instead of "behavioral pattern".

In another project class, every project is a reciprocation, where agents react on each other's actions. After analyzing what happens as such a reciprocation uncurls and which habits agents can choose to gain more utility, we put it in the context of shared effort, since most people have several interactions on their minds, not one. Here, unlike in co-authoring a paper, the value is not directly created and subsequently shared, but the agents obtain value while interacting.

We now present concrete results per chapter. In Chapter 2, we prove that shared effort games with linear project functions and equal sharing to those who contribute above the threshold always possess a mixed Nash equilibrium. For pure equilibria, we first provide sufficient existence conditions of an NE for continuous general shared effort games. For a thresholded game with linear project functions, we characterize the existence of an NE for two agents and provide several sufficient conditions for a general number of agents. Next, we analyze the efficiency of the NE. In order to analyze the case of more than two agents, we generalize the fictitious play, originally proposed by Brown [53], to the shared effort game, and simulate it to find Nash equilibria and their efficiency. To run this and to eventually check, whether a profile is an NE, we devise an $O(n \log n)$ best response algorithm for 2-project games. For two agents, we prove that the efficiency is at least half of optimum, so regulation is not really needed. For more agents, the efficiency drops sometimes to less than a half, so a regulation may be useful. This chapter is an extended version of paper [54], which was also presented at BNAIC'14 and at the 5th World Congress of the Game Theory Society. The full version is currently under submission to a journal. This chapter answers research question 1.

Next, in Chapter 3, we model competition between projects as either quota or success threshold and provide sufficiency results for the existence of an NE in this more refined model. We show that setting a success threshold is more powerful than setting a quota, in order to guarantee that an optimal profile can be an NE. We also see that the price of anarchy is low but the price of stability is high, so there are inefficient NE, while there exist also efficient ones, and therefore, regulation impelling the agents to act efficiently may be expedient. This chapter aims to answer research question 2.

The next step is to consider projects that are more complicated than those yielding a linear function of the total received contribution, which is equally divided to certain contributors. In Chapter 4, we analyze a public project of the development of a lengthy reciprocal interaction. We prove convergence, and in several cases we also find the limit of the actions. The results show that the interaction in the limit depends on agent's kindness if he is persistent in the following sense: reacts less to others, but acts according to her own will, and can act on many agents. The convergence results allude to behavioral styles and to cultures. This chapter appears as an extended abstract at [55] and was presented at MFSC'15 (collocated with AAMAS'15), and at BNAIC'15. Chapter 4 answers research question 3.

In Chapter 5, we define utilities in reciprocation and study which habits the participating agents can adopt to maximize their utilities. We characterize the NE of this game and find their efficiencies, expressed as prices of anarchy and stability. We show that when acting is easy enough, then the less kind agents should be more flexible and follow the behavior of the kinder ones, explaining why people often become more polite when they grow up. We also prove that when acting is easy enough, then an NE is

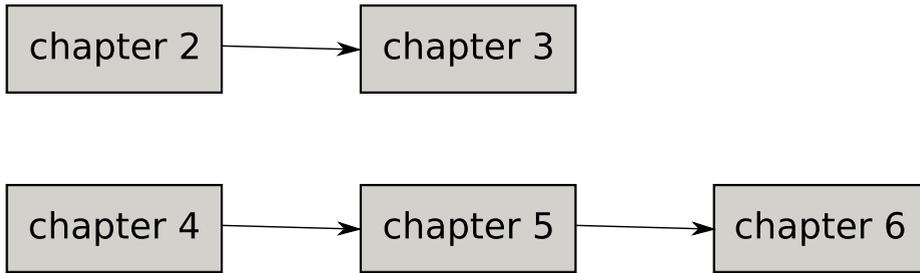


Figure 1.1: A chapter depends on another one, if it points there, directly or not.

optimal to the society, so selfishly reciprocating agents automatically benefit the whole group. This chapter is an extension of paper [56]. A part of this chapter was presented at AMEC/TADA'15 (collocated with AAMAS'15) and at BNAIC'16. This chapter answers research question 4.

Finally, in Chapter 6 we get back to our goal of analyzing shared effort. We model dividing own effort between several reciprocal interactions as a game. An agent's utility is what she obtains from the interactions where she participates. We prove that an equilibrium exists and find its efficiency, when no threshold for obtaining one's utility from reciprocation is present. With a threshold, partial existence and efficiency results are provided. We also consider the extended game where the agents first divide their time between interactions, and then choose the habits in every interaction. We provide sufficient conditions for the existence of a subgame perfect equilibrium (SPE). An SPE is a strategy profile that constitutes an NE at every state of an extensive game. We show that in the first game, without a threshold, any NE is optimal, so no regulation is required for the society. This chapter aims to answer research question 5.

We summarize the obtained results and discuss their implications in Chapter 7. We also propose some interesting directions for further work.

1.5. READING GUIDE

The best reading order is the appearance order in the thesis. However, the only real dependencies are depicted in the dag in Figure 1.1.

A reader in a hurry is advised to skip the following parts, because they are less central to the thesis:

1. The related work section, namely Section 6 and the simulation results, Section B, from the appendix of Chapter 2.
2. Chapter 3.
3. Section 4, Section 6, Section 7, and Section 8 from Chapter 4.
4. Section 5.2, Section 6.2, Section 10.2, and Section 7 from Chapter 5.
5. The thresholded cases, namely Section 5 and Section 6, and the extensive game, namely Section 7, from Chapter 6.

REFERENCES

- [1] *Linux website*, <https://www.linux.com/> (2016).
- [2] A. Forte and A. Bruckman, *Why do People Write for Wikipedia? Incentives to Contribute to Open-Content Publishing*, in *GROUP Workshop: Sustaining Community: The Role and Design of Incentive Mechanisms in Online Systems*. Sanibel Island, FL (2005) pp. 6–9.
- [3] R. Ganti, F. Ye, and H. Lei, *Mobile crowdsensing: current state and future challenges*, *IEEE Communications Magazine* **49**, 32 (2011).
- [4] J. Kleinberg and S. Oren, *Mechanisms for (mis)allocating scientific credit*, in *Proceedings of the 43rd annual ACM symposium on Theory of computing* (ACM, New York, NY, USA, 2011) pp. 529–538.
- [5] R. Axelrod, *The emergence of cooperation among egoists*, *American Political Science Review* **75**, 306 (1981).
- [6] W. J. Dixon, *Reciprocity in United States-Soviet relations: Multiple symmetry or issue linkage?* *American Journal of Political Science* **30**, pp. 421 (1986).
- [7] M. D. Ward, *Modeling the USA-USSR arms race*, *Transactions of The Society for Modeling and Simulation International* **43**, 196 (1984).
- [8] J. Gottman, C. Swanson, and J. Murray, *The mathematics of marital conflict: Dynamic mathematical nonlinear modeling of newlywed marital interaction*, *Journal of Family Psychology* **13**, 3 (1999).
- [9] A. Falk and U. Fischbacher, *A theory of reciprocity*, *Games and Economic Behavior* **54**, 293 (2006).
- [10] E. Fehr and S. Gächter, *Fairness and retaliation: The economics of reciprocity*, *Journal of Economic Perspectives* **14**, 159 (2000).
- [11] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*, MIT Press Books, Vol. 1 (The MIT Press, 1994).
- [12] A. Rubinstein, *Modeling Bounded Rationality*, 1st ed., Vol. 1 (The MIT Press, 1997).
- [13] J. Nash, *Non-Cooperative Games*, *The Annals of Mathematics Second Series*, **54**, 286 (1951).
- [14] Y. Chen, F. M. Harper, J. Konstan, and S. X. Li, *Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens*, *American Economic Review* **100**, 1358 (2010).
- [15] M. A. Raghunandan and C. A. Subramanian, *Sustaining cooperation on networks: An analytical study based on evolutionary game theory*, in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS'12 (2012) pp. 913–920.

- 1
- [16] W. Güth, R. Schmittberger, and B. Schwarze, *An experimental analysis of ultimatum bargaining*, *Journal of Economic Behavior & Organization* **3**, 367 (1982).
- [17] J. Sobel, *Interdependent preferences and reciprocity*, *Journal of Economic Literature* **43**, 392 (2005).
- [18] A. Begel, R. DeLine, and T. Zimmermann, *Social media for software engineering*, in *Proceedings of the FSE/SDP Workshop on Future of Software Engineering Research*, FoSER '10 (ACM, New York, NY, USA, 2010) pp. 33–38.
- [19] M. Allahbakhsh, B. Benatallah, A. Ignjatovic, H. R. Motahari-Nezhad, E. Bertino, and S. Dustdar, *Quality control in crowdsourcing systems: Issues and directions*, *IEEE Internet Computing* **17**, 76 (2013).
- [20] Y. Wang and D. R. Fesenmaier, *Assessing motivation of contribution in online communities: An empirical investigation of an online travel community*, *Electronic Markets* **13**, 33 (2003), <http://www.tandfonline.com/doi/pdf/10.1080/1019678032000052934>.
- [21] M. Bagnoli and M. McKee, *Voluntary contribution games: Efficient private provision of public goods*, *Economic Inquiry* **29**, 351 (1991).
- [22] M. Bagnoli and B. Lipman, *Provision of public goods: Fully implementing the core through private contributions*, *Review of Economic Studies* **56**, 583 (1989).
- [23] N. Kaufmann, T. Schulze, and D. Veit, *More than fun and money. worker motivation in crowdsourcing - a study on mechanical turk*, in AMCIS (2011).
- [24] F. M. Harper, S. X. Li, Y. Chen, and J. A. Konstan, *Social comparisons to motivate contributions to an online community*, in *Proceedings of the 2nd international conference on Persuasive technology* (Springer-Verlag, Berlin, Heidelberg, 2007) pp. 148–159.
- [25] P. J. Ludford, D. Cosley, D. Frankowski, and L. Terveen, *Think different: increasing online community participation using uniqueness and group dissimilarity*, in *Proceedings of the SIGCHI conference on Human factors in computing systems* (ACM, New York, NY, USA, 2004) pp. 631–638.
- [26] R. Siegel, *All-pay contests*, *Econometrica* **77**, 71 (2009).
- [27] L. S. Shapley, *A value for n-person games. Contribution to the Theory of Games*, *Annals of Mathematics Studies* **2**, 28 (1953).
- [28] H. Moulin, *On additive methods to share joint costs**, *Japanese Economic Review* **46**, 303 (1995).
- [29] Y. Bachrach, V. Syrgkanis, and M. Vojnović, *Efficiency and the Redistribution of Welfare*, Tech. Rep. (Microsoft Research, 2011).
- [30] M. R. Baye, D. Kovenock, and C. G. de Vries, *The all-pay auction with complete information*, *Economic Theory* **8**, 291 (1996).

- [31] B. Roberson, *The Colonel Blotto game*, *Economic Theory* **29**, 1 (2006).
- [32] E. Anshelevich and M. Hoefer, *Contribution games in networks*, *Algorithmica* **63**, 51 (2011).
- [33] R. Smorodinsky and M. Tennenholtz, *Overcoming free riding in multi-party computations – the anonymous case*, *Games and Economic Behavior* **55**, 385 (2006).
- [34] P. A. Bernstein, D. DeWitt, A. Heuer, Z. Ives, C. S. Jensen, H. Meyer, M. T. Özsu, R. T. Snodgrass, K.-Y. Whang, and J. Widom, *Database publication practices*, in *Proceedings of the 31st International Conference on Very Large Data Bases*, VLDB '05 (VLDB Endowment, 2005) pp. 1241–1245.
- [35] J. R. Douceur, *Paper rating vs. paper ranking*, *SIGOPS Oper. Syst. Rev.* **43**, 117 (2009).
- [36] R. Axelrod, *The evolution of cooperation*, Basic books (Basic Books, 1984).
- [37] W. Güth, M. Yaari, and U. Witt, *Explaining Reciprocal Behavior in Simple Strategic Games: An Evolutionary Approach*, (University of Michigan Press, Ann Arbor, 1992).
- [38] R. Sethi and E. Somanathan, *Preference evolution and reciprocity*, *Journal of Economic Theory* **97**, 273 (2001).
- [39] R. Axelrod and W. Hamilton, *The evolution of cooperation*, *Science* **211**, 1390 (1981), <http://www.sciencemag.org/content/211/4489/1390.full.pdf>.
- [40] J. A. Fletcher and M. Zwick, *Unifying the theories of inclusive fitness and reciprocal altruism*, *The American Naturalist* **168**, 252 (2006).
- [41] J. Berg, J. Dickhaut, and K. McCabe, *Trust, reciprocity, and social history*, *Games and Economic Behavior* **10**, 122 (1995).
- [42] C. Bicchieri, *The grammar of society: The nature and dynamics of social norms* (Cambridge University Press, 2006).
- [43] S. Van Segbroeck, J. M. Pacheco, T. Lenaerts, and F. C. Santos, *Emergence of fairness in repeated group interactions*, *Phys. Rev. Lett.* **108**, 158104 (2012).
- [44] R. L. Trivers, *The evolution of reciprocal altruism*, *The Quarterly Review of Biology* **46**, 35 (1971).
- [45] H. Gintis, *Game Theory Evolving: A Problem-centered Introduction to Modeling Strategic Behavior*, Economics / Princeton University Press (Princeton University Press, 2000).
- [46] E. Fehr, U. Fischbacher, and S. Gächter, *Strong reciprocity, human cooperation, and the enforcement of social norms*, *Human Nature* **13**, 1 (2002).
- [47] U. Segal and J. Sobel, *Tit for tat: Foundations of preferences for reciprocity in strategic settings*, *Journal of Economic Theory* **136**, 197 (2007).

- 1
- [48] J. C. Cox, D. Friedman, and S. Gjerstad, *A tractable model of reciprocity and fairness*, *Games and Economic Behavior* **59**, 17 (2007).
- [49] M. Dufwenberg and G. Kirchsteiger, *A theory of sequential reciprocity*, *Games and Economic Behavior* **47**, 268 (2004).
- [50] M. Rabin, *Incorporating fairness into game theory and economics*, *The American Economic Review* **83**, pp. 1281 (1993).
- [51] D. Kahneman, *Thinking, Fast and Slow* (Farrar, Straus and Giroux, 2011).
- [52] *Wikipedia:Ten Simple Rules for Editing Wikipedia*, https://en.wikipedia.org/wiki/Wikipedia:Ten_Simple_Rules_for_Editing_Wikipedia (2016).
- [53] G. W. Brown, *Iterative solution of games by fictitious play*, *Activity analysis of production and allocation* **13**, 374 (1951).
- [54] G. Polevoy, S. Trajanovski, and M. M. de Weerd, *Nash Equilibria in Shared Effort Games*, in *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14 (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2014) pp. 861–868.
- [55] G. Polevoy, M. de Weerd, and C. Jonker, *The convergence of reciprocation*, in *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '16 (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2016) pp. 1431–1432.
- [56] G. Polevoy, M. de Weerd, and C. Jonker, *The game of reciprocation habits*, in *Proceedings of the 2016 European Conference on Artificial Intelligence*, *Frontiers in Artificial Intelligence and Applications*, Vol. Volume 285: ECAI 2016 (2016) pp. 417–425.
- [57] E. Koutsoupias and C. Papadimitriou, *Worst-case equilibria*, in *16th Annual Symposium on Theoretical Aspects of Computer Science* (Trier, Germany, 1999) pp. 404–413.
- [58] C. Papadimitriou, *Algorithms, Games, and the Internet*, in *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing*, STOC '01 (ACM, New York, NY, USA, 2001) pp. 749–753.
- [59] A. S. Schulz and N. S. Moses, *On the performance of user equilibria in traffic networks*, in *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '03 (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2003) pp. 86–87.
- [60] E. Anshelevich, A. DasGupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden, *The price of stability for network design with fair cost allocation*, in *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on* (2004) pp. 295–304.
- [61] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, *Algorithmic Game Theory* (Cambridge University Press, 2007).

APPENDIX

A. A PRIMER ON GAME THEORY

We begin by introducing the general model, the model of non-cooperative games, followed by defining the Nash equilibria and their efficiency measures. Finally, we introduce the mixed Nash equilibrium.

A.1. NONCOOPERATIVE GAMES AND NASH EQUILIBRIUM

Game theory studies interaction between agents, like people, countries, companies, or robots. A game where several agents try to achieve something is a natural metaphor for an interaction. Concentrating on particular agents which may act on will and have their own interests usually expresses itself in an explicit modeling of agent's actions. This is treated in noncooperative game theory and modeled as a *noncooperative game*. Such a game consists of the following parts (see [11, Section 2.1.1]):

1. A set of players $N = \{1, 2, \dots, n\}$, which consists of the acting agents. We use “player” and “agent” interchangeably. A player symbolizes an independent entity.
2. Each player $i \in N$ has a set of possible strategies S_i ; a strategy $s_i \in S_i$ stands for a way of behavior.
3. Given the simultaneously choice of strategies by all the players s_1, s_2, \dots, s_n , the resulting situation may be more or less preferable to a given player. This preference is usually modeled by a personal utility function for each player, assigning a real value to every combination of the strategies of all the players. Formally, the utility of player i is $u_i: S_1 \times S_2 \times \dots \times S_n \rightarrow \mathbb{R}$. The larger a player's utility, the better for the player.

Analyzing such a game allows analyzing the modeled strategic situation. We concentrate on finding a Nash equilibrium (NE), in which rational players who do what they can in order to maximize own utility can stay stable. A Nash equilibrium [13] is a strategy profile $s = (s_1, s_2, \dots, s_n) \in S_1 \times S_2 \times \dots \times S_n = S$, such that no player can strictly improve her own utility by a unilateral deviation, when the others keep doing what they are doing. In formulas, the condition for $s \in S$ to be an NE is

$$\forall i \in N, \forall s'_i \in S_i: u_i(s) \geq u_i(s'_i, s_{-i}), \quad (1.1)$$

where⁴ $s_{-i} \triangleq (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$.

A.2. EFFICIENCY OF NASH EQUILIBRIA

We call the total utility of all the players at strategy profile $s \in S$ the “social welfare”, denoted $SW(s)$, meaning that $SW(s) \triangleq \sum_{i \in N} u_i(s)$. The Nash equilibria are the profiles that fulfill the condition from Formula (1.1). To quantify the loss in the social welfare, resulting from constraining the profiles to constitute an NE, the notions of price of anarchy (PoA) [57, 58] and price of stability (PoS) [59, 60] have been suggested. Formally, the

⁴We denote “is defined as” by \triangleq .

price of anarchy, PoA, of a game is defined as the ratio of the minimum social welfare in an NE to the maximum possible social welfare, i.e., $\text{PoA} \triangleq \frac{\min_{s \text{ is an NE}} \text{SW}(s)}{\max_{s \in S} \text{SW}(s)}$. To capture the other extreme, the price of stability is defined as the ratio of the social welfare in a best possible NE to the optimum possible social welfare, meaning that $\text{PoS} \triangleq \frac{\max_{s \text{ is an NE}} \text{SW}(s)}{\max_{s \in S} \text{SW}(s)}$. Intuitively, the price of anarchy quantifies the worse the society has to endure if its members do what they want, while the price of stability refers to the least evil the society has to take in order to be stable.

The fact that an NE can be inefficient is broadly known, such as in the famous example of the prisoner's dilemma [11, Example 16.2]. Since the introduction of price of anarchy [57] and price of stability [60], there have been many studies on the matter.

Roughgarden and Tardos [61, Chapter 17] discuss inefficiency of equilibria in non-cooperative games and consider the examples of network, load balancing and resource allocation games. The authors argue that understanding exactly when selfish behavior is socially profitable is important, since in many applications, implementing control is extremely difficult. Roughgarden and Tardos mention that the use of ratio of the objectives in an equilibrium and in the optimum to measure efficiency (PoA, PoS) constitute the two most popular approaches to choosing which equilibrium to use. Another possible approach is average-case analysis, being much more difficult to define and analyze. When defining the social good of an outcome, the authors mention that not only the sum of the costs, but also the maximum cost may be of interest. Roughgarden and Tardos also exemplify the potential function method for efficiency analysis.

A.3. MIXED NASH EQUILIBRIUM

A normal Nash equilibrium defined in Section A.1 is called *pure*. We now define a mixed NE. A *mixed extension* of a game is the same game where the strategies are all the probability distributions on the strategies of the original game. Formally, if player i 's original strategies are S_i , then her mixed strategies are $\{\alpha^j\}^{j \in S_i}$, where all α^j are nonnegative and $\sum_{j \in S_i} \alpha^j = 1$. One can think of the players picking their (pure) strategy according to this distribution. A Nash equilibrium of a mixed extension of a game is called a *mixed Nash equilibrium* of the original game.

The existence of a pure NE implies that of a mixed one, since any pure NE is also (naturally identified with) a mixed one. When a pure NE may not exist, the question of the existence of mixed NE becomes interesting. Therefore, we look into mixed NE, when a pure one may not exist.

2

SHARED EFFORT GAMES

Don't say you don't have enough time. You have exactly the same number of hours per day that were given to Helen Keller, Pasteur, Michelangelo, Mother Teresa, Leonardo da Vinci, Thomas Jefferson, and Albert Einstein.

H. Jackson Brown Jr., 1991

Shared effort games model strategic settings where people invest resources in public projects and the subsequent share of obtained profits is defined in advance. Such games model both projects like writing for Wikipedia, where everyone who knows the development environment shares the resulting benefits, and all-pay auctions such as contests and political campaigns, where only the winner obtains a profit. In θ -(equal) sharing (effort) games, the threshold θ for effort defines which contributors win and then receive their (equal) share. For public projects $\theta = 0$ and for all-pay auctions $\theta = 1$. Thresholds between 0 and 1 can model games such as paper co-authorship and shared homework assignments, where a minimum positive contribution is required before sharing in the profits. We constructively characterize the conditions for the existence of a pure equilibrium for two-player shared effort games with project value functions that are linear in the received contribution and find the prices of anarchy and stability. We provide some existence and efficiency results for more players as well. In the mixed case, we prove that an equilibrium always exists and provide results on their social welfare. For more players, generalized fictitious play simulations are used to show when a pure equilibrium exists and what its efficiency is. The found equilibria provide the likely strategy profiles and the socially preferred strategies regarding contributing to public projects. This facilitates setting socially efficient equilibria.

This chapter is an extended version of paper [1].

1. INTRODUCTION

We begin our journey by analyzing the common real-world situations that include a group of people investing resources in several public projects. The *revenues from these projects are typically divided* based on the investments. Examples of such situations include contributions to online communities [2], Wikipedia [3], political campaigns [4], paper co-authorship [5], or social exchange networks [6]. Projects can be coding Linux, cleaning the house, and arguing for an important decision. In this chapter we consider situations where the obtained revenue from such projects is shared equally, but possibly only among those who contribute at least a certain amount. This threshold can be either absolute or relative to the other investments. We concentrate on a relative threshold.

Assigning points for an exercise, where a percentage of the perfect work is required to obtain the (equal) homework's credits [7] is an absolute threshold example. In this example, participants' utilities (i.e., their gain) are typically equal for anyone who is above the threshold. An example with a relative threshold is the Colonel Blotto game (see e.g. [8]), where there are several battlefields over which the Colonel can distribute its forces, and the number of local victories determines the payoff. This example is "highly thresholded" because only the player whose effort per project (a battlefield) is maximum collects the complete revenue. Another example of a relative threshold is an employee that obtains decision power or a bonus if she has invested sufficient effort [9]. The following example is used later to further illustrate the model.

Example 1. *Consider two collaborating scientists in a narrow field. They can work on their papers alone or together. When they collaborate on a paper (a project, in this setting), an author has to contribute at least 0.2 of the work of the other one, in order to be considered a co-author. The reward, being the recognition, is equally divided among the authors. Author 1 has the time budget of 5 hours to work, and author 2 has 20 hours. The value of the reward of the first paper is 4 times the total contribution it receives, while the less "hot" second paper rewards the contributors with only twice the received contribution. This is illustrated in Figure 2.1. In the figure, the first paper receives the total contribution of $4 + 10 = 14$, creating the value of $4 \cdot 14 = 56$. Both contributors are authors, since $4 \geq 0.2 \cdot 10$, and the value is equally divided between them. The second paper receives $1 + 10 = 11$, and yields the value of $2 \cdot 11 = 22$. Here, only the contributor of 10 is an author, since $1 < 0.2 \cdot 10$, and he, thus, receives the whole reward of 22. This is not a Nash equilibrium, since the second contributor would benefit from moving the 1 hour contribution to the first paper. On the other hand, if both authors invest all their time in paper 1, the situation is stable. Indeed, moving a part to paper 2 would benefit nobody, since the paper is twice less profitable than paper 1, so sharing the value of paper 1 is as good as contributing alone to paper 2. The social welfare in this equilibrium is maximum possible, since everyone contributes to the most profitable project. In general, we would like to find stable contributions, and whether they will be efficient for both authors, relatively to the maximum possible division of the authors' time budgets.*

A similar example is that of scientists that are working in the same research area, but do not necessarily publish together. Their profit can be expressed in terms of publications, citations, and awards. This resembles Kleinberg and Oren [5], though they assume that a researcher may choose a single project, which revenue may be divided in various

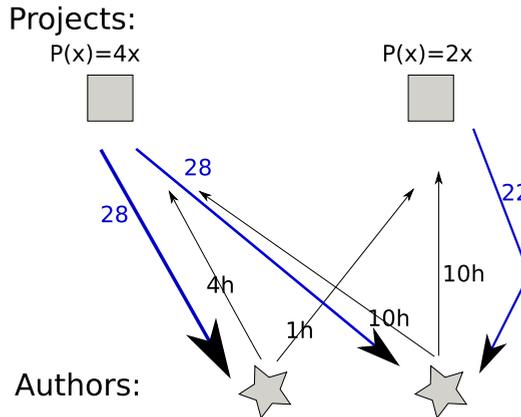


Figure 2.1: The co-authors invest what is shown in the arrows that go up, every project's revenue is defined as the P function of the total contribution, and it is equally shared among the contributors who contribute above the relative threshold of 0.2. The obtained shares are denoted by the arrows that go down.

manners.

People and organizations often invest resources in several projects and share the obtained revenues. It is thus important to predict stable contributions and suggest the efficient ones. To this end, we study the Nash equilibria of such games, the ratio of the least total utility of the players in an equilibrium to the optimum, called the price of anarchy (PoA) [10], and the ratio of the largest total utility in an equilibrium to the optimal total utility, called the price of stability (PoS) [11, 12]. If the price of anarchy is close to 1, then all equilibria are good, and we may suggest any equilibrium profile. If the price of anarchy is low, while the price of stability is high, then we have to regulate the play by suggesting the efficient equilibria, while if even the price of stability is low, the only way to make the play socially efficient is changing the game through subsidizing, etc.

This price of anarchy was bounded in [13], but assuming the k -approximate Vickrey condition, meaning that a player obtains at least $1/k$ th of her marginal contribution, which fails to hold in a positively thresholded model. There is no analysis of the existence of Nash Equilibria (NE) and their efficiency in *general* shared effort games.

This chapter aims to fill this gap by the following contributions:¹

1. a constructive characterization of the existence of pure strategy NE for two players and linear project values (Theorem 3),
2. the price of anarchy and stability of these equilibria (Theorem 5, Corollary 2),
3. sufficiency results on existence of pure NE for any number of players (Theorem 4),
4. the prices of anarchy and stability for any number of players (Theorem 6),

¹Relatively to [1], we extend the theory also for budgets not within a threshold factor from each other (part 1 in the list), substantially extend the simulations (part 9), and prove the existence of an NE in the mixed case, answering this natural question (part 5).

5. a proof that any shared effort game with linear project functions and equal (threshold) sharing has a mixed NE (Theorem 7),
6. providing some efficiency bounds on mixed equilibria,
7. a generalization of fictitious play² to shared effort (infinite) games (Definition 4),
8. an $O(n \log n)$ best response algorithm for 2-project multi-player games (Theorem 8),
9. simulation of fictitious play to find pure Nash Equilibria in 2-project multi-player games, and if an equilibrium is found, a report of its efficiency (Section 5.4).

We assume pure NE, unless explicitly mentioned otherwise. A pure shared effort game is already uncountably infinite and non-continuous. After defining shared effort games in the next section, we concentrate on pure strategies in Section 3. We theoretically treat the existence and efficiency of NE for the games with two players and linear project values in Section 3.1. Afterwards, we study games with any number of players. We treat the mixed extension in Section 4, proving the existence of a mixed NE and showing some efficiency bounds of the pure case generalize to the mixed case. Aiming to further investigate existence and efficiency of pure NE, we employ fictitious play simulations in Section 5. Concretely, we define an Infinite-Strategy Fictitious Play and simulate it till and if it practically converges within some time. We check whether we have found an NE and if that is the case, what its efficiency is. We describe the related work in Section 6. We conclude and discuss the future work in Section 7. A short primer on mixed NE appears in Section A.3.

2. MODEL

To model investing effort in shared projects, we define *shared effort games*, which also appear in [13]. The games consist of players who contribute to project, and share the value the projects obtains. Formally, players $N = \{1, \dots, n\}$ contribute to projects Ω . Each player $i \in N$ can contribute to projects in Ω_i , where $\emptyset \subsetneq \Omega_i \subseteq \Omega$; the contribution of player i to project $\omega \in \Omega_i$ is denoted by $x_\omega^i \in \mathbb{R}_+$; for any $\omega \notin \Omega_i$, we write $x_\omega^i = 0$. Each player i has a budget $B_i > 0$, and the strategy space of player i (i.e., the set of her possible actions) is³ $S^i \triangleq \left\{ x^i = (x_\omega^i)_{\omega \in \Omega} \in \mathbb{R}_+^{|\Omega|} \mid \sum_{\omega \in \Omega_i} x_\omega^i \leq B_i, \omega \notin \Omega_i \Rightarrow x_\omega^i = 0 \right\}$. Denote the vector of all the contributions by $x = (x_\omega^i)_{\omega \in \Omega}^{i \in N}$ and the strategies of all the players except i by x^{-i} .

To define the utilities, each project $\omega \in \Omega$ is associated with its *project function*, which determines its *value*, based on the total contribution vector $x_\omega = (x_\omega^i)_{i \in N}$ that the project receives; formally, $P_\omega(x_\omega): \mathbb{R}_+^n \rightarrow \mathbb{R}_+$, and every P_ω is increasing and differentiable, in every parameter. When the functions P_ω depend only on the $\sum (x_\omega^i)_{i \in N}$, we may write expressions like $P_\omega(x) = 2x$. The project's value is distributed among the players in $N_\omega \triangleq$

²The original fictitious play was proposed by Brown [14].

³We denote "defined as" by \triangleq .

$\{i \in N \mid \omega \in \Omega_i\}$ according to the following rule. From each project $\omega \in \Omega_i$, each player i gets a share $\phi_\omega^i(x_\omega): \mathbb{R}_+^n \rightarrow \mathbb{R}_+$ with free disposal:

$$\forall \omega \in \Omega: \sum_{i \in N_\omega} \phi_\omega^i(x_\omega) \leq P_\omega(x_\omega). \quad (2.1)$$

We assume that the sharing functions are non-decreasing.

The utility of a player $i \in N$ is defined to be

$$u^i(x) \triangleq \sum_{\omega \in \Omega_i} \phi_\omega^i(x_\omega).$$

The social welfare is defined as the total utility, i.e. $SW(x) \triangleq \sum_{i=1}^n u^i(x)$.

We now define a specific variant of a shared effort game, called a θ -sharing mechanism. This variant is relevant to many applications where a minimum contribution is required to share the revenue, such as paper co-authorship and homework, and we study predominantly such games. For any $\theta \in [0, 1]$, define the players who get a share as those who bid at least a θ fraction of the maximum bid size to ω ,

$$N_\omega^\theta \triangleq \left\{ i \in N_\omega \mid x_\omega^i \geq \theta \cdot \max_{j \in N_\omega} x_\omega^j \right\}.$$

The θ -equal sharing mechanism equally divides the project's value between all the users who contribute at least θ of the maximum bid to the project.

Definition 1. The θ -equal sharing mechanism, denoted by M_{eq}^θ , is

$$\phi_\omega^i(x_\omega) \triangleq \begin{cases} \frac{P_\omega(x_\omega)}{|N_\omega^\theta|} & \text{if } i \in N_\omega^\theta, \\ 0 & \text{otherwise.} \end{cases}$$

Reconsider the example from Section 1 to illustrate the above model.

Example 1 (Continued). The scientists N invest in the papers (projects) Ω . Assume that a paper's total value for the reputation of its authors is proportional to the total investment in the paper. That is, the project's functions P_ω are linear. In order to be considered an author, a minimum threshold θ of the maximum contribution is required, and a paper's total contribution to the authors' reputation is equally divided between all its authors. This is a shared effort game with a threshold $\theta \in (0, 1)$ and equal sharing.

3. PURE NASH EQUILIBRIUM

In this section, we analyze existence of stable profiles in shared effort games, i.e., the Nash equilibria. Having analyzed their existence, we analyze, how efficient they are for the society, relatively to the best possible profiles for the society. This allows predicting behavior and recommending what to do.

We begin with sufficiency results. Then, we concentrate on the case of linear project functions, characterizing the existence and efficiency of NE for two players. We follow with existence and efficiency results for any number of players. In this chapter, some proofs are deferred to Section A, to make the presentation flow.

Notice that our (pure) model is *not* a mixing of the model where a player may invest in at most one project, because a mixing would extend the profits linearly in the mixing coefficients, while this is not the case in our model with a positive threshold.

The following theorem proves an NE exists in the convex continuous case.

Theorem 1. *Let the strategy sets be non-empty, compact and convex. Then, if each ϕ_ω^i is continuous and concave, a pure NE exists. If we additionally assume that the strategy sets are equal to all the payers (in particular, all Ω_i s are the same) and that the utility functions are symmetric, then we also conclude that a symmetric NE exists.*

Proof. Immediate from Proposition 20.3 in [15] (and from Theorem 3 in [16], for the symmetric case). \square

For equal division without a threshold, the following existence result is stronger, besides that it does not consider symmetric NE.

Theorem 2. *The game with M_{eq}^0 admits a potential function. Therefore, if the functions P_ω are continuous and the strategy spaces are compact, then a pure NE exists.*

Proof. The strategy space of player i is S^i , and denote $S \triangleq S^1 \times \dots \times S^n$. Define $P: S \rightarrow \mathbb{R}$ by $P(x) \triangleq \sum_{\omega \in \Omega} \frac{P_\omega(x_\omega)}{|N_\omega|}$. This is a potential function, because it is equal to the utility of any player, and therefore, when player i changes her strategy, her utility changes exactly as the potential does.

The game possesses a pure NE, whenever the potential function admits the maximum. In our case, as these functions are continuous and the spaces are compact, they always achieve the maximum (see Lemma 4.3 in [17]). \square

Here, finding an NE would follow from finding a maximum of function $\sum_{\omega \in \Omega} P_\omega(x_\omega)$. This also maximizes the social welfare. For linear project functions, a profile is an NE if and only if all the agents contribute to the most profitable projects, dividing their budgets between those projects arbitrarily.

3.1. LINEAR PROJECT FUNCTIONS

We first study equilibrium existence, then moving to efficiency. We begin with two agents, subsequently generalizing it to any number of agents.

To be able to study the existence of Nash equilibria analytically, this section considers equal θ -sharing, where all the project functions are linear with coefficients $\alpha_m \geq \alpha_{m-1} \geq \dots \geq \alpha_1 > 0$ (the order is w.l.o.g.) We denote the number of projects with the largest coefficient project functions by $k \in \mathbb{N}$, i.e. $\alpha_m = \alpha_{m-1} = \dots = \alpha_{m-k+1} > \alpha_{m-k} \geq \alpha_{m-k-1} \geq \dots \geq \alpha_1$. We call those k projects *steep*. Assume w.l.o.g. that $B_n \geq \dots \geq B_2 \geq B_1 > 0$.

To proceed, we need some definitions. Given a strategy profile, we call a project that receives no contribution a *vacant* project. We define players that do not obtain a share from a given project as *dominated* at that project. We call them *suppressed* if they also contribute to that project. Formally,

Definition 2. *The dominated players at a project ω are $D_\omega \triangleq N_\omega \setminus N_\omega^\theta$, and the suppressed players at a project ω are $S_\omega \triangleq \{i \in N_\omega : x_\omega^i > 0\} \setminus N_\omega^\theta$.*

In an NE, a player is suppressed at a project if and only if it is suppressed at all the projects where it contributes. This holds since if a player is suppressed at project p but it also contributes to project $q \neq p$ and is not suppressed there, then it would like to move its contribution from p to project q .

We first assume 2 players, i.e., $n = 2$, and completely characterize this case. We introduce Lemmas 1, 2, and 3, before formulating and proving the characterization of an NE. These lemmas describe what must hold in any NE.

Lemma 1. *Consider an equal θ -sharing game with two players with $0 < \theta < 1$ and linear project functions.*

Then the following hold in any NE:

1. *At least one player contributes to a steep project.*
2. *Suppose that a non-suppressed player, contributing to a steep project, contributes to a non-steep project as well. Then, it contributes either alone or precisely the least amount it should contribute to achieve a portion in the project's value.*

The following lemma treats budgets that are close to each other.

Lemma 2. *Consider an equal θ -sharing game with two players with $0 < \theta < 1$ and linear project functions.*

If $B_1 \geq \theta B_2$, then the following hold in any NE.

1. *Each player contributes to every steep project.*
2. *A non-steep project receives the contribution of at most one player.*

We need another definition.

Definition 3. *A 2-steep project is a project that is most profitable among the non-steep ones.*

The following lemma treats budgets that are further from each other, than within the factor of the threshold.

Lemma 3. *Consider an equal θ -sharing game with two players with $0 < \theta < 1$ and linear project functions.*

If $B_1 < \theta B_2$, then the following hold in any NE where no player is suppressed.

1. *Player 1 contributes only to non-steep projects.*
2. *Each player receives a (strictly) positive utility, unless all projects are the same.⁴*
3. *Player 2 contributes alone to every steep project, and perhaps to a non-steep project together with 1, the threshold amount.*
4. *The non-steep and non-2-steep projects receive zero contribution.*

⁴That is, unless $k = m$.

5. If player 2 contributes to a 2-steep project, then there exists only a single 2-steep project.

We are finally ready to characterize the existence of an NE for two players.

Remark 1. From now on we assume that Ω_i includes all the steep projects, for each player $i \in N$, that is $\{m, m-1, \dots, m-k+1\} \subseteq \Omega_i, \forall i \in N$.

Theorem 3. Consider an equal θ -sharing game with two players with budgets B_1, B_2 . W.l.o.g., $B_2 \geq B_1$. Assume $0 < \theta < 1$, and linear project functions with coefficients $\alpha_m = \alpha_{m-1} = \dots = \alpha_{m-k+1} > \alpha_{m-k} \geq \alpha_{m-k-1} \geq \dots \geq \alpha_1$ (the order is w.l.o.g). This game has a pure strategy NE if and only if one of the following holds.⁵

1. $B_1 \geq \theta B_2$ and the following both hold.

- (a) $\frac{1}{2}\alpha_m \geq \alpha_{m-k}$,
 (b) $B_1 \geq k\theta B_2$;

2. $B_1 < \theta B_2$ and also at least one of the following holds.

- (a) $B_1 < \frac{\theta B_2}{k}$ and $\frac{\alpha_{m-k}}{\alpha_m} \leq \min\left\{\frac{1}{1+\theta}, \frac{2\theta}{1+\theta}\right\}$,
 (b) $B_1 < \frac{\theta B_2}{k+\theta^2}$ and $\alpha_{m-k} \geq 2\alpha_{m-k-1}$ and $\frac{2\theta}{1+\theta} \leq \frac{\alpha_{m-k}}{\alpha_m} \leq \frac{2(1-\theta)}{2-\theta}$ and $m-k$ is the only 2-steep project,
 (c) $B_1 < \frac{\theta}{|\Omega|} B_2$ and all the project functions are equal, i.e. $\alpha_m = \alpha_1$.

The idea of the proof is as follows. To show existence of an equilibrium under the assumptions of the theorem, we just provide a strategy profile and prove that no unilateral deviation is profitable. We show the other direction by assuming that a given profile is an NE and deriving the asserted conditions. To do this, we first use Lemmas 1, 2 and 3 that describe what holds in an equilibrium, in order to limit the possibilities for an equilibrium profile.

We prepend the following lemma with a technical statement.

Lemma 4. Consider an equal θ -sharing game with two players with budgets B_1, B_2 . W.l.o.g., $B_2 \geq B_1$. Assume $0 < \theta < 1$, and linear project functions with coefficients $\alpha_m = \alpha_{m-1} = \dots = \alpha_{m-k+1} > \alpha_{m-k} \geq \alpha_{m-k-1} \geq \dots \geq \alpha_1$ (the order is w.l.o.g). Assume that no player is suppressed anywhere, and player j does not contribute to a non-steep project p . Consider player $i \neq j$.

Then, the following hold.

1. If $\frac{1}{2}\alpha_m \geq \alpha_p$, then it is not profitable for i to move any budget $\delta > 0$ from any subset of the steep projects to p (or to a set of such projects).

⁵If α_{m-k-1} (and) or α_{m-k} does not exist, consider the containing condition to be vacuously true.

2. If $\frac{1}{2}\alpha_m > \alpha_p$, then it is (strictly) profitable for i to move any budget $\delta > 0$ from p to any subset of the steep projects. If j is suppressed after such a move, then requiring $\frac{1}{2}\alpha_m \geq \alpha_p$ is enough.
3. If $\frac{1}{2}\alpha_m < \alpha_p$ and it is possible to move $\delta > 0$ from any subset of the steep projects to p , such that i received and still receives half of the value of these steep projects, then it is (strictly) profitable for i .

We are now set to prove the theorem.

Proof. (\Rightarrow) We prove the existence of NE under the conditions of the theorem. We begin with case 1. Supposing that $B_1 \geq k\theta B_2$ and $\frac{1}{2}\alpha_{m-k+1} \geq \alpha_{m-k}$. Let both players allocate $1/k\theta$ of their respective budgets to each of the steep projects. We prove here that this is an NE. This profile provides each player with $k \cdot \frac{1}{2}\alpha_m \cdot \frac{B_1+B_2}{k} = \frac{1}{2}\alpha_m \cdot (B_1 + B_2)$. For any player i , moving $\delta > 0$ to some non-steep projects is not profitable, according to part 1 of Lemma 4. Another possible deviation is reallocating budget among the steep projects. Since $B_1 \geq k\theta B_2$, we conclude that $B_2 \leq \frac{B_1}{k\theta}$, so 2 is not able to suppress 1 (and the other way around is clearly impossible, even more so) and therefore, merely reallocating among the steep projects will not increase the profit. The only deviation that remains to be considered is simultaneously allocating $\delta > 0$ to some non-steep projects and reallocating the rest of the budget among the steep ones. Any such potentially profitable deviation can be looked at as two consecutive deviations: first allocating $\delta > 0$ to some non-steep projects, and then reallocating the rest of the budget among the steep ones. Part 1 of Lemma 4 shows that bringing back all $\delta > 0$ from non-steep projects to the steep ones, without getting suppressed anywhere (which is possible since $B_1 \geq \theta B_2$) will bear a non-negative profit. Therefore, we can ignore the last form of deviations. Therefore, this is an NE.

We now move to handle case 2. Case 2a: suppose that $B_1 < \frac{\theta B_2}{k}$ and $\frac{\alpha_{m-k}}{\alpha_m} \leq \min\left\{\frac{1}{1+\theta}, \frac{2\theta}{1+\theta}\right\}$. Let player 1 invest all its budget in $m-k$, and let 2 invest $\frac{B_2}{k}$ in each steep project. We prove this is an NE. The only possibly profitable deviation for player 1 is to invest in steep projects. However, since $B_1 < \frac{\theta B_2}{k}$, player 1 would obtain nothing from the steep projects. Also player 2 would not gain from a deviation, because first, from our assumption,

$$\frac{\alpha_{m-k}}{\alpha_m} \leq \frac{1}{\theta+1} \iff \alpha_m(B_1/\theta) \geq \alpha_{m-k}(B_1(1+1/\theta)),$$

and therefore, player 2 would not profit from suppressing player 1 at project $m-k$. Second, according to our assumption,

$$\frac{\alpha_{m-k}}{\alpha_m} \leq \frac{2\theta}{\theta+1} \iff \alpha_m(\theta B_1) \geq \frac{\alpha_{m-k}((1+\theta)B_1)}{2},$$

and therefore, player 2 would not profit from getting a half of the value of project $m-k$. Thus, no deviation is profitable. Therefore this is an NE.

Case 2b: suppose that $B_1 < \frac{\theta B_2}{k+\theta^2}$ and $\alpha_{m-k} \geq 2\alpha_{m-k-1}$ and $\frac{2\theta}{1+\theta} \leq \frac{\alpha_{m-k}}{\alpha_m} \leq \frac{2(1-\theta)}{2-\theta}$ and $m-k$ is the only 2-steep project. Let player 1 invest all its budget in $m-k$, and let 2 invest

θB_1 in $m - k$ and $\frac{B_2 - \theta B_1}{k}$ in each steep project. We prove that this is an NE. The possibly profitable deviations for player 1 is to invest in steep projects or in project $m - k - 1$. Here, we show them to be non profitable. First, since $B_1 < \frac{\theta B_2}{k + \theta^2} \iff B_1 < \theta \frac{B_2 - \theta B_1}{k}$, there is no profit for 1 from investing in a steep project. Second, according to our assumption,

$$\alpha_{m-k} \geq 2\alpha_{m-k-1} \iff \frac{\alpha_{m-k}(B_1(1+\theta))}{2} \geq \frac{\alpha_{m-k}(B_1(\theta^2 + \theta))}{2} + \alpha_{m-k-1}(B_1(1-\theta^2)),$$

and therefore, player 1 would not profit from investing in $m - k - 1$. Next, we show that also player 2 does not have incentives to deviate. Since we can assume that the non-2-steep projects do not receive any contribution and since the way how the contribution is divided between the steep projects does not influence the profit, the possible deviations to improve player 2's profit are transferring budget from the steep projects to $m - k$ or the other way around. We show now that they are not profitable. First, according to our assumption,

$$\frac{\alpha_{m-k}}{\alpha_m} \leq \frac{2(1-\theta)}{2-\theta} \iff \alpha_m(B_1/\theta - \theta B_1) + \frac{\alpha_{m-k}(B_1(1+\theta))}{2} \geq \alpha_{m-k}(B_1(1+1/\theta)),$$

and therefore, player 2 would not profit from suppressing 1 on project $m - k$. Second, according to our assumption,

$$\frac{\alpha_{m-k}}{\alpha_m} \geq \frac{2\theta}{1+\theta} \iff \frac{\alpha_{m-k}(B_1(1+\theta))}{2} \geq \alpha_m(\theta B_1),$$

and therefore, player 2 would not profit from moving θB_1 from $m - k$ to a steep project. We have shown that no deviation is profitable. Therefore this is an NE.

Case 2c: suppose that $B_1 < \frac{\theta}{|\Omega|} B_2$ and all the project functions are equal. Then, player 2 investing $\frac{B_2}{|\Omega|}$ in every project, and player 1 using any strategy is an NE. To see this, notice that player 2 obtains $\alpha_m(B_2 + \sum_{\omega \in \Omega} x_{\omega}^1)$, that is the maximum possible profit. Player 1 will be suppressed in any attempt to invest, and therefore has no incentive to deviate. Therefore this is an NE.

(\Leftarrow) We show the other direction now. We assume that a given profile is an NE and derive the conditions of the theorem. Assume that a given profile is an NE. We first suppose that $B_1 \geq \theta B_2$ and we shall derive that the conditions of 1 hold.

Since $B_1 \geq \theta B_2$, then according to Lemma 2, each player contributes to every steep project. Suppose to the contrary that $\frac{1}{2}\alpha_{m-k+1} < \alpha_{m-k}$. Let i be a player who contributes to m more than its threshold there, and let j be the other player. Then, by part 3 of Lemma 4, all non-steep projects with coefficients larger than $0.5\alpha_m$ must get a positive contribution from j , for otherwise i would profit by transferring there part of its budget from m . Therefore, the non-steep projects with coefficients larger than $0.5\alpha_m$ receive no contribution from i , according to Lemma 2.

Therefore, at all the steep projects, player j contributes exactly its threshold value, while i contributes above it. Also, i contributes nothing to any non-steep project: we have shown this for the non-steep projects with coefficients larger than $0.5\alpha_m$, now we show it for the rest. If i contributed to a non-steep project with coefficient at most $0.5\alpha_m$, he would benefit from deviating to a steep one, by part 2 of Lemma 4 (when the coefficient is exactly $0.5\alpha_m$, we use the fact that j would be suppressed by such a deviation).

We assume that $B_1 \geq \theta B_2$, and thus, for any $i \neq j$ we have

$$\theta B_j \leq B_i \iff B_j - \theta B_i \leq \frac{B_i}{\theta} - \theta B_i \iff B_j - \theta B_i \leq \frac{B_i - \theta^2 B_i}{\theta}.$$

Thus, a non-steep project with coefficients larger than $0.5\alpha_m$ receives from j at most $\frac{B_i - \theta^2 B_i}{\theta}$, and since i can transfer to that project $B_i - \theta^2 B_i$ without losing a share at the steep projects, i can transfer exactly θ -share of j 's contribution there and profit thereby, by part 3 of Lemma 4 (that lemma assumes j does not contribute to those non-steep projects, but contributing exactly the threshold to such a project is not worse than alone). This profitable deviation contradicts our assumption and we conclude that $\frac{1}{2}\alpha_m \geq \alpha_{m-k}$.

It is left to prove that $B_1 \geq k\theta B_2$. Part 2 of Lemma 4 implies there are no contributions to non-steep projects, since they would render the deviation to the steep projects profitable, unless $\frac{1}{2}\alpha_{m-k+1} = \alpha_{m-k}$, in which case a 2-steep project can get a positive investment from one player. Thus, the players' utility is at most the same as when each steep project obtains contributions from both players, and other projects receive nothing. Thus, each player's utility is at most $k \cdot (\alpha_m/2) \left(\frac{B_1+B_2}{k}\right) = (\alpha_m/2)(B_1 + B_2)$. If 2 could deviate to contribute all B_2 to a steep project while suppressing 1 there, player 2 would obtain $\alpha_m(B_2 + y)$, for some $y > 0$. This is always profitable, since

$$B_2 \geq B_1 \Rightarrow B_2 + 2y > B_1 \iff \alpha_m(B_2 + y) > (\alpha_m/2)(B_1 + B_2).$$

Thus, since we are in an NE, 2 may not be able to suppress i and therefore $B_2 \leq \frac{B_1}{k} \frac{1}{\theta} \Rightarrow B_1 \geq k\theta B_2$. Thus, we have proved that Conditions 1 hold.

Suppose now that $B_1 < \theta B_2$ and we shall derive that Conditions 2 hold.

We exhaust all the possibilities for an NE, as follows: 1 is suppressed, 1 is not suppressed and player 2 does not contribute to non-steep projects, and 1 is not suppressed and player 2 contributes to non-steep projects. We show that each of this options entails at least one of the sub-conditions of 2.

First, consider the case when 1 is suppressed.⁶ Then, 2 invests more than B_1/θ at each project. Therefore, $B_1 < \frac{\theta}{|\Omega|} B_2$. If not all projects were steep, then 2 would profitably transfer some amount to a steep project from the non-steep ones, while still dominating 1 everywhere. This deviation would contradict the profile being an NE. Therefore, all projects are steep and condition 2c holds.

Assume now that no player is suppressed. Therefore, according to Lemma 3, player 1 contributes only to the 2-steep projects, and player 2 contributes to all the steep ones, and perhaps to a 2-steep one as well.

First, we assume that player 2 does not contribute to non-steep projects and show that it entails condition 2a. Next, we assume that 2 does contribute to non-steep projects and show that this entails condition 2b.

First, assume that player 2 does not contribute to non-steep projects. Since player 1 does not prefer to deviate by contributing exactly the threshold at a steep project, $B_1 < \theta \frac{B_2}{k}$ is true. In an NE, player 2 would not profit from suppressing player 1 at a 2-steep project, and therefore

⁶See definition 2.

$$\alpha_m(B_1/\theta) \geq \alpha_{m-k}(B_1(1+1/\theta)) \iff \frac{\alpha_{m-k}}{\alpha_m} \leq \frac{1}{\theta+1}.$$

In addition, in an NE, player 2 would not profit from contributing exactly the threshold at a 2-steep project, and therefore

$$\alpha_m(\theta B_1) \geq \frac{\alpha_{m-k}(1+\theta)B_1}{2} \iff \frac{\alpha_{m-k}}{\alpha_m} \leq \frac{2\theta}{1+\theta}.$$

Thus, we have proved that condition 2a holds.

Next, assume that player 2 contributes to non-steep projects.

Now, according to Lemma 3, $m-k$ is the single 2-steep project, where player 1 contributes all B_1 , while player 2 contributes θB_1 there, and she divides the rest of her budget between all the steep projects, yielding a positive contribution to each such project.

Assume that a steep project receives $y > 0$ (from player 2, of course). If player 1 could achieve the threshold θy , it would deviate, for the following reasons. We have $\alpha_m(B_1+y)/2 > \alpha_{m-k}(B_1+\theta B_1)/2$, unless, perhaps, if $y < \theta B_1$. In such a case, however, 1 can suppress player 2 and obtain $\alpha_m(B_1+y)$, which is larger than $\alpha_{m-k}(B_1(1+\theta))/2$. Consequently, from the profile being an NE, we conclude that 1 is not able to achieve the threshold θy , and therefore

$$B_1 < \theta \frac{B_2 - \theta B_1}{k} \iff B_1 < \frac{\theta B_2}{k + \theta^2}.$$

In addition, since player 1 does not prefer to contribute to $m-k$ only the threshold $\theta^2 B_1$ and move the rest to $m-k-1$, it must hold that

$$\frac{\alpha_{m-k}(B_1(1+\theta))}{2} \geq \frac{\alpha_{m-k}(B_1(\theta^2+\theta))}{2} + \alpha_{m-k-1}(B_1(1-\theta^2)) \iff \alpha_{m-k} \geq 2\alpha_{m-k-1}.$$

Since player 2 does not want to suppress 1 at $m-k$, we conclude that

$$\alpha_m((B_1)/\theta - \theta B_1) + \frac{\alpha_{m-k}(B_1(1+\theta))}{2} \geq \alpha_{m-k}(B_1(1+1/\theta)) \iff \frac{\alpha_{m-k}}{\alpha_m} \leq \frac{2(1-\theta)}{2-\theta}.$$

Finally, since player 2 does not prefer moving θB_1 to a steep project over leaving it at $m-k$, it holds that

$$\frac{\alpha_{m-k}(B_1(1+\theta))}{2} \geq \alpha_m(\theta B_1) \iff \frac{\alpha_{m-k}}{\alpha_m} \geq \frac{2\theta}{1+\theta}.$$

Therefore, condition 2b holds. To conclude, at least one of the sub-conditions of 2 holds, thus finalizing the proof of the other direction of the theorem. \square

We conclude that besides the equilibria with $\alpha_m = \alpha_{m-k}$, there exists an NE if and only if α_{m-k} is at most a constant fraction of α_m .

Corollary 1. *Assume the conditions of Theorem 3 and that there exist two projects. Then, once all the parameters besides α_{m-k} and α_m are set, there exists a $C > 0$, such that an NE exists if and only if $\frac{\alpha_{m-k}}{\alpha_m} \leq C$, and, perhaps, if and only if $\frac{\alpha_{m-k}}{\alpha_m} = 1$.*

Proof. Consider the bounds on the possible values of $\frac{\alpha_{m-k}}{\alpha_m}$ such that at least one NE exists, besides case 2c. From Theorem 3, the only way that this corollary could be wrong would require the upper bound on $\frac{\alpha_{m-k}}{\alpha_m}$ from 2a to be strictly smaller than the lower bound from 2b, while the two bounds on $\frac{\alpha_{m-k}}{\alpha_m}$ from 2b gave a non-empty segment. These conditions mean that both $\frac{1}{1+\theta} < \frac{2\theta}{1+\theta}$ and $\frac{2\theta}{1+\theta} \leq \frac{2(1-\theta)}{2-\theta}$ should hold. The first inequality means $\theta > 0.5$, while the second one means $\theta \leq 0.5$. Since these conditions cannot hold simultaneously, the corollary is never wrong. \square

The proof of the necessity of the conditions of Theorem 3 relies on the lemmas that describe the structure of an NE, which are not easily generalized for $n > 2$. However, some of the sufficiency conditions, namely 1 and 2c, extend for a general n as follows.

Theorem 4. Consider an equal θ -sharing game with $n \geq 2$ players with budgets $B_n \geq \dots \geq B_2 \geq B_1$ (the order is w.l.o.g.), $0 < \theta < 1$, and linear project functions with coefficients $\alpha_m = \alpha_{m-1} = \dots = \alpha_{m-k+1} > \alpha_{m-k} \geq \alpha_{m-k-1} \geq \dots \geq \alpha_1$ (the order is w.l.o.g.).

This game has a pure strategy NE, if one of the following holds.⁷

1. $B_{n-1} \geq \theta B_n$ and the following both hold.

- (a) $\frac{1}{n} \alpha_m \geq \alpha_{m-k}$,
- (b) $B_1 \geq k\theta B_n$;

2. $B_{n-1} < \theta B_n$ and also the following holds.

- (a) $B_{n-1} < \frac{\theta}{|\Omega|} B_n$ and all the project functions are equal, i.e. $\alpha_m = \alpha_1$.

Proof. It is analogous to the proof for $n = 2$, with the following remarks. In case 1, all the players equally divide their budgets among all the steep projects. In case 2, player n dominates everyone else. \square

In order to facilitate decisions, it is important to analyze the efficiency of the various Nash Equilibria. We aim to find the famous price of anarchy (PoA), which is the ratio of a worst NE's efficiency to the optimum possible one, and the price of stability (PoS), which is the ratio of a best NE's efficiency to the optimum possible one.

We first completely handle two players.

Theorem 5. Consider an equal θ -sharing game with two players with budgets B_1, B_2 . W.l.o.g., $B_2 \geq B_1$. Assume $0 < \theta < 1$, and linear project functions with coefficients $\alpha_m = \alpha_{m-1} = \dots = \alpha_{m-k+1} > \alpha_{m-k} \geq \alpha_{m-k-1} \geq \dots \geq \alpha_1$ (the order is w.l.o.g.).⁸

1. Assume that $B_1 \geq \theta B_2$ and the following both hold.

- (a) $\frac{1}{2} \alpha_m \geq \alpha_{m-k}$,
- (b) $B_1 \geq k\theta B_2$;

⁷If α_{m-k} does not exist, consider the containing condition to be vacuously true.

⁸If α_{m-k} does not exist, consider the containing condition to be vacuously true.

Then, there exists a pure strategy NE and there holds: $\text{PoS} = \text{PoA} = 1$.

2. Assume that $B_1 < \theta B_2$ and also at least one of the following holds.

(a) $B_1 < \frac{\theta B_2}{k}$ and $\frac{\alpha_{m-k}}{\alpha_m} \leq \min \left\{ \frac{1}{1+\theta}, \frac{2\theta}{1+\theta} \right\}$.

Then, there exists a pure strategy NE and the following holds. $\text{PoS} = \frac{\alpha_m B_2 + \alpha_{m-k} B_1}{\alpha_m (B_1 + B_2)}$. If the case 2b holds as well, then $\text{PoA} = \frac{\alpha_m (B_2 - \theta B_1) + \alpha_{m-k} (B_1 (1 + \theta))}{\alpha_m (B_1 + B_2)}$; otherwise, $\text{PoA} = \text{PoS}$.

(b) $B_1 < \frac{\theta B_2}{k + \theta^2}$ and $\alpha_{m-k} \geq 2\alpha_{m-k-1}$ and $\frac{2\theta}{1+\theta} \leq \frac{\alpha_{m-k}}{\alpha_m} \leq \frac{2(1-\theta)}{2-\theta}$ and $m-k$ is the only 2-steep project.

Then, there exists a pure strategy NE and the following holds. If the case 2a holds as well, then $\text{PoS} = \frac{\alpha_m B_2 + \alpha_{m-k} B_1}{\alpha_m (B_1 + B_2)}$; otherwise, $\text{PoS} = \frac{\alpha_m (B_2 - \theta B_1) + \alpha_{m-k} (B_1 (1 + \theta))}{\alpha_m (B_1 + B_2)}$. In any case, $\text{PoA} = \frac{\alpha_m (B_2 - \theta B_1) + \alpha_{m-k} (B_1 (1 + \theta))}{\alpha_m (B_1 + B_2)}$.

(c) $B_1 < \frac{\theta}{|\Omega|} B_2$ and all the project functions are equal, i.e. $\alpha_m = \alpha_1$.

Then, there exist pure NE and $\text{PoS} = 1, \text{PoA} = \frac{B_2}{B_1 + B_2}$.

The proof is in Appendix A. Its idea is to show that some equilibria are the most efficient and some are the least efficient, using Lemmas 2, 3, and 4 to limit the possibilities for equilibria.

We derive the exact lower bound (infimum) and the maximum of the price of anarchy and stability.

Corollary 2. Consider an equal θ -sharing game with two players with budgets B_1, B_2 . W.l.o.g., $B_2 \geq B_1$. Assume $0 < \theta < 1$, and linear project functions with coefficients $\alpha_m = \alpha_{m-1} = \dots = \alpha_{m-k+1} > \alpha_{m-k} \geq \alpha_{m-k-1} \geq \dots \geq \alpha_1$ (the order is w.l.o.g.).⁹ Then, the infimum of PoS over all the cases is $\frac{k}{k+\theta}$ (> 0.5), and the maximum is 1. The same holds for PoA.

We now generalize the efficiency results 1 and 2c of Theorem 5 for a general $n \geq 2$.

Theorem 6. Consider an equal θ -sharing game with $n \geq 2$ players with budgets $B_n \geq \dots \geq B_2 \geq B_1$, $0 < \theta < 1$ (the order is w.l.o.g.), and linear project functions with coefficients $\alpha_m = \alpha_{m-1} = \dots = \alpha_{m-k+1} > \alpha_{m-k} \geq \alpha_{m-k-1} \geq \dots \geq \alpha_1$ (the order is w.l.o.g.).¹⁰

1. Assume that $B_{n-1} \geq \theta B_n$ and the following both hold:

(a) $\frac{1}{n} \alpha_m \geq \alpha_{m-k}$,

(b) $B_1 \geq k\theta B_n$;

Then, there exists a pure strategy NE and there holds: $\text{PoS} = 1$.

2. Assume that $B_{n-1} < \theta B_n$, $B_{n-1} < \frac{\theta}{|\Omega|} B_n$ and all the project functions are equal, i.e. $\alpha_m = \alpha_1$.

Then, there exist pure NE and the following holds: $\text{PoS} = 1, \text{PoA} = \frac{B_n}{\sum_{i \in \{1, 2, \dots, n\}} B_i}$.

⁹If α_{m-k} does not exist, consider the containing condition to be vacuously true.

¹⁰If α_{m-k} does not exist, consider the containing condition to be vacuously true.

To summarize the section, for linear project functions and the equal θ -sharing model, we completely analyze the existence and efficiency of equilibria for 2 agents and provide some results for $n \geq 2$ agents as well. Some other cases are simulated in the next section.

4. MIXED NASH EQUILIBRIUM

As we show in the previous section, a game may not possess a pure NE. Therefore, we naturally turn to mixed extensions¹¹ and ask whether a mixed extension always has a NE. At first, this is unclear. As the game is infinite, the theorem by Nash [18] about the existence of a mixed NE in finite games is irrelevant. Since the game is not continuous, even the theorem by Glicksberg [19] about the existence of a mixed NE in continuous games is not applicable. Fortunately, we answer affirmatively employing a more general existence theorem by Maskin and Dasgupta [20].

Theorem 7. *Any shared effort game with linear project functions and θ -equal sharing has a mixed Nash equilibrium.*

The existence result automatically extends to the solution concepts that include mixed Nash equilibria, such as correlated [21] and coarse correlated [22] equilibria. Luckily, not only existence results but also some bounds on the social welfare of solution concepts extend to the other equilibria as well. An important preliminary observation is that the maximum social welfare stays the same even when (correlated) randomization is allowed; it is always $\alpha_m \sum_{i=1}^n B_i$.

Consider the results of Theorem 6. Its lower bounds on the price of anarchy stem from the utility that certain players can always achieve, and the bounds, therefore, hold for mixed, correlated and coarse correlated equilibria as well. Since any pure NE is also a mixed/correlated/coarse-correlated NE, the rest of the efficiency results, based on presenting an NE, also extend to the other solution concepts.

To conclude, we prove the existence of NE in the mixed case. Then, we show that the efficiency bounds from Theorem 6 apply to the mixed case as well. This is the only section dealing with not only pure NE; so the next section already considers pure NE.

5. SIMULATIONS

Theory covers shared effort games with more than two players only partially. To explore the existence and efficiency of NE in these games, we simulate a variation of fictitious play [23], which comes to mind for its well-known convergence properties (see Section 6.1). At each time step of a fictitious play, each player best-responds to the cumulative strategy of the others.

Since the classical fictitious play is defined for finite games, and we are dealing with an infinite game, we may not apply fictitious play as is. Therefore, We first adapt the fictitious play [14] to our infinite game. Danskin [24] defines the best response to maximize the average utility against all the previous strategies of the other players, while we, as well as the original fictitious play, best respond to the cumulative strategy of the others.

¹¹A mixed extension has strategies that are distributions on the pure strategies and the respective utilities are the expected utilities under these distributions.

After generalizing the fictitious play, we want to run it. In fictitious play [14], every player finds a best response to the current strategies of the others at each step, if it exists; if not, the play is undefined.

We suggest an algorithm for finding a best response, if it exists, for the case of two projects. This allows implementing the adapted fictitious play.

5.1. INFINITE-STRATEGY FICTITIOUS PLAY FOR SHARED EFFORT GAMES

Since the game is infinite, we do not need mixing to average the strategies. Denote the set of all the best responses of player i to profile x^{-i} of the others by $\text{BR}(x^{-i})$.

Definition 4. Given a shared effort game with players N , budget-defined strategies $S^i = \{x^i = (x_\omega^i)_{\omega \in \Omega} \in \mathbb{R}_+^{|\Omega|} \mid \sum_{\omega \in \Omega_i} x_\omega^i \leq B_i, \omega \notin \Omega_i \Rightarrow x_\omega^i = 0\}$ and utilities $u^i(x) \triangleq \sum_{\omega \in \Omega_i} \phi_\omega^i(x_\omega)$, define an Infinite-Strategy Fictitious Play (ISFP) as the following set of sequences. Consider a (pure) strategy in this game at time 1, i.e. $(x^i(1))^{i \in N} = ((x^i(1)_\omega)_{\omega \in \Omega})^{i \in N}$, define $X^i(1) \triangleq \{x^i(1)\}$, and define recursively, for each $i \in N$ and $t \geq 0$, the set of the possible strategies at time $t+1$:

$$X^i(t+1) \triangleq \left\{ \frac{tx^i(t) + \text{br}(x^{-i}(t))}{t+1} \mid x^i(t) \in X^i(t), \text{br}(x^{-i}(t)) \in \text{BR}(x^{-i}(t)) \right\}. \quad (2.2)$$

Thus, a fictitious play is a best response to the arithmetic average of the others' actions till now.

We say that an ISFP converges to $x^* \in \mathbb{R}_+^n$ if at least one of its sequences converges to x^* in every coordinate.

Since $\text{BR}(x^{-i}(t))$ is a set, there may be multiple ISFP sequences. For an ISFP to be defined, we need that $\text{BR}(x^{-i}(t)) \neq \emptyset$, that is the utility functions attain a maximum. Since the utility functions are, generally speaking, not upper semi-continuous, they may sometimes not attain a maximum, rendering the ISFP undefined.

In ISFP, all the plays have equal weights in the averaging. In the other extreme, a player just best-responds to the previous strategy profile of other players, thereby attributing the last play with the weight of 1 and all the other plays with 0. In general, we define, for an $\alpha \in [0, \infty]$, an α -ISFP play as in Definition 4, but with the following formula instead¹²

$$X^i(t+1) \triangleq \left\{ \frac{\alpha tx^i(t) + \text{br}(x^{-i}(t))}{\alpha t + 1} \mid x^i(t) \in X^i(t), \text{br}(x^{-i}(t)) \in \text{BR}(x^{-i}(t)) \right\}. \quad (2.3)$$

Here, the last play's weight is $\frac{1}{\alpha t + 1}$.

We do not know whether and when any convergence property can be proven for the generalized fictitious play in shared effort games. In simulations, our generalized fictitious play often converges to a NE. Roughgarden and Tardos [25, Chapter 17] say that a highly probable convergence to an equilibrium in real life dynamics bolsters the importance of its efficiency.

Next, we solve the algorithmic problem of finding whether a best response exists, and if it does, what it is.

¹²For $\alpha = \infty$, we just obtain a constant sequence.

5.2. BEST RESPONSE IN A 2-PROJECT GAME

Let the projects be $\Omega = \{\psi, \omega\}$. We would like to find a best response for a player $i \in N$, all the other players' strategies $x^{-i} \in S^{-i}$ being fixed. From the weak monotonicity of the share functions, we may assume w.l.o.g. that a best responding player contributes all her budget. Then, a strategy is uniquely determined by the contribution to project ψ and we shall write x^i for x_{ψ}^i , meaning that $x_{\omega}^i = B_i - x^i$.

We now state the conditions for the following theorem. Consider M_{eq}^{θ} sharing and convex project value functions. Let $D_0^i < D_1^i < \dots < D_m^i$ and $W_0^i < W_1^i < \dots < W_l^i$ be the jumps of ϕ_{ψ}^i and ϕ_{ω}^i respectively. (The first points in each list are the minimum contributions to projects ψ, ω , respectively, required for i to obtain a share. The other points are the points at which another player becomes suppressed at the respective project.) The possible discontinuity points of the total utility of i are thus $D_0^i < D_1^i < \dots < D_m^i$ and $B_i - W_l^i < \dots < B_i - W_1^i < B_i - W_0^i$. Denote the distinct points of these lists merged in the increasing order by L . Let all the contributions of the players in $N \setminus \{i\}$ be fixed as they are given, and consider x^i as the only variable. Let L_{B_i} denote the points of the list L that are on $[0, B_i]$, together with 0 and B_i , and let M_{B_i} be L_{B_i} with an arbitrary point added between each two consecutive points.

Theorem 8. *The maximum of the one-sided limits at the points of L_{B_i} and of the values at the points of M_{B_i} yields the utility supremum¹³ of the responses of player i . This supremum is a maximum (and in particular, a best response exists) if and only if it is achieved at a point of M_{B_i} .*

When finding the one-sided limit at a point of L_{B_i} takes constant time, the resultant algorithm runs in $O(n \log n)$ time and in linear space.

5.3. THE SIMULATION METHOD

We consider the θ -equal 2-project case with linear project functions, where Theorem 8 supports best responding. For each of the considered shared effort games, we run several α -ISFPs, for several α s. If at least once in the simulation process no best response exists, we drop this attempt. While running an α -ISFP, we stop after a predefined number of iterations, or if an NE has been found.

We choose an initial belief state about all the players and run the ISFP from this state on, updating this common belief state at each step by finding a best response of each player to the current belief state. To increase the chances of finding an existing NE, for each game, we generate 45 fictitious plays by randomly and independently generating the initial belief state on each player's actions, uniformly over the possible histories. We could have tried more than 45, but it takes too much time, while not helping much; for example, if we ran 60 fictitious plays, which is 33% more than 45, the leftmost figure on the first row of Figure 2.3 would change as shown in Figure 2.2, which is not much different for such a large increase in the run time. While simulating, when a player has multiple best responses, we choose a closest one to the current belief state of the fictitious play, in the sense of minimizing the maximum distance from the last action's components. If we find an NE in at most 50 iterations, we give a positive answer, and otherwise,

¹³The supremum is the exact upper bound; it always exists.

this attempt does not solve this game. For each found NE, we calculate its efficiency by dividing its total profit by the optimum possible total profit and plot it using shades of gray. When no NE is found, we plot it black.

2

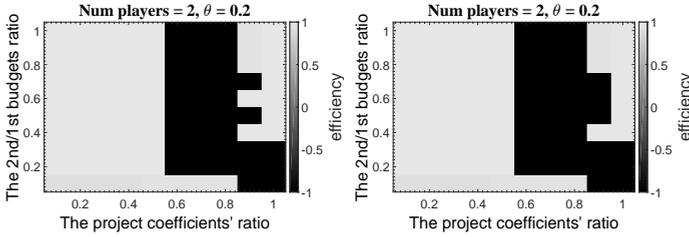


Figure 2.2: The existence and efficiency of NE for 2 players as a function of the ratio of the project functions coefficients and the ratio of the two largest budgets. The right figure is obtained from increasing the number of plays by 33%, relatively to the left one. Only 1 data point has changed.

5.4. RESULTS AND CONCLUSIONS

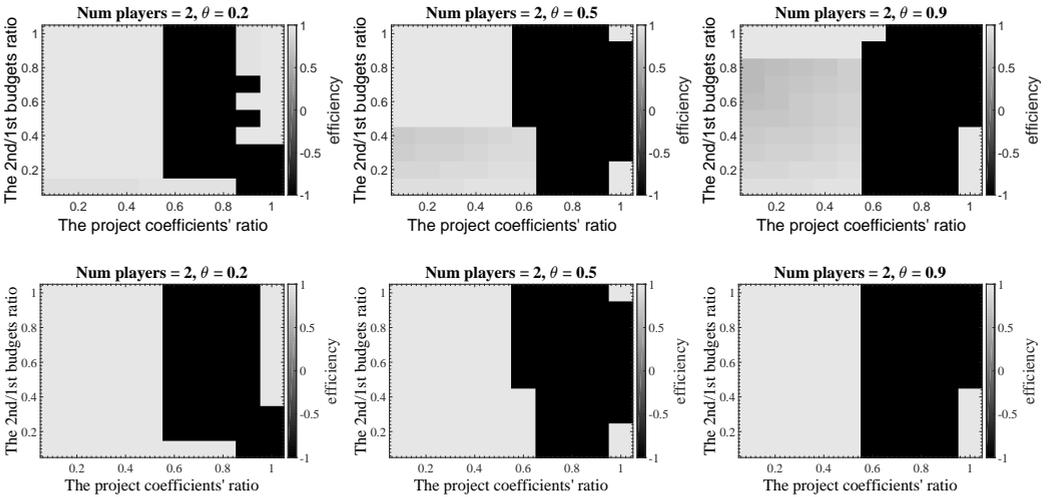


Figure 2.3: The existence and efficiency of NE for 2 players as a function of the ratio of the project functions coefficients and the ratio of the two largest budgets. The first row plots the results of the simulations, and the second row shows the theoretical predictions. In the results of the simulations, black color means that Nash Equilibrium has not been found. For all the other cases, the efficiency, a value in (0, 1], is shown by the shade (from dark gray = 0.0 to white = 1.0). In the plot of the theoretical predictions, black color means that no NE exists, while white means that a Nash Equilibrium exists.

We present the results representing the trends in the body of the chapter; the reader can see also other, but similar, results in Section B.

First, to validate our simulations, we compare the results of the simulations to the theoretically known case of two players. This can be seen in Figure 2.3. We see that the

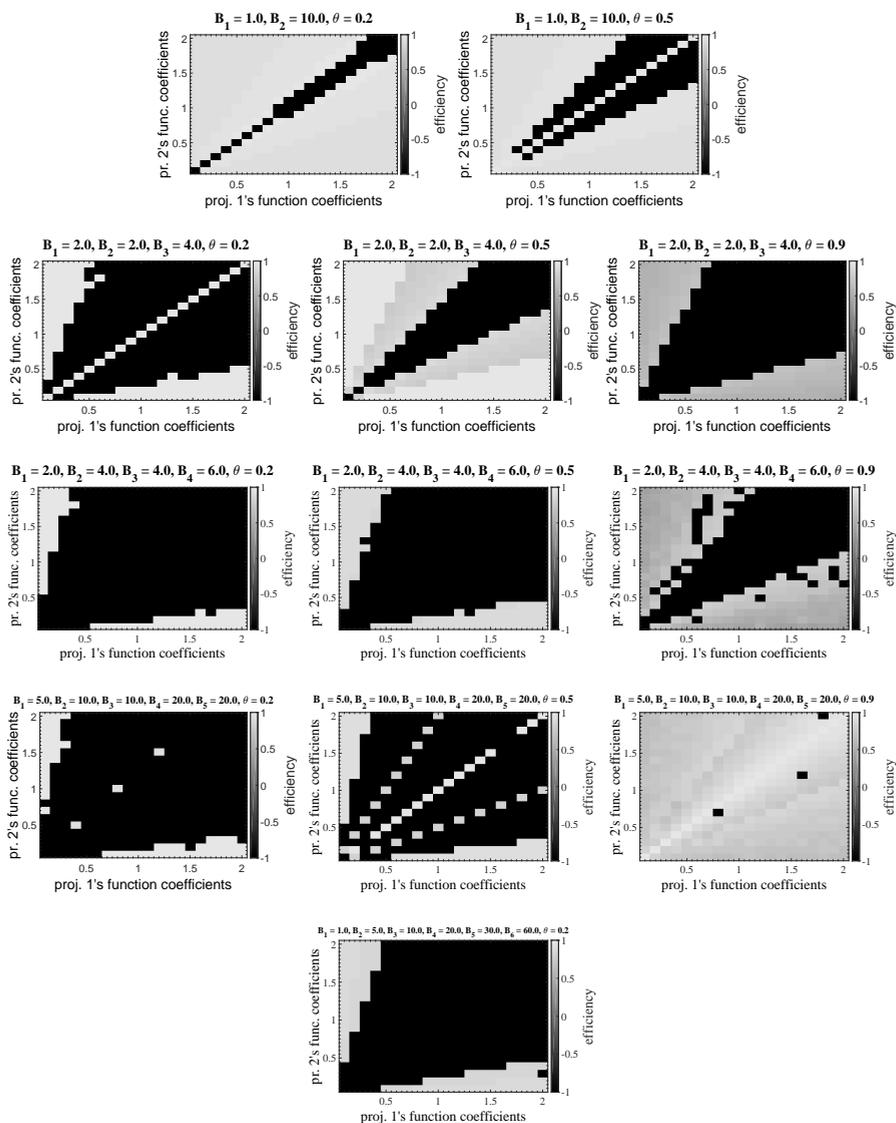


Figure 2.4: The existence and efficiency of NE as function of project functions for 2, 3, 4, 5 and 6 players. Black color means that Nash Equilibrium has not been found. For all the other cases, the efficiency, a value in $(0, 1)$, is shown by the shade (from dark gray = 0.0 to white = 1.0).

2

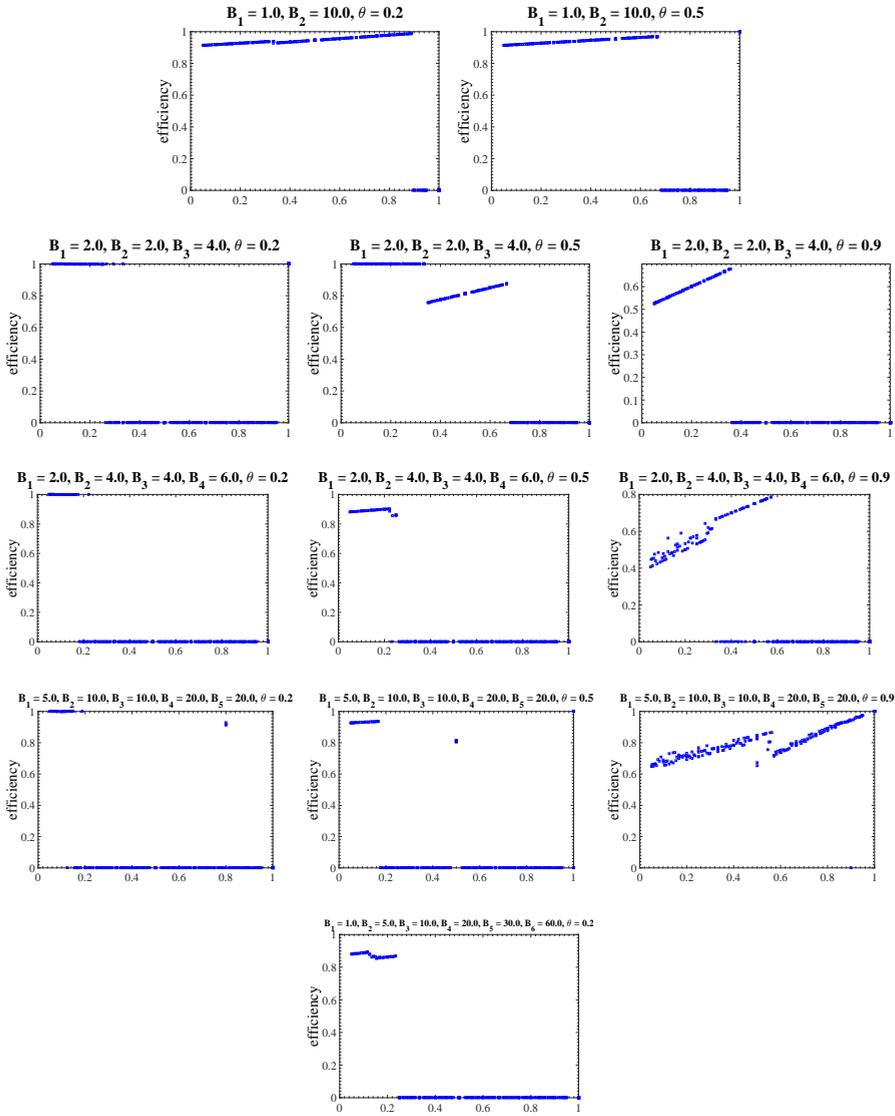


Figure 2.5: The efficiency as a function of project value functions' ratio for 2,3,4,5 and 6 players. Efficiency of 0 means that Nash Equilibrium has not been found. For all the other cases, the efficiency, a value in (0, 1), is shown.

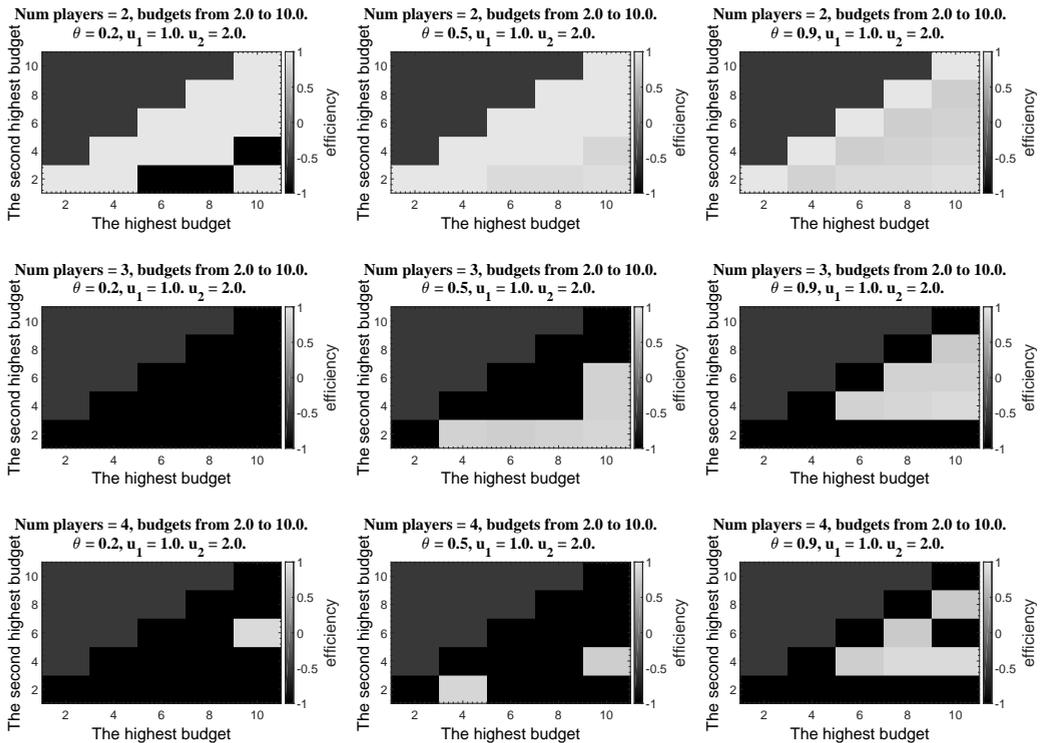


Figure 2.6: The existence and efficiency of NE as function of the largest and the second largest budgets. Black color means that Nash Equilibrium has not been found, and very dark gray color means the area is not defined, since the second highest budget may not be larger than the highest one. For all the other cases, the efficiency, a value in $(0, 1]$, is shown by the shade (from dark gray = 0.0 to white = 1.0).

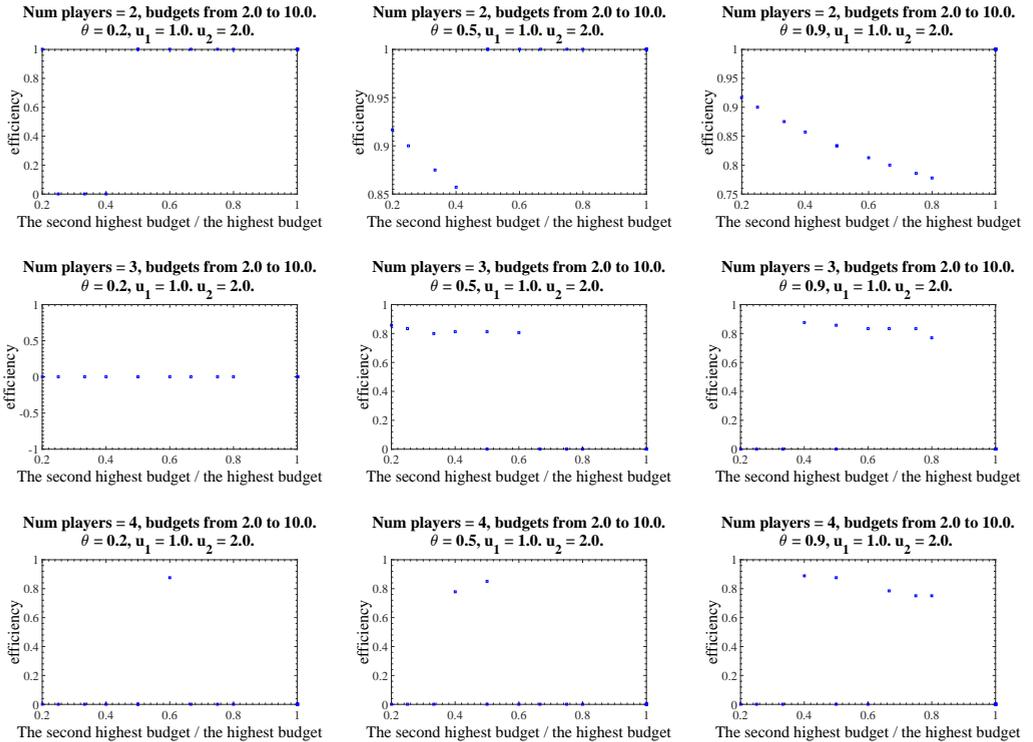


Figure 2.7: The efficiency as a function of the ratio of the two largest budgets. Efficiency of 0 means that Nash Equilibrium has not been found, and -0.5 means the the domain is not defined, since the second highest budget may not be larger than the highest one. For all the other cases, the efficiency, a value in $(0, 1]$, is shown.

simulations are imperfect, since they do find an equilibrium in several cases when no equilibrium exists. In Figure 2.8 in Section B, we also observe that simulations can fail to find an existing NE. Nonetheless, they are quite accurate, and therefore, useful for obtaining an impression on the existence and the efficiency of NE. In particular, for 2 players, where we have a perfect theoretical prediction, we see that the simulations err for $\theta = 0.2$ in 5 from $10 * 10 = 100$ points, which constitute 5%. For $\theta = 0.5$, no errors exist, and for $\theta = 0.9$, the error percentage is 1%.

Besides Figure 2.3, Figure 2.8, and Figure 2.9 in Section B describe the existence and efficiency of the NE as a function of the ratio of the coefficients of the project functions and of the ratio of the second highest budget to the highest one, and compare this to the theoretical predictions. Notice the difference between the areas in Figure 2.3 that correspond to $B_1 \geq \theta B_2$ and to $B_1 < \theta B_2$, fitting Theorem 3. For more than two players, a similar but less sharp difference appears. For more than two players, our theory is not a complete characterization, and this reason for lack of fit can be exacerbated by the possible simulation errors, so we do not expect a complete correspondence. We notice that the area of equilibrium existence is more complicated for multiple players than the sufficient conditions from Theorem 4 predict, meaning that the simulations find that the sufficient conditions are not necessary.

For two agents, except for the NE when the two projects are equal, an NE exists for a budget ratio if and only if the value functions ratio is below a certain value. This is proven in Corollary 1.

When the project function coefficients are the independent variables, Figure 2.4 presents the NE. Each line of the simulation plots corresponds to a setting for a given number of players, and within a line, the plots are generated for an increasing sequence of θ . Mostly, an NE exists except a cone where the project functions are quite close to each other. Interestingly, sometimes an NE exists also when the project functions are nearly the same (at ratio 1), or at another constant ratio with each other. In all cases, the ratio of project functions determines existence of an NE, which, for two players, is predicted by Theorem 3. Usually, the more there are players, the less settings with an NE we find.

More results with similar trends appear in Figure 2.10, Figure 2.11 and Figure 2.12 in Section B.

We plot efficiency as a function of the ratio of project functions' coefficients in Figure 2.5. Mostly, the efficiency is uniquely determined by the ratio of the project functions, and the function is piecewise linear. Each linear piece is non-decreasing. For two players, it is linear, in the spirit of Theorem 5. (Though not directly predicted by it, since the theorem considers extremely efficient or inefficient equilibria, and has several cases, which imply piecewise linearity.) The larger is the θ , the steeper becomes the piecewise linear dependency. More results with similar trends appear in Figure 2.13, Figure 2.14 and Figure 2.15 in Section B.

When the largest and the second largest budgets are the independent variables, Figure 2.6 presents the equilibria. (The budgets of the other players are spread on equal intervals). Each line of the plot corresponds to a setting of project function coefficients for a given number of players, and within a line, the plots are generated for an increasing sequence of θ . Whether we find an NE almost always depends on the ratio of the second

highest budget to the highest one. For two players, this fits Theorem 3. More results with similar trends appear in Figure 2.16 in Section B.

Figure 2.7 plots efficiency as a function of the ratio of the two largest budgets. Efficiency is determined by the ratio of the second highest to the highest budgets. For two agents, this is in the spirit of Theorem 5. More results with similar trends appear in Figure 2.17 in Section B.

To summarize, the existence of NE is related to the ratio of project function coefficients and the budget ratio being in some limits, limits that in particular depend on the threshold. This is in the spirit of Theorems 3 and 4. Based on Theorem 3, Corollary 1, and the simulation results, we hypothesize that an NE exists if and only if at least one of several sets of conditions on the ratio of the budgets holds, and for every such condition, several conditions of being smaller or equal to a function of the threshold (and not the budgets) on the ratio of project function coefficients hold together. Therefore, the more players exist, the more conditions on the budgets have to hold simultaneously, and therefore, the less NE exist. We also observe that when only the project function coefficients change, there can sometimes be an NE in the case of project coefficients being at a single given ratio, and besides this case, an NE exists if and only if the ratio of the smaller coefficient to the larger one is at most some value.

The efficiency of the NE that we find depends on the ratio of the project function coefficients and on the ratio of the budgets, in the spirit of Theorems 5 and 6. Based on Theorem 5 and the simulation results, we hypothesize that the price of anarchy and stability of a shared effort game depends piecewise linearly on the project function coefficients ratio and (in some other manner) on the budget ratio. In each linear piece, the dependency on the project function coefficients is non-decreasing with θ .

6. RELATED WORK

Most of the relevant related work is presented in Section 1.1.1. We do repeat some parts here, to keep this chapter self-contained, and we also elaborate on some points, relevant to this chapter. We first present a model, resembling ours, but quite different. Then, we present the area of efficiency of Nash equilibria, and proceed to what has been said about the existence and efficiency of NE for shared effort, concluding that no analysis of the general setting has taken place, a gap which we partially fill. In Section 6.1, we motivate our usage of the fictitious play.

Zick, Elkind and Chalkiadakis [26, 27] suggest a model, which is reminiscent of ours, but expressed in the terms of cooperative games, where every coalition of players has a value. Contributing to a coalition can be considered as contributing to a project in a shared effort game, and in both cases players have budgets and obtain revenues. Despite these similarities, the model of Zick et al. is not a game in our sense, since it does not define profits. Moreover, even if we consider a shared effort game as a particular case of their model, a positive threshold can vitiate individual rationality, since a player who obtains a positive profit when she is the only contributor to a project may obtain nothing when others contribute to the same project.

An NE can be inefficient, such as in the famous example of the prisoner's dilemma [15, Example 16.2]. Since the introduction of price of anarchy [28] and price of stability [12], there have been many studies on the matter.

Roughgarden and Tardos [25, Chapter 17] discuss inefficiency of equilibria in non-cooperative games and consider the examples of network, load balancing and resource allocation games. They argue that understanding exactly when selfish behavior is socially profitable is important, since in many applications, implementing control is extremely difficult. Roughgarden and Tardos mention that the use of ratio of the objectives in an equilibrium and in the optimum to measure efficiency (PoA, PoS) constitute the two most popular approaches to choosing which equilibrium to use. The price of anarchy measures the best guarantee on an NE, while the price of stability measures the cost of leading the game to a specific equilibrium. Another possible approach is average-case analysis, being much more difficult to define and analyze. We now present what has been said about the efficiency of games, similar to ours.

Let us present the related models. Shared effort games where only the contributor with the highest contribution obtains the project's value, while everyone pays, are called all-pay auctions, and their equilibria have been studied, for instance, by Baye, Kovenock and de Vries [29]. All-pay auctions model lobbying, single-winner contests, political campaigns, striving for a job promotion (see e.g. [4]) and Colonel Blotto games with two players [8]. In a Colonel Blotto game, two generals divide their armies between battlefields, and at every battlefield, the larger force wins. The overall number of the won battlefields defines the utility of a general. Roberson [8] analyzes the equilibria of this game and their expected payoffs. Any outcome is socially optimum, since this is a constant-sum game.

For shared effort games, under very specific conditions (obtaining at least a constant share of one's marginal contribution to the project's value and $\theta = 0$), Bachrach et al. [13] have shown that the price of anarchy (PoA) is at most the number of players. This work also shows upper bounds on the PoA for the case of convex project functions, where each player receives at least a constant share of its marginal contribution to the project's value. In this paper, we study more general $\theta \in [0, 1]$ sharing mechanisms without these conditions, provide precise conditions for existence of NE, and find their efficiency. Anshelevich and Hoefer [30] considered an undirected graph model, where the nodes are the players and each player divides its budget between its adjacent edges in minimum effort games (where the edges are the 2-player projects), each of which equally rewards both sides by measure of the project's success (i.e., duplication instead of division). Anshelevich and Hoefer prove the existence of equilibria, find the complexity of finding an NE, and find that the PoA is at most 2. A related setting of *multi-party computation games* appeared in [31]. There, the players are computing a common function that requires them to compute a costly private value, motivating free-riding. The work suggests a mechanism, where honest computation is an NE. This differs from our work, since Smorodinsky and Tennenholtz consider cost minimization, and the choice of the players is either honestly computing or free riding, no choice of projects.

To conclude, there has been no research of the NE of our problem in the general case.

6.1. FICTITIOUS PLAY

We generalize and employ fictitious play, introduced by Brown [14], to find equilibria in simulations. It has been widely researched. In this play, each player best-responds to the product of cumulative marginal histories of the others' actions at every time step. It

is a myopic learning process, where each player always best responds to the other players' strategies at every time step. If the game is finite, then if a fictitious play converges, then the distribution in its limit is an NE [32]. Conversely, a game is said to possess the fictitious play property if every fictitious play approaches equilibrium in this game [33]. Many researchers show games that possess this property, for example, two-person zero-sum games [23] and finite weighted potential games [33]. A famous example for a game without such property is a 3×3 game by Shapley [34]. In this game, there is a cyclic fictitious pay that plays each strategy profile for at least an exponentially growing number of times, and therefore, does not converge at all. Moreover, even its subsequences do not converge to an NE.

7. CONCLUSIONS AND FUTURE WORK

This chapter considers shared effort games where the players contribute to the given projects, and subsequently share the profits of these projects, conditionally on the allocated effort. We study existence and efficiency of the NE, arriving at the following.

A pure NE exists if the utility functions are continuous and concave, and the strategy sets are non-empty, compact and convex. We first characterize the existence and efficiency of pure NE for shared effort games with two players and linear project functions. When a NE exists and the budgets are close to each other, all the NE are socially optimal. When the budgets are further apart, in the sense that smallest budget is less than threshold times the largest one, the efficiency depends on the ratio of the budgets and of the two or three largest projects' coefficients and is always greater than half of the optimum.

When the budgets are close, we demonstrate an optimal NE where everyone equally spreads her budgets between the most profitable projects. This motivates the organizers of a project to make it most profitable possible. Even second best can be no good.

For arbitrarily many players, we find socially optimal pure equilibria in some cases. When all the projects are equivalent and the largest budget is much larger than the rest together, then every NE is nearly optimal. This bound also holds for the mixed extension, where we show that an NE always exists.

For more than two players, we simulate fictitious play, to study the existence of NE. To this end, we generalize fictitious play to infinite strategy spaces and describe some of the best responses of a player to the other players' strategies. In the cases where we find a NE, we also estimate its efficiency. All the theoretical predictions about the simulated cases have been corroborated. For more than three players, the efficiency of a NE can be suboptimal, as we see in Section 5.4. The most important factor for existence and efficiency of an equilibrium is the ratio of the largest to the second largest project function coefficients and of the largest to the second largest budgets. Therefore, to influence the projects and the agents to some extent, one should better influence the projects with the highest revenues and the agents with the largest abilities (budgets).

Consider some directions for future work. Mixed equilibria model randomization over pure strategies. Since randomization can be undertaken in practice, we would like to find concrete mixed equilibria to be able to advise on playing them, like we do here for the pure equilibria. Second, we would like to extend our complete theoretical characterization of the existence of (pure) NE to more than two players and to non-linear project

functions. Next, extending simulations to more than two projects would improve our understanding of the various NE that are not yet analyzed analytically.

The theoretical analysis of efficiency implies that for two players with close budgets, no coordination is needed, since the price of anarchy is 1. The price of anarchy is close to 1 also for two players with budgets that are far from each other, and it is more than half. For three or more players, some coordination may improve the total utility, though we have seen many cases with efficiency above 0.75. We have provided conditions for a general number of players where every equilibrium is almost optimal, so no coordination is required. To conclude, we have analyzed when contributions to public projects are in equilibrium and what is lost in these equilibria relatively to the best possible contribution profiles. In the scenarios where much is lost, coordination may improve efficiency. We have assumed every project yields some value. In reality, however, a project often needs to meet certain requirements to obtain any utility altogether. These additional requirements are modeled in the following chapter.

REFERENCES

- [1] G. Polevoy, S. Trajanovski, and M. M. de Weerd, *Nash Equilibria in Shared Effort Games*, in *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems, AAMAS '14* (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2014) pp. 861–868.
- [2] F. M. Harper, S. X. Li, Y. Chen, and J. A. Konstan, *Social comparisons to motivate contributions to an online community*, in *Proceedings of the 2nd international conference on Persuasive technology* (Springer-Verlag, Berlin, Heidelberg, 2007) pp. 148–159.
- [3] A. Forte and A. Bruckman, *Why do People Write for Wikipedia? Incentives to Contribute to Open-Content Publishing*, in *GROUP Workshop: Sustaining Community: The Role and Design of Incentive Mechanisms in Online Systems. Sanibel Island, FL* (2005) pp. 6–9.
- [4] R. Siegel, *All-pay contests*, *Econometrica* **77**, 71 (2009).
- [5] J. Kleinberg and S. Oren, *Mechanisms for (mis)allocating scientific credit*, in *Proceedings of the 43rd annual ACM symposium on Theory of computing* (ACM, New York, NY, USA, 2011) pp. 529–538.
- [6] J. Kleinberg and E. Tardos, *Balanced outcomes in social exchange networks*, in *Proceedings of the 40th annual ACM symposium on Theory of computing* (ACM, New York, NY, USA, 2008) pp. 295–304.
- [7] *Course grading policy*, <http://www.physics.umd.edu/courses/Phys121/Griffin/grading.html> (2016).
- [8] B. Roberson, *The Colonel Blotto game*, *Economic Theory* **29**, 1 (2006).
- [9] B. D. Nguyen and K. M. Nielsen, *What death can tell: Are executives paid for their contributions to firm value?* AFA Denver Meetings Paper (2011).

- [10] E. Koutsoupias and C. H. Papadimitriou, *Worst-case equilibria*, *Computer Science Review* **3**, 65 (2009).
- [11] A. S. Schulz and N. S. Moses, *On the performance of user equilibria in traffic networks*, in *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '03 (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2003) pp. 86–87.
- [12] E. Anshelevich, A. DasGupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden, *The price of stability for network design with fair cost allocation*, in *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on* (2004) pp. 295–304.
- [13] Y. Bachrach, V. Syrgkanis, and M. Vojnović, *Efficiency and the Redistribution of Welfare*, Tech. Rep. (Microsoft Research, 2011).
- [14] G. W. Brown, *Iterative solution of games by fictitious play*, *Activity analysis of production and allocation* **13**, 374 (1951).
- [15] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*, MIT Press Books, Vol. 1 (The MIT Press, 1994).
- [16] S.-f. Cheng, D. M. Reeves, Y. Vorobeychik, and M. P. Wellman, *Notes on Equilibria in Symmetric Games*, in *Proceedings of the 6th International Workshop On Game Theoretic And Decision Theoretic Agents* (2004) pp. 1–5.
- [17] D. Monderer and L. S. Shapley, *Potential games*, *Games and Economic Behavior* **14**, 124 (1996).
- [18] J. Nash, *Non-Cooperative Games*, *The Annals of Mathematics Second Series*, **54**, 286 (1951).
- [19] I. L. Glicksberg, *A Further Generalization of the Kakutani Fixed Point Theorem, with Application to Nash Equilibrium Points*, *Proceedings of the American Mathematical Society* **3**, 170 (1952).
- [20] E. Maskin and P. Dasgupta, *The existence of equilibrium in discontinuous economic games, part i (theory)*, *Review of Economic Studies* **53**, 1 (1986), reprinted in K. Binmore and P. Dasgupta (eds.), *Economic Organizations as Games*, Oxford: Basil Blackwell, 1986, pp. 48-82.
- [21] R. J. Aumann, *Subjectivity and correlation in randomized strategies*, *Journal of Mathematical Economics* **1**, 67 (1974).
- [22] H. Moulin and J. P. Vial, *Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon*, *International Journal of Game Theory* **7**, 201 (1978).
- [23] J. Robinson, *An iterative method of solving a game*, *Annals of Mathematics Second Series*, **54**, pp. 296 (1951).

- [24] J. M. Danskin, *Fictitious play for continuous games revisited*, International Journal of Game Theory **10**, 147 (1981).
- [25] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, *Algorithmic Game Theory* (Cambridge University Press, 2007).
- [26] Y. Zick and E. Elkind, *Arbitrators in overlapping coalition formation games*, in AA-MAS, edited by L. Sonenberg, P. Stone, K. Tumer, and P. Yolum (IFAAMAS, 2011) pp. 55–62.
- [27] Y. Zick, G. Chalkiadakis, and E. Elkind, *Overlapping coalition formation games: charting the tractability frontier*, in AAMAS, edited by W. van der Hoek, L. Padgham, V. Conitzer, and M. Winikoff (IFAAMAS, 2012) pp. 787–794.
- [28] E. Koutsoupias and C. Papadimitriou, *Worst-case equilibria*, in 16th Annual Symposium on Theoretical Aspects of Computer Science (Trier, Germany, 1999) pp. 404–413.
- [29] M. R. Baye, D. Kovenock, and C. G. de Vries, *The all-pay auction with complete information*, Economic Theory **8**, 291 (1996).
- [30] E. Anshelevich and M. Hoefer, *Contribution games in networks*, Algorithmica **63**, 51 (2011).
- [31] R. Smorodinsky and M. Tennenholtz, *Overcoming free riding in multi-party computations – the anonymous case*, Games and Economic Behavior **55**, 385 (2006).
- [32] S. Mannor, *Advanced topics in systems, learning and control*, Lecture notes (2008).
- [33] D. Monderer and L. S. Shapley, *Fictitious play property for games with identical interests*, Journal of Economic Theory **68**, 258 (1996).
- [34] L. Shapley and R. C. S. M. CALIE, *Some Topics in Two-person Games*, Memorandum (Rand Corporation) (Rand Corporation, 1963).

APPENDIX

A. OMITTED PROOFS

We give the previously omitted proofs here, in the order of appearance.

We shall now prove Theorem 7. To remain self-contained, before proving the theorem, we bring here the necessary definitions used by [20]. Given player i with the strategy set $A_i \subseteq \mathbb{R}^m$, define $A \triangleq A_1 \times \dots \times A_n$.

Definition 5. For each pair of agents $i, j \in 1, \dots, n$, let $D(i)$ be a positive natural, and for a $d \in \{1, \dots, D(i)\}$, let $f_{i,j}^d: \mathbb{R} \rightarrow \mathbb{R}$ be continuous, such that $(f_{i,j}^d)^{-1} = f_{j,i}^d$. For every player i , we define

$$A^*(i) \triangleq \{(a_1, \dots, a_n) \in A \mid \exists j \neq i, \exists k \in \{1, \dots, m\}, \exists d \in \{1, \dots, D(i)\}, \\ \text{such that } a_{j,k} = f_{i,j}^d(a_{i,k})\}. \quad (2.4)$$

We now define weakly lower semi-continuity, which intuitively means that there is a set of directions, such that approaching a point from any of these directions never causes the function to jump upward.

Definition 6. Let $B^m \triangleq \{z \in \mathbb{R}^m \mid \sum_{i=1}^m z_i^2\} = 1$, meaning it is the surface of the unit sphere centered at zero. Let $e \in B^m$ and $\theta > 0$. Function $g_i(a_i, a_{-i})$ is weakly lower semi-continuous in the coordinates of a_i if for all \hat{a}_i there exists an absolutely continuous measure ν on B^m , such that for all a_{-i} , we have

$$\int_{B^m} \left\{ \liminf_{\theta \rightarrow 0} g_i(\hat{a}_i + \theta e, a_{-i}) d\nu(e) \right\} \geq g_i(\hat{a}_i, a_{-i}).$$

Finally, we are ready to prove Theorem 7.

Proof. We show now that all the conditions of Theorem 5* from [20] hold. First, the strategy set of player i is simplex, and as such, it is non-empty, convex and compact. The utility function of player i is discontinuous only at a threshold of one of the projects. These point belong to the set $A^*(i)$, defined in Formula (2.4), if we take

$$\begin{aligned} D(i) &\triangleq 2; \\ f_{i,j}^1(y) &\triangleq y \begin{cases} \theta & \text{if } i < j, \\ 1/\theta & \text{if } i > j; \end{cases} \\ f_{i,j}^2(y) &\triangleq y \begin{cases} \theta & \text{if } j < i, \\ 1/\theta & \text{if } j > i. \end{cases} \end{aligned}$$

The sum of all the utilities is a continuous function. In addition, the utility of player i is bounded by the largest project's value when all the players contribute their budgets there. It is also weakly lower semi-continuous in i 's contribution, since if we take the measure ν to be $\nu(S) \triangleq \lambda(S \cap B^m_+)$, where λ is the Lebesgue measure on B^m and $B^m_+ \triangleq \{z \in B^m \mid z_l \geq 0, \forall l = 1, \dots, m\}$, we obtain an absolutely continuous measure ν , such that the integral sum only the convergences to a point from the positive directions, and such convergences will never cause a jump up, since our model causes jumping only down when reducing contributions.

We have shown that all the conditions of Theorem 5* hold, which implies that a mixed NE exists. \square

The proof of Lemma 1 follows.

Proof. First, at least one player contributes somewhere, since otherwise any positive contribution would be a profitable deviation for every player (Recall that all α_i 's and B_j 's are positive.) Moreover, at least one of the players contributes to a steep project, for the following reasons. If only the non-steep projects receive a contribution, then take any such project p . If a single player contributes there, then this player would benefit from moving to contribute to a vacant steep project. If both players contribute to p , then if one is suppressed, it would like to deviate to any project where it would not be suppressed, and if no-one is suppressed, then a player who contributes not less would like to contribute to a vacant steep project instead.

We prove part 2 now. Let $i \in N$ be any non-suppressed player among those who contribute to a steep project, w.l.o.g. – to project m . Assume first that player $j \neq i$ is not suppressed. Then, for any non-steep project where i contributes, i contributes either alone or precisely the least amount it should contribute to achieve a portion in the project's value, because otherwise i would like to increase its contribution to m on the expense of decreasing its contribution to the considered non-steep project.

Now, consider the case where j is suppressed. Then, even if j contributes to a non-steep project where i contributes (and is suppressed there), i still will prefer to move some budget from this project to m , since i receives the whole value of m as well. Thus, this cannot be an NE. □

The proof of Lemma 2 appears now.

Proof. Since $B_1 \geq \theta B_2$, no player is suppressed, because any player prefers not being suppressed, and at any project, a player who concentrates all its value there is not suppressed.

Every steep project receives a positive contribution from each player, for the following reasons. If only a single player contributes to a steep project, then the player who does not contribute there will profit from contributing there exactly the threshold value, while leaving at least the threshold values at all the projects where it contributed. There is always a sufficient surplus to reach the threshold because $B_1 \geq \theta B_2$. If no player contributes to a steep project p , then there exists another steep project q , where two players contribute, according to part 1 of Lemma 1 and what we have just described. The player who contributes there strictly more than the threshold would profit from moving some part of his contribution from q to p , still remaining not less than the threshold on p , contradictory to having an NE.

We next prove the second part of the lemma. Since both players are non-suppressed contributors to steep projects, then, according to part 2 in Lemma 1, we conclude that there exist no non-steep projects where j and i contribute together. □

We now present the proof of Lemma 3.

Proof. We prove part 1 first. Consider an NE profile. Assume to the contrary that player 1 contributes to a steep project, w.l.o.g., to project m . Since $B_1 < \theta B_2$ and no player is suppressed, player 2 could transfer to m budget from other projects, such that at each project, where 2 was obtaining a share of a revenue, 2 still obtains a share, and 2 suppresses 1 at m . This would increase 2's utility, contrary to the assumption of an NE.

Now, we prove parts 2 and 3. Since 1 does not contribute to steep projects, part 1 of Lemma 1 implies that 2 contributes to a steep project (say, 2 contributes $y > 0$ to project m), and part 2 of Lemma 1 implies that if 2 contributes to a non-steep project, it contributes there either alone or precisely the least amount it should contribute to achieve a portion in the project's value. Since contributing alone is strictly worse than contributing this budget to a steep project, 2 may only contribute together with 1, the threshold amount. Therefore, 1 receives a positive value in this profile, and we have part 2. The only thing left to prove here is that 2 contributes to each steep project. If not, 1

would prefer to move some of its contribution there, in contradiction to the assumption of an NE.

We prove part 4 now. Assume to the contrary that a non steep and non 2-steep project receives a contribution. We proved in part 3 that 1 contributes there, alone or not. For her, moving a small enough utility to a 2-steep project would increase her utility, regardless whether 2 contributes to any of those projects. This is so because if 2 contributes together with 1, it contributes precisely the threshold amount, according to part 3. This incentive to deviate contradicts the assumption of an NE.

We prove part 5 now. If 2 contributes to a 2-steep project p , then there may not exist another 2-steep project, since otherwise 1 would like to transfer a small amount from p to another 2-steep project q , such that without losing a share of the value of p , player 1 gets the whole value of project q . □

The proof of Lemma 4 appears now.

Proof. Before moving, player i 's utility is $\sum_{q \in \Omega} (\frac{1}{2} \text{ or } 1) \alpha_q \cdot (x_q^1 + x_q^2)$.

We begin by proving part 1. Assume $\frac{1}{2} \alpha_m \geq \alpha_p$. If i moves $\delta > 0$ from the steep projects to p , then its utility from the steep projects decreases by at least $0.5 \alpha_m \delta$, and its utility from p increases by $\alpha_p \delta$. The total change is $(-0.5 \alpha_m + \alpha_p) \delta$, and since $\frac{1}{2} \alpha_m \geq \alpha_p$, this is non-positive.

We prove part 2 now. Moving δ from p to a subset of the steep projects decreases the utility of i by $\alpha_p \delta$ and increases it by at least $0.5 \alpha_m \delta$. Since $\frac{1}{2} \alpha_m > \alpha_p$, the sum of these is (strictly positive). If this move suppresses j , then the increase is more than $0.5 \alpha_m \delta$, thus requiring $\frac{1}{2} \alpha_m \geq \alpha_p$ is enough.

To prove part 3, assume that $\frac{1}{2} \alpha_m < \alpha_p$ and we can take $\delta > 0$ from some of the steep projects where i receives half of the value so as to keep receiving a half of the new value. Then, moving this δ to p decreases i 's utility from the steep projects by $0.5 \alpha_m \delta$ and its utility from p increases by $\alpha_p \delta$. The total change is $(-0.5 \alpha_m + \alpha_p) \delta$, and since $\frac{1}{2} \alpha_m < \alpha_p$, this is (strictly) positive. □

We prove Theorem 5 now.

Proof. We first prove case 1. Consider any NE. By Lemma 2, each player contributes to all steep projects and if it contributes to a non-steep project, then it is the only contributor there. Take any non-steep project p , where someone contributes, say player i . Consider moving all what player i contributes to p to a steep project. If the other player gets suppressed as a result of this move, then, according to part 2 of Lemma 4, there could be no contribution to p , since this move is profitable. If the other player does not get suppressed as a result of this move, then according to part 1 of Lemma 4, this move is weakly profitable. Since we began at an NE, this move does not change the social welfare. This can be done for any contribution to non-steep projects, and therefore, PoA = 1. We have fully proven case 1.

Consider case 2a now. From the proof of the existence on an NE in this case, we know that 1 investing all its budget in $m - k$ and 2 dividing its budget equally between the steep project constitute an NE. Thus, PoS $\geq \frac{\alpha_m B_2 + \alpha_{m-k} B_1}{\alpha_m (B_1 + B_2)}$. Since $\frac{\alpha_{m-k}}{\alpha_m} \leq \frac{1}{1+\theta}$, not

all the projects have equal value functions. Therefore, in an NE no player is suppressed, since if 2 dominated 1, then 2 would have to invest more than B_1/θ in each project, and 2 would like to deviate to contribute to steep projects more. Since no player is suppressed, we conclude from part 1 of Lemma 3 that player 1 never contributes to a steep project in an NE, and thus $\text{PoS} = \frac{\alpha_m B_2 + \alpha_{m-k} B_1}{\alpha_m (B_1 + B_2)}$.

Next, let us approach the price of anarchy. According to Lemma 3, the only way to reduce the efficiency relatively to the price of stability is for 2 to invest a 2-steep project. If this happens, then we obtain that case 2b must hold, exactly as it is done in the proof of the other direction of Theorem 3. Therefore, if this case does not hold, then $\text{PoS} = \text{PoS}$. If it does, then we have the NE when 2 invests θB_1 in project $m - k$ (and 1 invests all its budget there, and 2 equally divides the rest of its budget between the steep projects), which yields the price of anarchy of $\frac{\alpha_m (B_2 - \theta B_1) + \alpha_{m-k} (B_1 (1 + \theta))}{\alpha_m (B_1 + B_2)}$.

We prove case 2b now. We show in the proof of case 2b of Theorem 3, player 1 investing all its budget in $m - k$ and 2 investing θB_1 in $m - k$ and uniformly dividing the rest between the steep projects is an NE. Thus, $\text{PoS} \geq \frac{\alpha_m (B_2 - \theta B_1) + \alpha_{m-k} (B_1 (1 + \theta))}{\alpha_m (B_1 + B_2)}$. Since $\frac{\alpha_{m-k}}{\alpha_m} \leq \frac{2(1-\theta)}{2-\theta}$, we conclude analogously to what we did in the proof of the previous case that no player is suppressed. Thus, Lemma 3 implies that the only way to achieve a more efficient NE is for 2 to contribute only to the steep projects. If this is an NE, then we obtain that case 2a must hold, exactly as it is done in the proof of the other direction of Theorem 3. Therefore, if this does not hold, we have $\text{PoS} = \frac{\alpha_m (B_2 - \theta B_1) + \alpha_{m-k} (B_1 (1 + \theta))}{\alpha_m (B_1 + B_2)}$. If case 2a does hold, then we know that the profile where 1 invests all its budget in $m - k$ and 2 divides its budget between the steep projects is an NE, and thus $\text{PoS} = \frac{\alpha_m B_2 + \alpha_{m-k} B_1}{\alpha_m (B_1 + B_2)}$.

We turn to the price of anarchy now. According to Lemma 3, the NE with player 1 investing all its budget in $m - k$ and 2 investing θB_1 in $m - k$ and uniformly dividing the rest between the steep projects is the worst possible NE, and thus $\text{PoA} = \frac{\alpha_m (B_2 - \theta B_1) + \alpha_{m-k} (B_1 (1 + \theta))}{\alpha_m (B_1 + B_2)}$.

Finally, in the case 2c we know that 2 dividing its budget equally and 1 contributing all its budget is an NE, and therefore $\text{PoS} = 1$. To find the price of anarchy, recall that if 2 does as before while 1 invests nothing at all, it still is an NE, and thus $\text{PoA} \leq \frac{\alpha_m B_2}{\alpha_m (B_1 + B_2)}$. Since 2 always gets at least $\alpha_m B_2$ in any NE, the price of anarchy cannot decrease below it, and thus $\text{PoA} = \frac{\alpha_m B_2}{\alpha_m (B_1 + B_2)}$. \square

Let us now prove Corollary 2.

Proof. The maxima are obtained in case 1 of Theorem 5.

To find the infima, find the infimum in every case, substituting the extreme values in the expressions for PoS and PoA. We begin with the PoS. In case 2a, the infimum of the PoS is obtained for $\alpha_{m-k} = 0$ and $B_1 = \frac{\theta B_2}{k}$, and it is $\frac{k}{k+\theta}$. In case 2b, the infimum of the PoS is the minimum of the following two expressions

1. $\frac{\alpha_m B_2 + \alpha_{m-k} B_1}{\alpha_m (B_1 + B_2)}$ when $\alpha_{m-k} = \frac{2\theta}{1+\theta} \alpha_m$ and $B_1 = \frac{\theta B_2}{k+\theta^2}$, which is $\frac{k+\theta^2 + \frac{2\theta^2}{1+\theta}}{k+\theta+\theta^2}$.
2. $\frac{\alpha_m (B_2 - \theta B_1) + \alpha_{m-k} (B_1 (1 + \theta))}{\alpha_m (B_1 + B_2)}$ when $\alpha_{m-k} = \frac{2\theta}{1+\theta} \alpha_m$ and $B_1 = \frac{\theta B_2}{k+\theta^2}$, which is $\frac{k+2\theta^2}{k+\theta+\theta^2}$.

The minimum of these expressions is $\frac{k+2\theta^2}{k+\theta+\theta^2}$. Finally, the infimum of the price of stability in case 2c is 1. The absolute infimum is the minimum of these three expressions, which is $\frac{k}{k+\theta}$.

We consider now the infimum of the price of anarchy. In case 2a, the infimum of the PoA is obtained as follows:

1. If also case 2b holds, then it is the value of $\frac{\alpha_m(B_2-\theta B_1)+\alpha_{m-k}(B_1(1+\theta))}{\alpha_m(B_1+B_2)}$ when $\alpha_{m-k} = \frac{2\theta}{1+\theta} \alpha_m$ and $B_1 = \frac{\theta B_2}{k+\theta^2}$, which is $\frac{k+2\theta^2}{k+\theta+\theta^2}$.
2. Otherwise, PoA = PoS, and so the infimum is $\frac{k}{k+\theta}$.

The minimum of these two expressions is $\frac{k}{k+\theta}$. In case 2b, the infimum of the PoA is $\frac{\alpha_m(B_2-\theta B_1)+\alpha_{m-k}(B_1(1+\theta))}{\alpha_m(B_1+B_2)}$ when $\alpha_{m-k} = \frac{2\theta}{1+\theta} \alpha_m$ and $B_1 = \frac{\theta B_2}{k+\theta^2}$, which is $\frac{k+2\theta^2}{k+\theta+\theta^2}$. In case 2c it is $\frac{B_2}{B_1+B_2}$ when $B_1 = \frac{\theta}{m} B_2$, which is $\frac{m}{m+\theta}$. Therefore, the infimum of the price of anarchy is $\frac{k}{k+\theta}$. \square

We now present the proof of Theorem 6.

Proof. We first prove case 1. According to proof of case 1 in Theorem 4, equally dividing all the budgets between the steep projects is an NE. Therefore, PoS = 1.

We consider case 2 now. We know that n dividing its budget equally and all the other players contributing all their budgets is an NE, and therefore PoS = 1. To find the price of anarchy, recall that if n does as before while all the other players invest nothing at all, it still is an NE, and thus $\text{PoA} \leq \frac{\alpha_m B_n}{\alpha_m (\sum_{i \in \{1, 2, \dots, n\}} B_i)}$. Since n always gets at least $\alpha_m B_n$ in any NE, the price of anarchy cannot decrease below it, and thus $\text{PoA} = \frac{\alpha_m B_n}{\alpha_m (\sum_{i \in \{1, 2, \dots, n\}} B_i)}$. \square

The proof of Theorem 8 is presented now.

Proof. The utility of i is $u^i(x^i) = \phi_\psi^i(x^i) + \phi_\omega^i(B_i - x^i)$. Consider the open intervals between the consecutive points of L_{B_i} . On each of these intervals, the function $\phi_\psi^i(x^i)$ is convex, being proportional to the convex project value function, and $\phi_\omega^i(B_i - x^i)$ is convex because the function $B_i - x$ is convex and concave and ϕ_ω^i is convex and weakly monotone. Therefore, the utility is also convex, as the sum of convex functions.

Therefore, the utility's supremum of the closure of such a convexity interval is attained as the one-sided limit of at least one of its edge points. This supremum can be a maximum if and only if it is not larger than the maximum of the utility at an interval edge point or at an internal point of an interval (in the last case, the convexity implies that the utility is constant on this interval). \square

B. SIMULATION RESULTS

We now present some additional simulation results; they further support the conclusions of Section 5.4. First, we run over the ratios of the project function coefficients and the ratios of the two largest budgets. Then, we go over the the project coefficients, and finally, we go over the two largest budgets.

First, we plot the existence and efficiency of NE as function of the ratio of the project function coefficients and of the ratio of the two largest budgets. We compare these predictions with the theoretical predictions, which are precise only for two players, and otherwise, they just show when we definitely have an NE, while stating nothing about the other cases. The first row is contains the simulation results, and the second row plots the theoretical predictions. The case of 3 players is shown in Figure 2.8, and 4 are given in Figure 2.9.

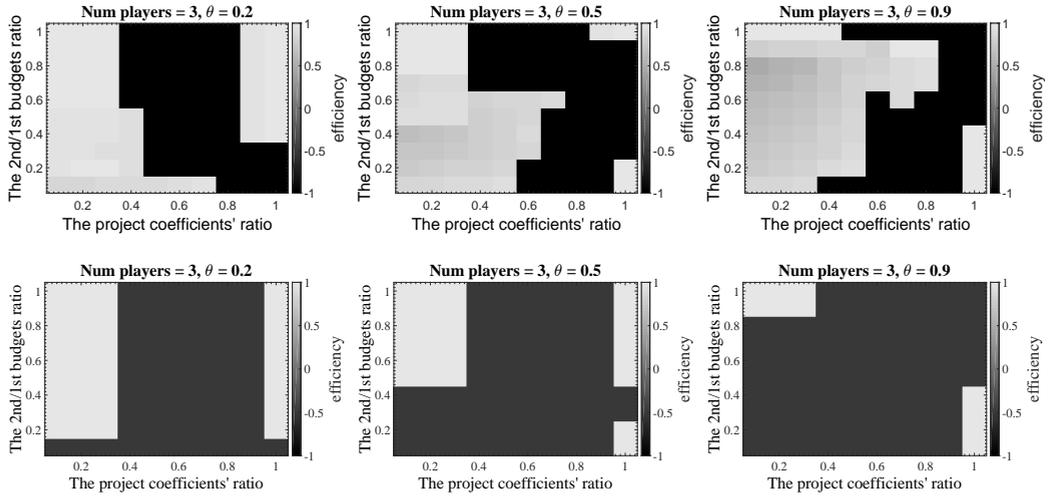


Figure 2.8: The existence and efficiency of NE for 3 players as a function of the ratio of the project function coefficients and the ratio of the two largest budgets. The first row plots the results of the simulations, and the second row shows the theoretical predictions. In the results of the simulations, black color means that Nash Equilibrium has not been found. For all the other cases, the efficiency, a value in (0, 1], is shown by the shade (from dark gray = 0.0 to white = 1.0). In the plot of the theoretical predictions, black color means that we do not know, while white means that a Nash Equilibrium exists.

We now present some more experiments where project coefficients are being run over. The existence and efficiency is shown in Figure 2.10, Figure 2.11, and Figure 2.12, and the efficiency is plotted as a function of the ratio of the project coefficients in Figure 2.13, Figure 2.14, and Figure 2.15.

We now show some more experiments where the two largest budgets are being run over, and the rest are spread on equal distances. The existence and efficiency is shown in Figure 2.16, and the efficiency is plotted as a function of the ratio of the two largest budgets in Figure 2.17.

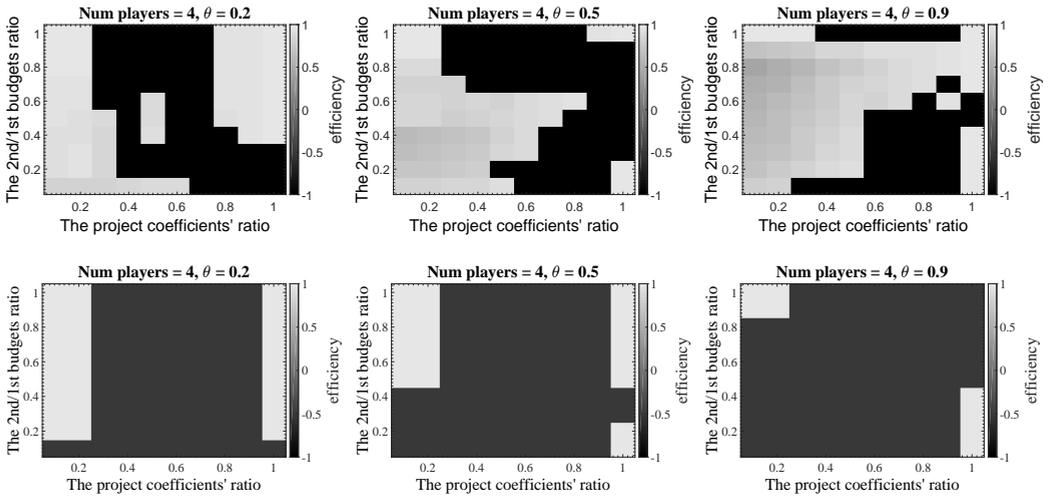


Figure 2.9: The existence and efficiency of NE for 4 players as a function of the ratio of the project function coefficients and the ratio of the two largest budgets. The first row plots the results of the simulations, and the second row shows the theoretical predictions. In the results of the simulations, black color means that Nash Equilibrium has not been found. For all the other cases, the efficiency, a value in (0, 1), is shown by the shade (from dark gray = 0.0 to white = 1.0). In the plot of the theoretical predictions, black color means that we do not know, while white means that a Nash Equilibrium exists.

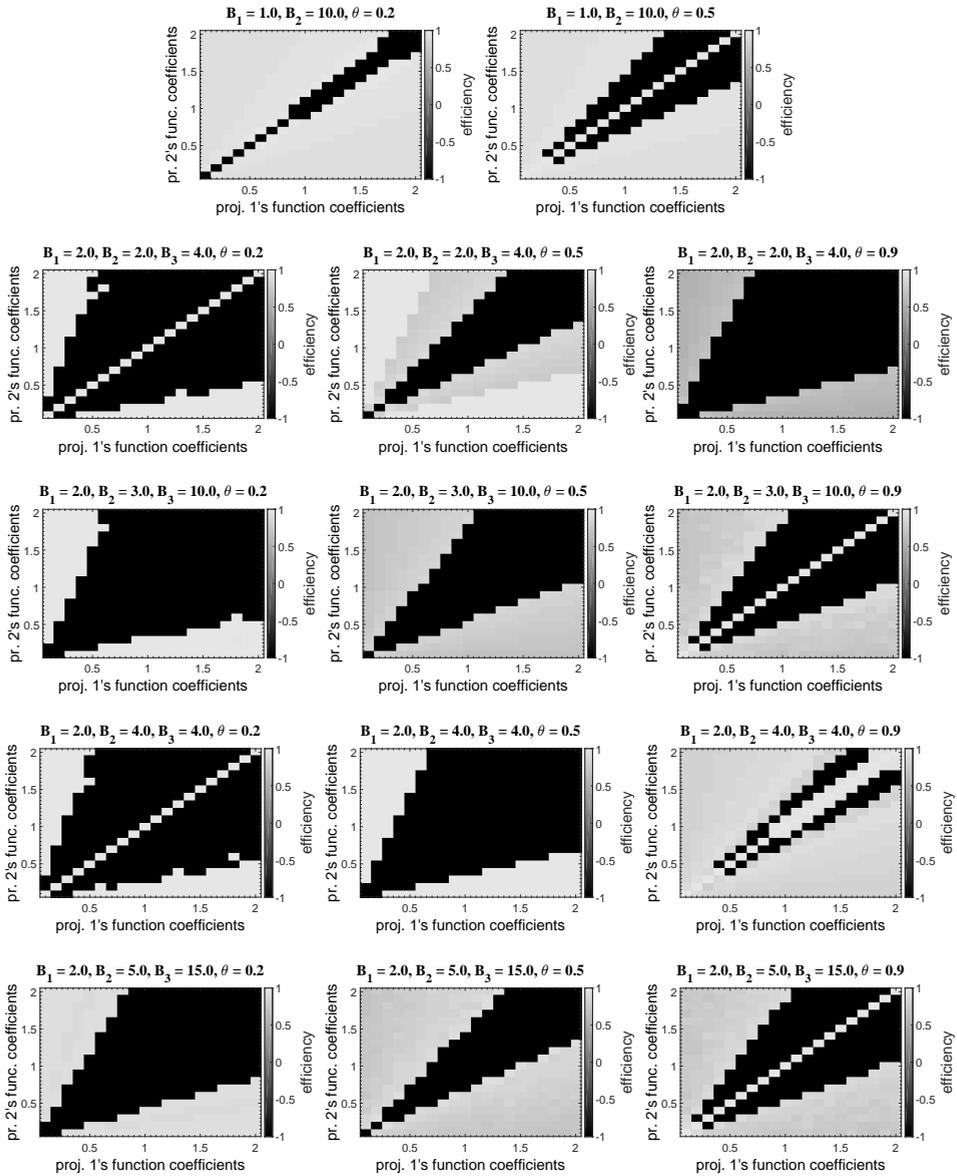


Figure 2.10: The existence and efficiency of NE as function of project functions for 2 and 3 players. Black color means that Nash Equilibrium has not been found. For all the other cases, the efficiency, a value in $(0, 1]$, is shown by the shade (from dark gray = 0.0 to white = 1.0).

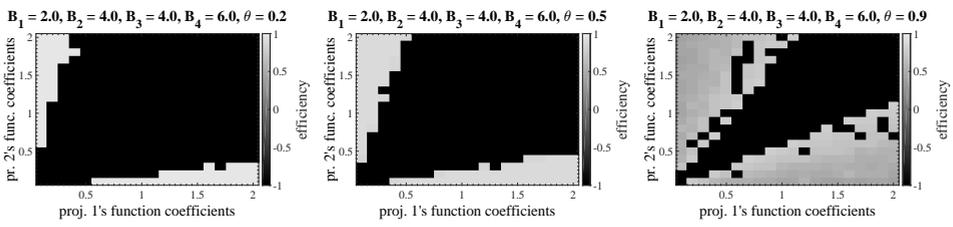


Figure 2.11: The existence and efficiency of NE as function of project functions for 4 players.

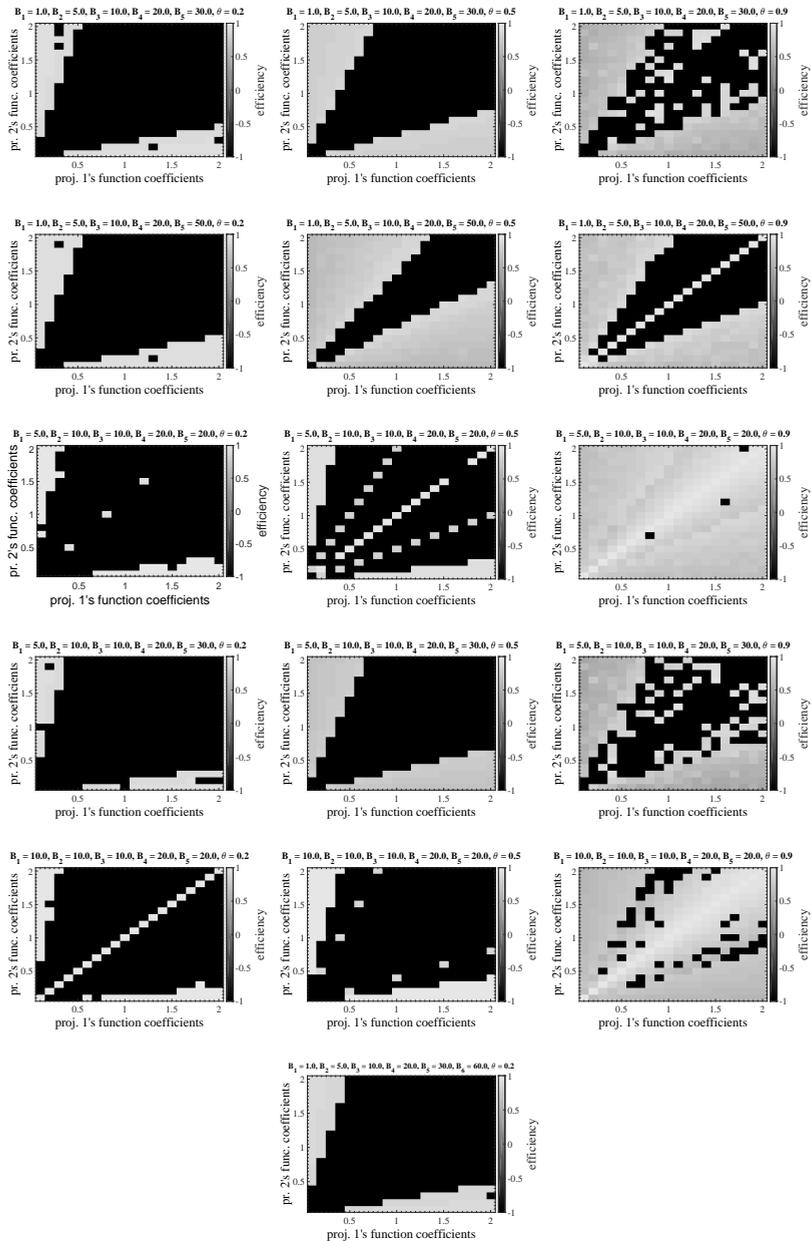


Figure 2.12: The existence and efficiency of NE as function of project functions for 5 and 6 players.

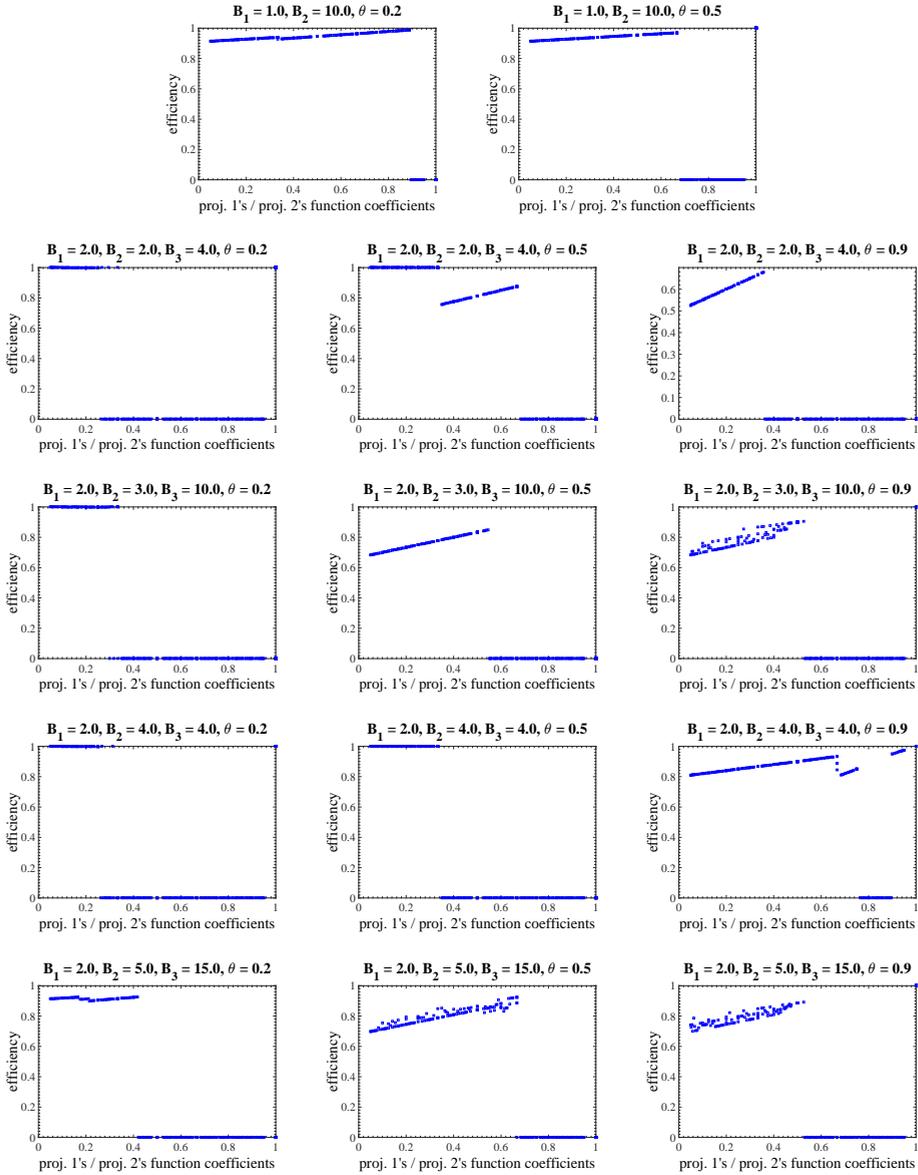


Figure 2.13: The efficiency as a function of project values' ratio for 2 and 3 players. Efficiency of 0 means that Nash Equilibrium has not been found. For all the other cases, the efficiency, a value in (0, 1], is shown.

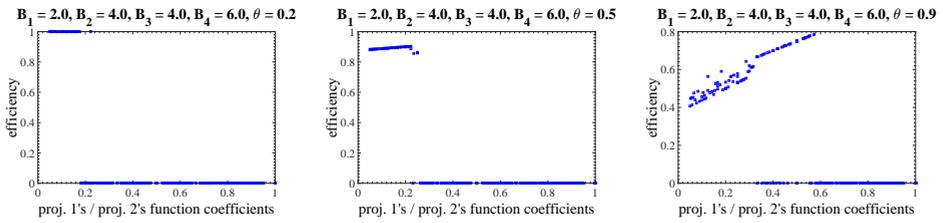


Figure 2.14: The efficiency as a function of project values' ratio for 4 players.

2

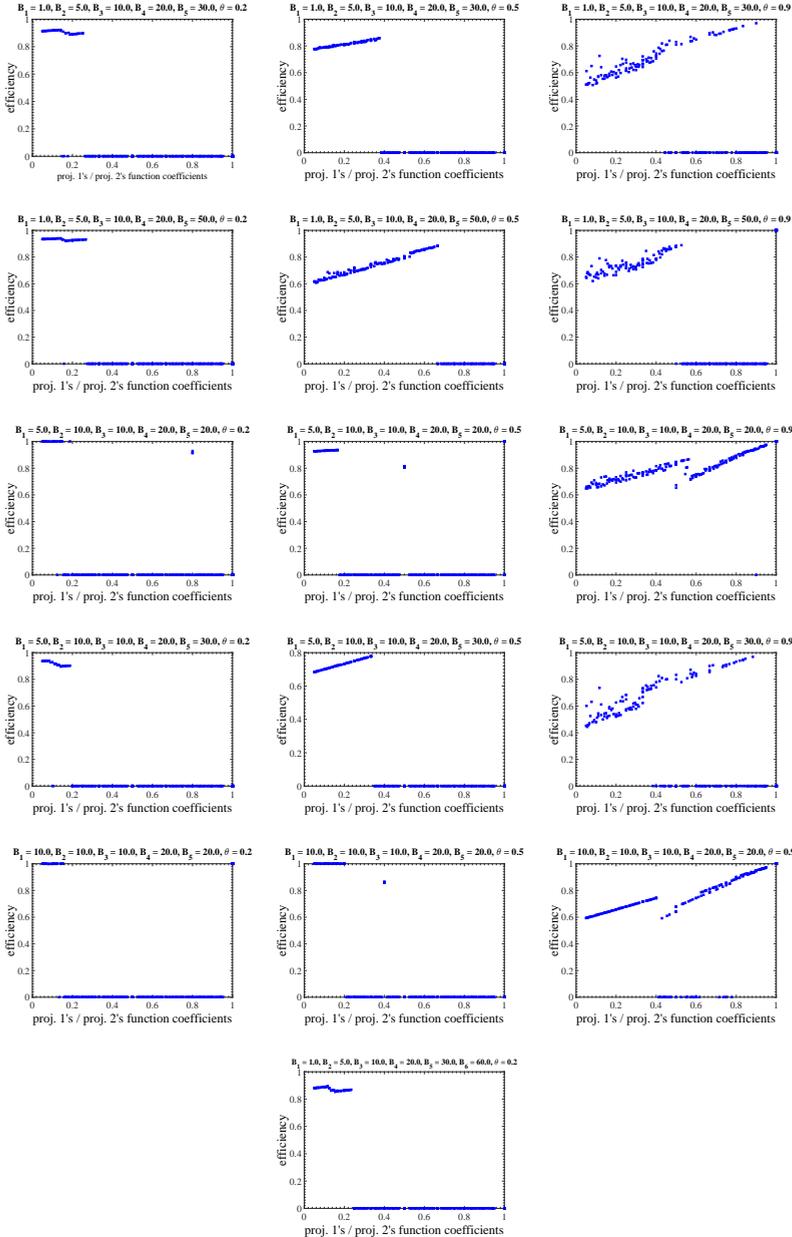


Figure 2.15: The efficiency as a function of project values' ratio for 5 and 6 players.

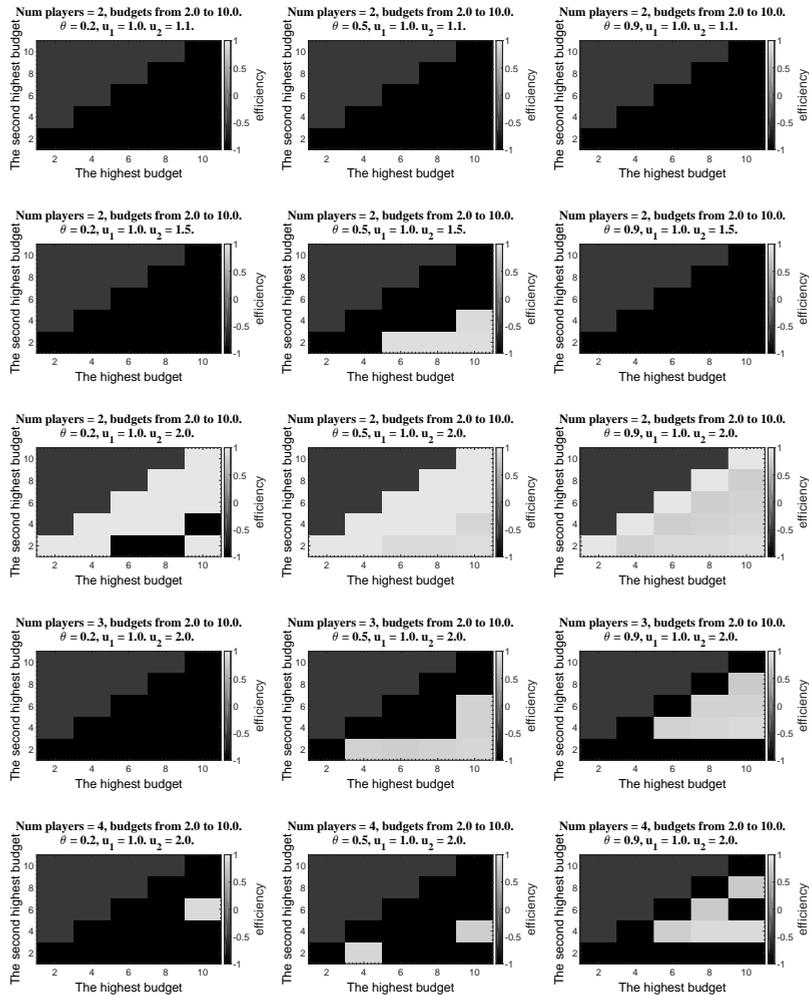


Figure 2.16: The existence and efficiency of NE as function of the largest and the second largest budgets. Black color means that Nash Equilibrium has not been found, and very dark grey color means the area is not defined, since the second highest budget may not be larger than the highest one. For all the other cases, the efficiency, a value in $(0, 1]$, is shown by the shade (from dark grey = 0.0 to white = 1.0).

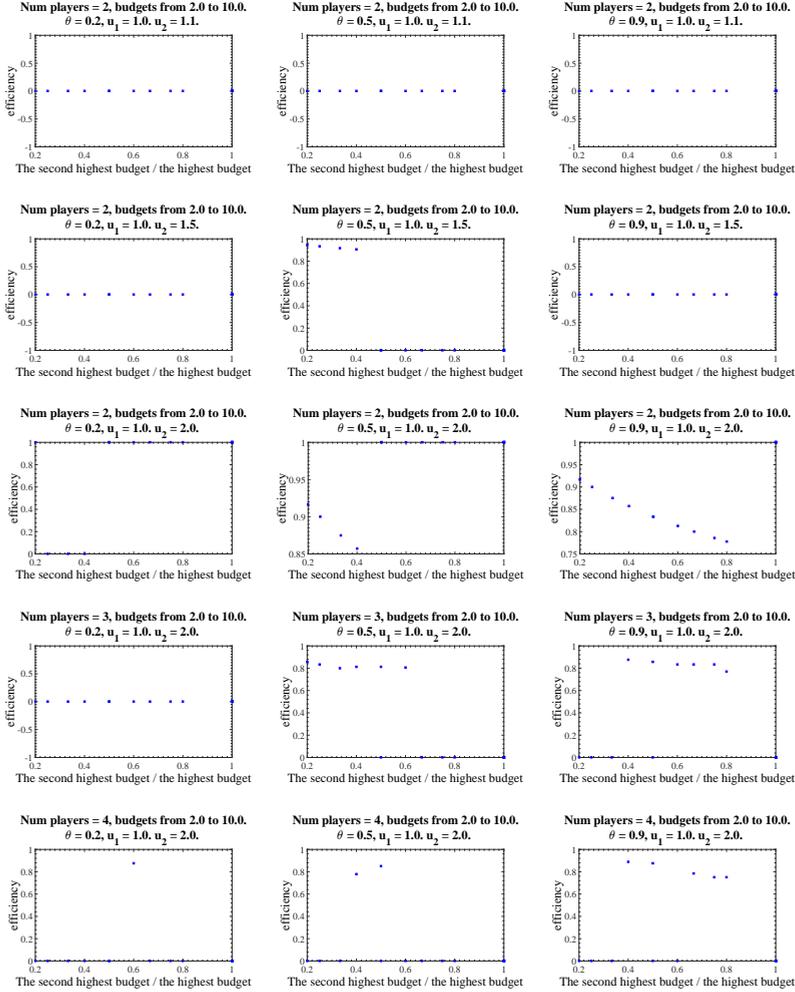


Figure 2.17: The efficiency as a function of the ratio of the two largest budgets. Efficiency of 0 means that Nash Equilibrium has not been found, and -0.5 means the the domain is not defined, since the second highest budget may not be larger than the highest one. For all the other cases, the efficiency, a value in $(0, 1)$, is shown.

3

COMPETITION BETWEEN COOPERATIVE PROJECTS

Time is what we want most, but what we use worst.

William Penn, 1693

A paper needs to be good enough to be published; a grant proposal needs to be sufficiently convincing compared to the other proposals, in order to get funded. Papers and proposals are examples of cooperative projects that compete with each other and require effort from the involved agents, while often these agents need to divide their efforts across several such projects. We aim to provide advice how an agent can improve her utility and how the designer of such a competition (e.g., the program chairs or funding agency) can create the conditions under which a socially optimal outcome can be obtained. We therefore extend a model for dividing effort across projects with two types of competition: a quota or a success threshold. For these two types of games we prove conditions for equilibrium existence and efficiency. Additionally we find that competitions using a success threshold can more often have an efficient equilibrium than those using a quota. We also show that when a socially optimal Nash equilibrium exists, there exist inefficient equilibria as well. Therefore, regulation may be needed to choose the efficient equilibrium.

1. INTRODUCTION

The previous chapter considers shared effort in projects. In this chapter we extend the model by incorporating the requirements that such projects face. Now, a project obtains its value only if it stands up to certain requirements, such as being one of the best projects or being at least at a given level.

Cooperative projects often are in competition with each other. For example, a paper needs to have a certain quality, or to be among a certain number of the best papers to be published, and a grant needs to be one of the best to be awarded. Generally speaking, either only projects that achieve a certain *minimum level*, or those that are among a certain *quota* of the best projects attain their value. A quota can be expressed in other ways, such as a success rate and the number of the total projects. Agents endowed with a resource budget (such as time) need to divide this resource across several such projects. In this chapter we consider so-called public projects where agents may contribute resources to create something together. If such a project stands up to competition, its rewards are typically divided among the contributors based on their individual investments.

We often see agents who divide effort across competing projects. In addition to co-authoring articles or books [1–3] and research proposals, examples include participating in crowdsensing projects [4] and online communities [5]. Examples of quotas for successful projects include supporting politicians being elected to a fixed-size committee [6], or investing effort in manufacturing several products, where the market becomes saturated with a certain amount of products. More examples of success thresholds are investing in start-ups, where a minimum investment is needed to survive, or contributing to educational projects, where a minimum level is required for a project to succeed.

The ubiquity and the complexity of such competing projects calls for a decision-support system, helping agents to divide their efforts wisely. Assuming rationality of all the others, an agent would like to know how to behave, and the designer of the competition would like to know which rules lead to better results. In the language of game theory, for each setting, we are interested in the equilibria and their efficiency.

Chapter 2 uses the model of Bachrach, Syrgkanis and Vojnović [7] of a shared effort game, where each player has a budget to divide among a given set of projects. The game possesses a contribution threshold θ , and the project's value is equally shared among the players who invest above this threshold. They analyzed Nash equilibria (NE) and their price of anarchy (PoA) and stability (PoS) for such games.¹ However, they ignored that there may be competition between the projects. In order to model this aspect, we extend their model by allowing the projects only to obtain their modeled value if they stand up to a competition.

Another related model, called an all-pay auction, is a constrained model where only one contributor benefits from the project. Its equilibria are analyzed by Baye, Kovenock and de Vries [12] and studied by many others. A famous example is the colonel Blotto game with two players [13], where these players spread their forces among several battlefields, winning a battle if allocating it more forces than the opponent does. The relative

¹The price of anarchy [8, 9] is the ratio of the minimum social welfare in a Nash equilibrium to the maximum possible social welfare. The price of stability [10, 11] is the ratio of the maximum social welfare in a NE relative to the maximum possible social welfare.

number of won battles determines the player's utility. Anshelevich and Hoefer [14] generalize this by modeling many of such two-player games by an undirected graph where nodes are the players, and edges are the projects. A project, being an edge, obtains contributions from two players. They especially concentrated on minimum-effort projects. Their work proves the existence of an NE and shows that the price of anarchy is at most 2 in such bilateral projects.

In the context of publishing papers, Batchelor [15] states that high standards are important to the scientific progress, and that "an appreciable proportion of those that do find their way into the scientific literature are not worthy of publication", that is, that the standards would better be increased. However, in addition to maximizing the total value of the published papers, he considers goals such as reducing the amount of noise (low quality publications). In order to characterize the influence of a quota or a success threshold on the efficient strategies for the individual agents and their society, we analyze the NE and their efficiency in the modeled games. Discovering the structure of equilibria in various quotas or success thresholds can bear significant impacts on the efficiency of investing time and effort in the mentioned enterprises and, therefore, on a sizable part of life.

Compared to Kleinberg and Oren [3], we model contributing to more than one project by an agent, and concentrate on the competition, rather than on sharing a project's utility. We also emphasize that unlike devising division rules to make people contribute properly, studied in cooperative game theory (see *Shapley value* [16] for a prominent example), we model given division rules and analyze the obtained game.

We allow any number of agents to contribute to projects, and we introduce two models of competition, a quota or a success threshold.

1. Given a quota q , only q projects receive their value. This models the limit on the number of papers to be accepted for a presentation, the number of politicians in a city council, or the number of projects an organization can fund.
2. There exists a success threshold δ , such that only the projects that have a value of at least δ actually receive their value. This models a paper or proposal acceptance process that is purely based on quality.

For these two types of games we prove conditions for equilibrium existence. A crucial question that we then answer is how adjusting such a quota or a success threshold can influence the efficiency of the investments, and thereby the social welfare of the participants. We can then conclude that competitions using a success threshold have efficient equilibria more often than those using a quota.

In order to make things more clear, we use the following running example:

Example 2. Consider scientists (agents) investing time and effort from their time/effort budget in writing papers. A paper attains its value (representing acknowledgment and all the related rewards) if it stands up to the competition with other papers. The competition can mean either being one of the q best papers, or achieving at least the minimum level of δ , depending on the circumstances. A scientist is rewarded by a paper by becoming its co-author if she has contributed enough to that paper.

We formally define the above-mentioned models in Section 2, analyze the Nash equilibria of the first model and their efficiency in Section 3, and analyze the second model in Section 4. Having analyzed both models of competition between projects, we compare their characteristics, the possibility to influence authors' behavior through tuning the acceptance criteria or to influence the political supporters by changing the rules of elections, and draw further conclusions in Section 5.

2. MODEL

We build our model on that from Section 2 of Chapter 2, since that is a model of investment in common projects with a general threshold. We, therefore, first present the model from Section 2 of Chapter 2 for *shared effort games*, which also appears in [7]. From Definition 8 on, we introduce competition between the projects.

There are n players $N = \{1, \dots, n\}$ and a set of projects Ω . Each player $i \in N$ can contribute to any of the projects in Ω_i , where $\emptyset \subsetneq \Omega_i \subseteq \Omega$; the contribution of player i to project $\omega \in \Omega_i$ is denoted by $x_\omega^i \in \mathbb{R}_+$. Each player i has a budget $B_i > 0$, so that the strategy space of player i (i.e., the set of her possible actions) is defined as $\{x^i = (x_\omega^i)_{\omega \in \Omega_i} \in \mathbb{R}_+^{|\Omega_i|} \mid \sum_{\omega \in \Omega_i} x_\omega^i \leq B_i\}$. Denote the strategies of all the players except i by x^{-i} .

The next step to define a game is defining the utilities. Associate each project $\omega \in \Omega$ with its *project function*, which determines its *value*, based on the total contribution vector $x_\omega = (x_\omega^i)_{i \in N}$ that it receives; formally, $P_\omega(x_\omega): \mathbb{R}_+^n \rightarrow \mathbb{R}_+$. The assumption is that every P_ω is both increasing in every parameter. The increasing part stems from the idea that receiving more effort does not make a project worse off. When we write a project function as a function of a single parameter, like $P_\omega(x) = \alpha x$, we assume that project functions P_ω depend only on the $\sum_{i \in N} (x_\omega^i)$, which is denoted by x_ω as well, when it is clear from the context. However, by default, $P_\omega(x_\omega)$ is a function of the whole contribution vector, meaning that $P_\omega(x_\omega): \mathbb{R}_+^n \rightarrow \mathbb{R}_+$. The project's value is distributed among the players in $N_\omega \triangleq \{i \in N \mid \omega \in \Omega_i\}$ according to the following rule. From each project $\omega \in \Omega_i$, each player i gets a share $\phi_\omega^i(x_\omega): \mathbb{R}_+^n \rightarrow \mathbb{R}_+$ with free disposal:

$$\forall \omega \in \Omega: \sum_{i \in N_\omega} \phi_\omega^i(x_\omega) \leq P_\omega(x_\omega). \quad (3.1)$$

We assume that the sharing functions are non-decreasing. The non-decreasing assumption fits the intuition that contributing more does not get the players less.

Denote the vector of all the contributions by $x = (x_\omega^i)_{\omega \in \Omega}^{i \in N}$. The utility of a player $i \in N$ is defined to be

$$u^i(x) \triangleq \sum_{\omega \in \Omega_i} \phi_\omega^i(x_\omega).$$

Consider the numerous applications where a minimum contribution is required to share the revenue, such as paper co-authorship and homework. To concentrate on these applications and to analyze them, we define a specific variant of a shared effort game, called a *θ -sharing mechanism*. This variant is relevant to many applications, including co-authoring papers and participating in crowdsensing projects. For any $\theta \in [0, 1]$, the

players who get a share are defined to be

$$N_\omega^\theta \triangleq \left\{ i \in N_\omega \mid x_\omega^i \geq \theta \cdot \max_{j \in N_\omega} x_\omega^j \right\},$$

which are those who bid at least θ fraction of the maximum bid size to ω . Now, define the θ -equal sharing mechanism as equally dividing the project's value between all the users who contribute to the project at least θ of the maximum bid to the project.

The θ -equal sharing mechanism, denoted by M_{eq}^θ , is

$$\phi_\omega^i(x_\omega) \triangleq \begin{cases} \frac{P_\omega(x_\omega)}{|N_\omega^\theta|} & \text{if } i \in N_\omega^\theta, \\ 0 & \text{otherwise.} \end{cases}$$

Similarly to Section 2 of Chapter 2, to enable theoretical analysis, we now consider θ -equal sharing, where all the project functions are linear with coefficients $\alpha_m \geq \alpha_{m-1} \geq \dots \geq \alpha_1$. We denote the number of projects with the largest coefficient project functions by $k \in \mathbb{N}$, i.e. $\alpha_m = \alpha_{m-1} = \dots = \alpha_{m-k+1} > \alpha_{m-k} \geq \alpha_{m-k-1} \geq \dots \geq \alpha_1$. We call those projects *steep*. Assume w.l.o.g. that $B_n \geq \dots \geq B_2 \geq B_1$.

We shall need the following definitions from Chapter 2. We call a project that receives no contribution in a given profile a *vacant* project.

Definition 7. A player is dominated at a project ω , if it belongs to the set $D_\omega \triangleq N_\omega \setminus N_\omega^\theta$. A player is suppressed at a project ω , if it belongs to the set $S_\omega \triangleq \{i \in N_\omega : x_\omega^i > 0\} \setminus N_\omega^\theta$. That is, a player who is contributing to a project but is dominated there.

We now depart from Section 2 of Chapter 2 and model competition in two different ways.

Definition 8. In the quota model, given a natural number $q > 0$, only the q highest valued projects actually obtain a value to be divided between their contributors. The rest obtain zero. In the case of ties, all the projects that would have belonged to the highest q under some tie breaking rule receive their value; therefore, more than q projects can receive their value in this case. We say a project is in the quota, if its value is among the q highest valued projects; a project is out of the quota otherwise.

The second model is called the success threshold model.

Definition 9. In the success threshold model, given a threshold δ , only the projects with value at least δ obtain a value, while the rest obtain zero.

We use Example 2 to illustrate this model.

Example 2 (Continued). Figure 3.1 depicts a success threshold model, where paper C does not make it to the success threshold, and is, therefore, unpublished, and the contribution it has received is lost. The other two papers are above the success threshold, and get published; such a paper's recognition is equally divided between the contributors who contribute at least θ of the maximum contribution to the paper, and become co-authors. The stars denote scientists that contribute to papers A, B or C. We assume that the success

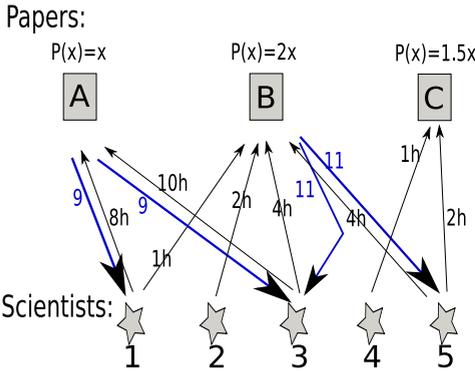


Figure 3.1: Scientists contribute time to papers (in black), and share the value of the accepted ones (in blue). Paper C does not reach the success threshold of 5 and therefore has no value to be shared.

threshold for a paper to be published is 5, and that the minimum contribution ratio to become a co-author of a paper is $\theta = 0.75$. Given this, paper A receives the total contribution of 18, yielding the value of 18 and thereby passing the success threshold, i.e. getting published. Its value is equally divided between both contributors, since both become co-authors, because $8 \geq \theta \cdot 10$. Paper B receives the contribution of 11, yielding the value of $2 \cdot 11 = 22$, which lies above the success threshold, causing B's publication. This value is equally divided between the two agents who contribute 4, since 1 and 2 are below $\theta \cdot 4$, and thus are not co-authors. As to paper C, it receives 3 and, therefore, yields $1.5 \cdot 3 = 4.5$, which is below the success threshold, so C is not published at all.

Let us apply the other definitions. In this example, B is the only steep project, thus $k = 1$. No project is vacant, since every paper receives some contribution. At paper A, scientists 2, 4 and 5 are dominated, but no-one is suppressed, because only the non-contributing to paper A scientists are dominated. At paper B, on the other hand, scientists 1, 2 and 4 are dominated, while scientists 1 and 2, who contribute there, are also suppressed.

3. THE QUOTA MODEL

In this section, we study the equilibria of shared effort games with a quota and their efficiency. We first give an example of an NE, and generalize it to a general theorem. Then, we use Theorem 4 on page 29 and Theorem 6 on page 30 from Chapter 2 as a basis for new existence and efficiency theorems for the quota model. We analyze the implications of these results on achieving efficiency at the end.

We need the following definition, generalizing Definition 2 on page 22 from Chapter 2 of a suppressed player.

Definition 10. A player is wasted at a project if it is suppressed there or if the project is out of the quota, but the player contributes there positively.

We demonstrate this on Example 2.

Example 2 (Continued). Assuming quota $q = 2$, no scientist is wasted at paper A, since no-one is suppressed there and A is in the quota. B is also in the quota, but scientists 1

and 2 are suppressed there, and thus wasted there. C is outside the quota, and therefore, its contributors, namely 4 and 5, are wasted there.

In an NE, a player is wasted at a project if and only if it is wasted at any project where it contributes. This is true since if a player is wasted at project p but it also contributes to project $q \neq p$ and is not wasted there, then it would like to move its contribution from p to project q .

First, we provide sufficient conditions for having an NE, which lets us estimate the efficiency. Then, we analyze, how the organizers can choose the quota to influence the behavior in NE. The intuition is that the additional condition enables more NE to exist, which increases the price of stability but decreases the price of anarchy. On the other hand, a profile that is an NE in the model from Section 2 of Chapter 2 can cease being so in our model, since some projects may obtain no value because of the quota.

Now, we show that having a quota can lead to counter-intuitive results. In the following example, there can be an NE where no steep project obtains a contribution. The idea is that any deviation from the project where everyone contributes is non-profitable, because it would still leave the other projects out of quota.

Example 3. Given projects 1 and 2, such that $\alpha_2 > \alpha_1$, assume that all the players contribute all their budgets to project 1. If $\alpha_2 B_n < \alpha_1 \sum_{i=1}^{n-1} B_i$ and $q = 1$, then no player can deviate to project 2, as this would still leave that project out of the quota, and therefore, this profile is an NE.

In this NE, the social welfare is equal to $\alpha_1 \sum_{i \in N} B_i$. The optimal social welfare, achieved if and only if all the players contribute all their budgets to project 2, is equal to $\alpha_2 \sum_{i \in N} B_i$. The ratio between the social welfare in this NE and the optimal one is $\frac{\alpha_1}{\alpha_2}$. That ratio is an upper bound on the price of anarchy of this game. In addition, since the optimal profile is also an NE, the price of stability is 1.

The price of anarchy is smaller than $\frac{\alpha_1}{\alpha_2}$ if and only if some agents do not contribute all their budgets. This can only happen in an NE if θ is positive, and if this is the case, then we can have arbitrarily low price of anarchy, down to the case when only agent n contributes, if $\theta B_n > B_{n-1}$, and then, $\text{PoA} = \frac{\alpha_1 B_n}{\alpha_2 \sum_{i \in N} B_i}$.

The ideas of this example yield the following theorem about possible Nash equilibria.

Theorem 9. Consider a θ -equal sharing game with $n \geq 2$ players with budgets $B_n \geq \dots \geq B_2 \geq B_1$ (the order is w.l.o.g.), $0 < \theta < 1$, linear project functions with coefficients $\alpha_m = \alpha_{m-1} = \dots = \alpha_{m-k+1} > \alpha_{m-k} \geq \alpha_{m-k-1} \geq \dots \geq \alpha_1$ (the order is w.l.o.g.), and quota q .

This game has a pure strategy NE, if $q = 1$ and $\alpha_m B_n < \alpha_1 \sum_{i=1}^{n-1} B_i$. In addition, $\text{PoA} \leq \frac{\alpha_1}{\alpha_m}$ and $\text{PoS} = 1$. The exact expression for the price of anarchy is $\frac{\alpha_1 \sum_{i: B_i \geq \theta B_n} B_i}{\alpha_m \sum_{i \in N} B_i}$.

Proof. If all the players contribute to a single project, then since $\alpha_m B_n < \alpha_1 \sum_{i=1}^{n-1} B_i$, no player can deviate to any project, because this would still leave this project out of the quota. Therefore, this profile is an NE.

In particular, when all the players invest all their budgets in project m , it is an NE, and thus, $\text{PoS} = 1$. When all the players invest in 1, it also is an NE, showing that $\text{PoA} \leq \frac{\alpha_1}{\alpha_m}$.

To find the exact expression for the price of anarchy, notice that the worst equilibrium for the social welfare is when everyone contributes to the least profitable project, and

only those who have a reason to do so contribute. Having an incentive means being not below the threshold amount, θB_n . This equilibrium yields $\alpha_1 \sum_{i: B_i \geq \theta B_n} B_i$. \square

This theorem, in accord with the intuition above, shows that reducing the quota can either facilitate an optimal NE, or a very inferior NE. Actually, every efficiency of the form $\frac{\alpha_i}{\alpha_m}$ is possible at equilibrium.

We now present an existence theorem, extending Theorem 4 on page 29 from Chapter 2 to the quota model. The theorem presents possible equilibria, providing advice on possible stable states. The theorem afterwards compares the efficiencies of these states.

Theorem 10. *Consider an equal θ -sharing game with $n \geq 2$ players with budgets $B_n \geq \dots \geq B_2 \geq B_1$ (the order is w.l.o.g.), $0 < \theta < 1$, linear project functions with coefficients $\alpha_m = \alpha_{m-1} = \dots = \alpha_{m-k+1} > \alpha_{m-k} \geq \alpha_{m-k-1} \geq \dots \geq \alpha_1$ (the order is w.l.o.g.), and quota q .*

This game has a pure strategy NE, if one of the following holds.²

1. $B_1 \geq k\theta B_n$ and all of the following hold.

(a) if $k \leq q$, then $\frac{1}{n} \alpha_{m-k+1} \geq \alpha_{m-k}$,

(b) if $k > q$, then $B_n < \sum_{j=1}^{n-1} B_j / q$;

2. $B_{n-1} < \frac{\theta}{|\Omega|} B_n$ and all the project functions are equal, i.e. $\alpha_m = \alpha_1$.

The proof provides an equilibrium profile and shows that no deviation is profitable.

Proof. To prove part 1, distinguish between the case where $k \leq q$ and $k > q$. If $k \leq q$, then the profile where all the players allocate $1/k$ th of their respective budgets to each of the steep projects is an NE for the same reasons that were given for the original model, since here, the quota's existence can only reduce the motivation to deviate.

If, on the other hand, $k > q$, consider the profile where all the players allocate $1/q$ th of their respective budgets to each of the q steep projects $m, m-1, \dots, m-q+1$. This is an NE, since the only deviation that is possibly profitable, besides reallocating between the non vacant projects, is a player moving all of her contributions from some projects to one or more of the vacant projects. This cannot bring profit, because these previously vacant projects will be outside of the quota, since $B_n < \sum_{j=1}^{n-1} B_j / q$. As for reallocating between the non-vacant projects, this is not profitable, since $B_1 \geq k\theta B_n$ means that suppressing is impossible. Therefore, this is an NE.

We now prove part 2. Let every player divide her budget equally among all the projects. No player wants to deviate, for the following reasons. All the projects obtain equal value, and therefore are in the quota. Player n suppresses all the rest and obtains her maximum possible profit, $\alpha_m (\sum_{i \in N} B_i)$. The rest obtain no profit, since they are suppressed whatever they do. \square

We now prove an efficiency result, similar to Theorem 6 on page 30 from Chapter 2. The result is based on the equilibria from Theorem 10.

²If α_{m-k} does not exist, consider the containing condition to be vacuously true.

Theorem 11. Consider an equal θ -sharing game with $n \geq 2$ players with budgets $B_n \geq \dots \geq B_2 \geq B_1$, $0 < \theta < 1$ (the order is w.l.o.g.), linear project functions with coefficients $\alpha_m = \alpha_{m-1} = \dots = \alpha_{m-k+1} > \alpha_{m-k} \geq \alpha_{m-k-1} \geq \dots \geq \alpha_1$ (the order is w.l.o.g.)³, and quota q .

1. Assume that $B_1 \geq k\theta B_n$ and all of the following hold.

- (a) if $k \leq q$, then $\frac{1}{n} \alpha_{m-k+1} \geq \alpha_{m-k}$,
- (b) if $k > q$, then $B_n < \sum_{j=1}^{n-1} B_j / q$;

Then, there exists a pure strategy NE and there holds: PoS = 1.

2. Assume that $B_{n-1} < \frac{\theta}{|\Omega|} B_n$ and all the project functions are equal, i.e. $\alpha_m = \alpha_1$.

Then, there exists a pure strategy NE and the following holds: PoS = 1, PoA = $\frac{B_n}{\sum_{i \in \{1, 2, \dots, n\}} B_i}$.

Proof. We first prove part 1. According to proof of part 1 of Theorem 10, equally dividing all the budgets among $\min\{k, q\}$ steep projects is an NE. Therefore, PoS = 1.

Part 2 is proven analogously to how it is proven for Theorem 6 from Chapter 2. We repeat the proof to be self-contained. For part 2, recall that in the proof of part 2 of Theorem 10, we show that everyone equally dividing the budgets between all the projects is an NE. This is optimal for the social welfare, and so PoS = 1. We turn to find the price of anarchy now. If player n acts as just mentioned, while the other players do not contribute anything, then this is an NE, since all the projects are equal and therefore, in the quota, and players $1, \dots, n-1$ will be suppressed at any contribution. An NE cannot have a lower social welfare, since n gets at least $\alpha_m B_n$ in any NE, since this is obtainable alone. Therefore, the fraction between the two social welfare values, namely $\frac{\alpha_m B_n}{\alpha_m \sum_{i \in \{1, 2, \dots, n\}} B_i}$, is the PoA. \square

For practical purposes, such as organizing elections, we would like to know how quota can be chosen to improve the behavior of the players (supporters, in the elections example). Theorem 9 implies that if no player has a budget as large as the total budget of all the other players times α_1 / α_m , then by choosing $q = 1$, many equilibria, including an optimal one, can be achieved. A possible problem with such an optimal equilibrium is that all the players will invest in the same project, which is unrealistic in some applications, such as conferences. It is still realistic in elections or funding large projects. From the proof of part 1 of Theorem 10 follows that applying a quota can force the agents to concentrate on less projects than they would concentrate on without the quota, if they follow the NE from the proof. Theorem 11 provides optimal Nash equilibria, though they may be not the only possible ones.

Additionally, the condition " $k > q \Rightarrow B_n < \sum_{j=1}^{n-1} B_j / q$ " in Theorem 11 implies, by taking the contrapositive, that if the largest budget can be much larger than the rest, then at least k projects have to be successful, if one wants our optimum NE to be guaranteed. When there are many equally glorious projects to contribute to, meaning that k is large, this constraint becomes non-trivial to implement. The condition

³If α_{m-k} does not exist, consider the containing condition to be vacuously true.

“ $k \leq q \Rightarrow \frac{1}{n}\alpha_{m-k+1} \geq \alpha_{m-k}$ ” in Theorem 11 implies, by looking at the contrapositive, that if the difference between the two largest projects is not big enough, then the quota has to be less than k if one wants our optimum NE to be guaranteed. This becomes non-trivial when the number of the most glorious (steep) projects is small.

4. THE SUCCESS THRESHOLD MODEL

In this section, we consider shared effort games with a success threshold, studying the existence and efficiency of their Nash equilibria. We assume that the success threshold δ is always at most the sum of all the budgets times α_m , to allow for at least one project to be accepted, in at least one strategy profile. Similarly to the previous section, we begin by giving an example, which inspires a theorem, and then we provide existence and efficiency theorems, based on such theorems from Chapter 2.

We call a project that has a value of at least the threshold in a given profile an *accepted* project, and we call it *unaccepted* otherwise. We shall need the following definition, generalizing Definition 2 on page 22 from Chapter 2 of a suppressed player.

Definition 11. *A player is blank at a project if it is suppressed there or if the project is unaccepted, but the player contributes there positively.*

Let us refer to Example 2 again.

Example 2 (Continued). *The accepted papers are A and B. Therefore, the blank scientists at each paper are those we describe as wasted after Definition 10.*

In an NE, a player is blank at a project if and only if it is blank at any project where it contributes. This is true since if a player is blank at project p but it also contributes to project $q \neq p$ and is not blank there, then it would like to move its contribution from p to project q .

The goal is to estimate the efficiency of the possible equilibria. Then, we analyze how organizers can choose the success threshold to influence the behavior in NE. The main intuitive difference between a quota and a success threshold is that a quota is about relative project values, while a success threshold is absolute.

Similarly to what we saw with quota, success threshold can cause counter-intuitive results. In the following example, there can be an NE where no steep project obtains a contribution.

Example 4. *Given the projects 1 and 2, such that $\alpha_2 > \alpha_1$, assume that all the players contribute all their budgets to project 1. If $\delta > \alpha_2 B_n$, then no player can deviate to project 2, as this would leave that project unaccepted, and therefore, this profile is an NE.*

The conclusions about the prices of anarchy and stability are the same as in Example 3.

The ideas of this example yield the following theorem.

Theorem 12. *Consider an equal θ -sharing game with $n \geq 2$ players with budgets $B_n \geq \dots \geq B_2 \geq B_1$ (the order is w.l.o.g.), $0 < \theta < 1$, linear project functions with coefficients $\alpha_m = \alpha_{m-1} = \dots = \alpha_{m-k+1} > \alpha_{m-k} \geq \alpha_{m-k-1} \geq \dots \geq \alpha_1$ (the order is w.l.o.g.), and success threshold δ .*

This game has a pure strategy NE, if $\alpha_m B_n < \delta$. In addition, $\text{PoA} \leq \frac{\alpha_1}{\alpha_m}$ and $\text{PoS} = 1$. If $\alpha_1 \sum_{i=1}^n B_i < \delta$, then $\text{PoA} = 0$.

Proof. If all the players contribute to a single project, then since $\alpha_m B_n < \delta$, no player can deviate to any project, because this would still leave that project unaccepted. Therefore, this profile is an NE.

In particular, when all the players invest all their budgets in project m , it is an NE, and thus, $\text{PoS} = 1$. When all the players invest in 1, it also is an NE, showing that $\text{PoA} \leq \frac{\alpha_1}{\alpha_m}$, and if $\alpha_1 \sum_{i=1}^n B_i < \delta$, then $\text{PoA} = 0$. \square

This theorem, in accord with the intuition above, shows that increasing the success threshold can either facilitate an optimal NE, or an inferior NE. Actually, every efficiency of the form $\frac{\alpha_j}{\alpha_m}$, for $j \geq \min \{i : \alpha_i \sum_{l=1}^n B_l \geq \delta\}$, is possible at an equilibrium.

Next, we provide sufficient conditions for the existence of an NE, extending Theorem 4 on page 29 from Chapter 2 to the success threshold model as follows.

Theorem 13. Consider an equal θ -sharing game with $n \geq 2$ players with budgets $B_n \geq \dots \geq B_2 \geq B_1$ (the order is w.l.o.g.), $0 < \theta < 1$, linear project functions with coefficients $\alpha_m = \alpha_{m-1} = \dots = \alpha_{m-k+1} > \alpha_{m-k} \geq \alpha_{m-k-1} \geq \dots \geq \alpha_1$ (the order is w.l.o.g.), and success threshold δ .

This game has a pure strategy NE, if one of the following holds.⁴

Define $p \triangleq \left\lfloor \frac{\alpha_m \sum_{i \in N} B_i}{\delta} \right\rfloor$; intuitively, it is the number of the projects that can be accepted.

1. $B_1 \geq k\theta B_n$ and all of the following hold.

(a) If $k \leq p$, then $\frac{1}{n} \alpha_{m-k+1} \geq \alpha_{m-k}$,

(b) If $k > p \geq 1$, then $\alpha_m B_n < \delta$;

2. $B_{n-1} < \frac{\theta}{|\Omega|} B_n$, all the project functions are equal, i.e. $\alpha_m = \alpha_1$, and $\theta \leq \alpha_m$.

Proof. To prove part 1, we distinguish between the case where $k \leq p$ and $k > p$. If $k \leq p$, then the profile where all the players allocate $1/k$ th of their respective budgets to each of the steep projects is an NE for the same reasons that were given for the original model, since here, the requirement to be not less than the success threshold can only reduce the motivation to deviate.

If, on the other hand, $k > p$, consider the profile where all the players allocate $1/p$ th of their respective budgets to each of the p steep projects $m, m-1, \dots, m-p+1$. This is an NE, since the only deviation that is possibly profitable, besides moving budgets between the non vacant projects, is a player moving all of her contributions from some projects to one or more of the vacant projects. This cannot bring profit, because these previously vacant projects will be unaccepted, since $\alpha_m B_n < \delta$. Additionally, any reallocating between the non-vacant projects is not profitable, since $B_1 \geq k\theta B_2$ means that suppressing is impossible. Therefore, the current profile is an NE.

We now prove part 2. The proof distinguishes between the case where the condition $p \geq |\Omega|$ holds or not. If $p \geq |\Omega|$, then the proof continues as in the case of part 2 of Theorem 10, where every player divides her budget equally among all the projects. All the projects are accepted, so no new deviations become profitable.

⁴If α_{m-k} does not exist, consider the containing condition to be vacuously true.

In the case that $p < |\Omega|$, consider the profile where all the players allocate $1/p$ th of their respective budgets to each of the p projects $m, m-1, \dots, m-p+1$. This is an NE, since the only deviation that is possibly profitable is some player $j < n$ moving all her budget to a vacant project. However, this is not profitable, since the project would be unaccepted, because $B_j \leq B_{n-1} < \frac{\theta}{|\Omega|} B_n < \theta\delta/\alpha_m \leq \delta$. The penultimate inequality stems from $p < |\Omega| \iff \frac{\alpha_m \sum_{i \in N} B_i}{|\Omega|} < \delta$, and the final one results from the assumption that $\theta \leq \alpha_m$. Therefore, this is an NE. \square

We now prove an efficiency result, inspired by Theorem 6 on page 30 from Chapter 2.

Theorem 14. Consider an equal θ -sharing game with $n \geq 2$ players with budgets $B_n \geq \dots \geq B_2 \geq B_1$, $0 < \theta < 1$ (the order is w.l.o.g.), linear project functions with coefficients $\alpha_m = \alpha_{m-1} = \dots = \alpha_{m-k+1} > \alpha_{m-k} \geq \alpha_{m-k-1} \geq \dots \geq \alpha_1$ (the order is w.l.o.g.)⁵, and success threshold δ .

Define $p \triangleq \left\lfloor \frac{\alpha_m \sum_{i \in N} B_i}{\delta} \right\rfloor$, as in Theorem 13.

1. Assume that $B_1 \geq k\theta B_n$ and all of the following hold.

- (a) If $k \leq p$, then $\frac{1}{n} \alpha_{m-k+1} \geq \alpha_{m-k}$,
- (b) If $k > p \geq 1$, then $\alpha_m B_n < \delta$;

Then, there exists a pure strategy NE and there holds: PoS = 1.

2. Assume $B_{n-1} < \frac{\theta}{|\Omega|} B_n$, all the project functions are equal, i.e. $\alpha_m = \alpha_1$, and $\theta \leq \alpha_m$.

Then, there exists a pure strategy NE and PoS = 1, If, an addition, $\alpha_m B_n \geq \delta$, then $\text{PoA} = \frac{B_n}{\sum_{i \in \{1, 2, \dots, n\}} B_i}$.

Proof. We first prove part 1. According to proof of part 1 in Theorem 13, equally dividing all the budgets among $\min\{k, p\}$ steep projects is an NE. Therefore, PoS = 1.

Part 2 is proven as follows. Since all the players dividing their budgets equally between any $\min\{p, m\}$ projects constitutes an NE, we have PoS = 1.

To discuss of the price of anarchy, we define the number of projects that player n can make accepted on her own, $r \triangleq \left\lfloor \alpha_m \frac{B_n}{\delta} \right\rfloor$, and distinguish between the case where $m \leq r$ and $m > r$. If $m \leq r$, consider the profile where player n divides her budget equally between all the projects, while the other players contribute nothing at all. This is an NE, because all the projects are accepted, player n cannot increase her profit and any other player will be suppressed, if she contributes anything anywhere. On the other hand, if $m > r$, consider the profile where player n divides her budget equally between $m, m-1, \dots, m-r+1$, while the other players contribute nothing at all. This is an NE, since the only possible deviation is player $j < n$ contributing to a vacant project. However, we have $B_j \leq B_{n-1} < \frac{\theta}{|\Omega|} B_n < \theta\delta/\alpha_m \leq \delta$. This means that the project would be unaccepted. Therefore, this is an NE.

Therefore, $\text{PoA} \leq \frac{\alpha_m B_n}{\alpha_m (\sum_{i \in N} B_i)}$. Since $\alpha_m B_n \geq \delta$, at any NE, player n receives at least $\alpha_m B_n$, and therefore, $\text{PoA} = \frac{B_n}{\sum_{i \in \{1, 2, \dots, n\}} B_i}$. \square

⁵If α_{m-k} does not exist, consider the containing condition to be vacuously true.

For practical purposes, such as organizing elections or a conference, we would like to know how to choose the success threshold to improve the behavior of the players. Theorem 12 implies that by choosing a success threshold that disables any player to make a project successful on her own, the optimum social welfare can be achieved in equilibrium. The first problem of this approach is that it also allows very inefficient profiles constitute equilibria, and the second problem is that the discussed equilibria have all the players investing in the same project, which is unreasonable to applications like conferences, though possible in other applications, such as sponsorship of elections. From the proof of both parts of Theorem 13 follows that applying a success threshold can force the agents to concentrate on less projects than without the threshold, if they follow the NE from the proof. Theorem 14 provides optimal NE, though they are not the only ones, so regulation may be needed to ensure that those equilibria are indeed chosen.

In addition, the condition “If $k \leq p \Rightarrow \frac{1}{n} \alpha_{m-k+1} \geq \alpha_{m-k}$ ” in Theorem 14 implies, by its contrapositive, that if the second best project is close to a best one, then the threshold should be big enough, for at least our optimum NE to be guaranteed. The contrapositive of the condition “ $k > p \Rightarrow \alpha_m B_n < \delta$ ” implies that if the biggest player is able to make a project succeed on her own, then the threshold should be small enough so that p is at least the number of the most profitable projects, for at least our optimum NE to be guaranteed.

5. CONCLUSIONS AND FURTHER RESEARCH

This chapter analyzes the possible stable investments in projects, where a project has to comply to certain requirements to obtain its value. The goal is to advise which investments are individually and socially preferable. In order to model common resource allocation to competing projects, such as paper co-authorship and investment in firms, we model agents contributing to several projects. Each agent has a budget of effort, which she freely divides between the projects. A project that succeeds in the competition obtains a value, which is divided between the contributors who have contributed at least a given fraction of the maximum contribution to the project. Although in practice the exact thresholds are usually not known, the results we obtain are general and therefore meaningful in such unknown settings.

We model succeeding in a competition either by a quota of projects that actually obtain their value, or by a success threshold on the value of projects that do. Comparing these models, we see from Theorems 9 and 12 that the success threshold allows ensuring that there exists a socially optimal equilibrium while the quota requires also assuming that the largest effort budget is less than the sum of the other ones times α_1/α_m . In addition, comparing Theorems 11 and 14 shows that provided the smallest budget is at least a certain fraction of the largest one, the following holds: Large enough a threshold guarantees that an optimal profile will be an equilibrium, while choosing small enough a quota guarantees the existence of an optimum equilibrium provided that the largest budget is less than the sum of the other ones. Unlike in the described cases, where success threshold seems stronger than quota, we notice that the second part of Theorem 14 actually contains an additional condition, relatively to the second part of Theorem 11, but since the second parts refer to the case of a single agent being able to dominate everyone everywhere and all the projects being equally rewarding, this is less practical

anyway. To conclude the comparison, choosing success threshold has more power, since choosing quota needs to assume an additional relation between the budgets, in order to guarantee that socially optimal Nash equilibria exist. Intuitively, this stems from a quota needing an assumption on what the players are able to do to increase their utility, given the quota, while providing a success threshold can be done already with the budgets in mind.

Interestingly, both a quota and a success threshold create equilibria where the agents contribute to less projects than without any of these conditions. This is a concentrating effect.

We conclude that both models allow to guarantee that an optimum profile is an NE, though no guarantee is provided as to inefficient profiles being an NE as well. Therefore, some coordination may be required to actually achieve an efficient profile.

There are many interesting directions to expand this research. First, real papers, books, and many other common projects have an upper bound on the maximal number of participants. Analogously, a person has an upper bound on the maximal number of projects she can contribute to. The model should account for these bounds. Another point is that competition can be of many sorts. For instance, a project may need to have a winning coalition of contributors, in the sense of cooperative games. The fate of the projects that fail the competition can also vary; for example, their value can be distributed between the winning projects. We have managed to extend the sufficiency results for existence from Section 3.1 of Chapter 2, but the necessity seems hard for analytical analysis. Other analytical approaches or simulations may be tried to delineate the set of Nash equilibria more clearly. Naturally, project functions do not have to be linear, so there is a clear need to model various non-linear functions. Such a more general model will make the conclusions on scientific investments, paper co-authorship, elections, and the many other application domains more precise, and enable us to further improve the advice to participants as well as organizers.

Having modeled and analyzed projects that yield a revenue to be shared, we turn next to reciprocal interactions. This kind of projects is quite complex, so we will study a single reciprocation in the following chapter, as a preparation for considering sharing effort between such projects.

REFERENCES

- [1] P. A. Bernstein, D. DeWitt, A. Heuer, Z. Ives, C. S. Jensen, H. Meyer, M. T. Özsu, R. T. Snodgrass, K.-Y. Whang, and J. Widom, *Database publication practices*, in *Proceedings of the 31st International Conference on Very Large Data Bases*, VLDB '05 (VLDB Endowment, 2005) pp. 1241–1245.
- [2] J. R. Douceur, *Paper rating vs. paper ranking*, *SIGOPS Oper. Syst. Rev.* **43**, 117 (2009).
- [3] J. Kleinberg and S. Oren, *Mechanisms for (mis)allocating scientific credit*, in *Proceedings of the 43rd annual ACM symposium on Theory of computing* (ACM, New York, NY, USA, 2011) pp. 529–538.
- [4] R. Ganti, F. Ye, and H. Lei, *Mobile crowdsensing: current state and future challenges*, *IEEE Communications Magazine* **49**, 32 (2011).

- [5] F. M. Harper, S. X. Li, Y. Chen, and J. A. Konstan, *Social comparisons to motivate contributions to an online community*, in *Proceedings of the 2nd international conference on Persuasive technology* (Springer-Verlag, Berlin, Heidelberg, 2007) pp. 148–159.
- [6] R. Siegel, *All-pay contests*, *Econometrica* **77**, 71 (2009).
- [7] Y. Bachrach, V. Syrgkanis, and M. Vojnović, *Efficiency and the Redistribution of Welfare*, Tech. Rep. (Microsoft Research, 2011).
- [8] E. Koutsoupias and C. Papadimitriou, *Worst-case equilibria*, in *16th Annual Symposium on Theoretical Aspects of Computer Science* (Trier, Germany, 1999) pp. 404–413.
- [9] C. Papadimitriou, *Algorithms, Games, and the Internet*, in *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing*, STOC '01 (ACM, New York, NY, USA, 2001) pp. 749–753.
- [10] A. S. Schulz and N. S. Moses, *On the performance of user equilibria in traffic networks*, in *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '03 (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2003) pp. 86–87.
- [11] E. Anshelevich, A. DasGupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden, *The price of stability for network design with fair cost allocation*, in *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on* (2004) pp. 295–304.
- [12] M. R. Baye, D. Kovenock, and C. G. de Vries, *The all-pay auction with complete information*, *Economic Theory* **8**, 291 (1996).
- [13] B. Roberson, *The Colonel Blotto game*, *Economic Theory* **29**, 1 (2006).
- [14] E. Anshelevich and M. Hoefer, *Contribution games in networks*, *Algorithmica* **63**, 51 (2011).
- [15] G. K. Batchelor, *Preoccupations of a journal editor*, *Journal of Fluid Mechanics* **106**, 1 (1981).
- [16] L. S. Shapley, *A value for n -person games. Contribution to the Theory of Games*, *Annals of Mathematics Studies* **2**, 28 (1953).

4

TOWARDS DECISION SUPPORT IN RECIPROICATION

Life cannot subsist in society but by reciprocal concessions.

Samuel Johnson, 1774

People often interact repeatedly: with relatives, through file sharing, in politics, etc. Many such interactions are reciprocal: reacting to the actions of the other. In order to facilitate decisions regarding reciprocal interactions, we analyze the development of reciprocation over time. To this end, we propose a model for such interactions that is simple enough to enable formal analysis, but is sufficient to predict how such interactions will evolve. Inspired by existing models of international interactions and arguments between spouses, we suggest a model with two reciprocating attitudes where an agent's action is a weighted combination of the others' last actions (reacting) and either i) her innate kindness, or ii) her own last action (inertia). We analyze a network of repeatedly interacting agents, each having one of these attitudes, and prove that their actions converge to specific limits. Convergence means that the interaction stabilizes, and the limits indicate the behavior after the stabilization. For two agents, we describe the interaction process and find the limit values. For a general connected network, we find these limit values if all the agents employ the second attitude, and show that the agents' actions then all become equal. In the other cases, we study the limit values using simulations. We discuss how these results predict the development of the interaction and can be used to help agents decide on their behavior.

An extended abstract is published at [1] as "The Convergence of Reciprocation". Most parts of this chapter appear at <http://arxiv.org/abs/1601.07965>.

1. INTRODUCTION

In the previous chapters we analyzed participating in projects that yield values. Now, we consider another kind of projects, namely, *interaction*. Interaction is central in human behavior, e.g., at school, in file sharing, in business cooperation and political struggle. We aim at facilitating decision support for the interacting agents (how to act) and for the outside observers (how to influence the acting agents). To this end, we want to predict interaction. In this and the next chapter we study a single interaction project in detail, and in Chapter 6, we model agents dividing their efforts between multiple interaction projects.

Instead of being economically rational, people tend to adopt ways of behavior [2, 3] not necessarily maximizing some utility function. Furthermore, people tend to *reciprocate*, i.e., react on the past actions of others [4–7]. Reciprocation is also important for engineering computer systems that represent people or interact with people, such as service robots owned by people or software agents using clouds. Since reciprocation is ubiquitous, predicting such behavior will allow predicting many real-life interactions and advising on how to improve them. Therefore, we need a model for agents reciprocating for a long time with certain reciprocal habits that is amenable to theoretical analysis and precise enough to predict such interactions. Understanding such a model would also help understanding how to improve personal and public value of such interactions.

A broad overview of the existing models appears in Section 1.1.2. Here we take a quicker look, and conclude that none of the existing work fits our goals. The literature has two major kinds of models of (sometimes repeated) reciprocation: models that explain the existence of reciprocation and those analyzing its consequences. We begin by looking at the first kind of models. There are several kinds of models that explain why reciprocation has come to being. First, there are the models explaining how reciprocation could have directly evolved, such as the famous research of Axelrod [8, 9], showing that reciprocating is expedient to self-interested agents, or papers that take into account helping genetic relatives as well [10, 11]. Other works advocate less evident ways of evolution of reciprocation. For example, Van Segbroeck et al. [12] argue that fairness motivates reciprocation, and investigate the evolution of fairness. Trivers [13] demonstrates how emotions sustaining reciprocation, like moralistic aggression and guilt, have evolved. Reacting to past actions without any hope to gain from it, called *strong reciprocity*, is modeled and analyzed by Gintis [14, Chapter 11]. Some explanation for it appears in Fehr, Fischbacher and Gächter [15]. Segal and Sobel [16] assume that agents care not only about the outcomes but also about the strategies, and provide conditions for the utility being represented by a unique linear combination of the outcome dependent utilities of the agents.

On another research avenue, given that reciprocal tendencies exist, the following articles analyze what ways it makes interactions develop. Falk and Fischbacher, Rabin's and others [4, 17–19] model and analyze games where the utility function of rational agents positively depends on showing reciprocation. The role of institutions in how well reciprocation pays off is studied in [5].

None of the above research directions satisfies our need to model a lengthy reciprocal process with actions from a continuous domain, given that reciprocation takes place. Reciprocity is seen as an inborn quality [13, 15], which has probably been evolved from

rationality of agents, as was shown by Axelrod [8]. In addition, predicting reciprocal processes would be in the spirit of the call to consider various repercussions of reciprocity from [20]. For lack of analysis of non-discrete lengthy interactions, caused by inborn reciprocation, we model and study how reciprocity makes interaction evolve with time.

Since we are interested in the extent of actions, we represent actions by *weight*, where a bigger value means a more desirable contribution to its recipient or, in the interpersonal context, investment in the relationship. In our model, agents reciprocate both to the agent they are acting on and to their whole neighborhood. We model reciprocity by two reciprocation attitudes, an action's weight being a convex combination of i) one's own kindness or ii) one's own last action, and the other's and neighborhood's last actions on the acting agent. A convex combination, having its weights nonnegative and sum up to 1, represents a whole being assembled from fractions. When determining an action, the whole past should be considered, but to facilitate analysis, we assume that the last actions represent the history well enough. The motivation for defining an action (or how much it changes) or a state by a linear combination of the other side's actions and own actions and qualities comes from similar models of arms race [21, 22] and spouses' interaction [23] (piecewise linear in this case). Attitude i) depending on the (fixed) kindness is called *fixed*, and ii) depending on one's own last action is called *floating*. Given this model, we study its behavioral repercussions.

We now demonstrate reciprocal interaction in daily life.

Example 5. Consider $n = 4$ colleagues 1, 2, 3, 4, who can help or harm each other. Let the possible actions be: giving bad work, showing much contempt, showing little contempt, supporting emotionally a little, supporting emotionally a lot, advising, and let their respective weight be a point in $[-1, -0.5]$, $[-0.5, -0.2]$, $[-0.2, 0]$, $(0, 0.4]$, $[0.4, 0.7]$, $[0.7, 1]$. Assume that each person knows what the other did to him last time. The social climate, meaning what the whole group did, also influences behavior. However, we may just concentrate on a single pair of even-tempered colleagues who reciprocate regardless of the others.

Our major contributions are proving when the reciprocation process converges and finding the limit of convergence. A limit when time approaches infinity describes what actions will take place once they have stabilized. These results predict reciprocation and explain the above mentioned phenomena. In particular, exponential convergence means a rapid stabilizing, and it explains acquiring personal behavioral styles, which is often seen in practice [24]. We prove that when at most one agent is *fixed*, the limits of the actions of all agents are the same; this explains formation of organizational subcultures, known in the literature [25]. We also find that only the kindness values of the *fixed* agents influence the limits of the various actions, thereby explaining that persistence (i.e., being faithful to one's inner inclination) makes interaction go one's own way. The fact that persistence allows determining the extent of interaction is known in daily life; for instance, the recommendations to reject undesired requests by firmly repeating the reasons for rejection [26, Chapter 1] and [27, Chapter 8] basically recommend to persist.

We first consider two agents in Section 3 and Section 4, assuming their interaction is independent of other agents, or that the total influence of the others on the pair is negligible. The main convergence theorems we use when analyzing strategic choices of reciprocation habits in Chapter 5 are Theorems 15, 16 and Corollary 3. In the synchronous

case, the limits for two agents will also follow from a general convergence result that is presented later, unless both agents are *fixed*. We still present them with the other results for two agents for the completeness of Section 3. We study interaction of many agents, where the techniques we used for two agents are not applicable, in Section 5. In the central convergence Theorem 20, we find the limit when all the agents act synchronously and at most one has the *fixed* reciprocation attitude. Section 6 simulates cases that we have not analyzed analytically.

2. MODELING RECIPROCATION

2.1. BASICS

We begin by modeling agents, actions and times, providing the notation necessary for the following definition of reciprocation, such as the inner parameters of an agent. We conclude the section by sharpening the model and providing explanatory examples.

Let $N = \{1, 2, \dots, n\}$ be $n \geq 2$ interacting agents. There is a given undirected interaction graph $G = (N, E)$, such that agent i may act on j and vice versa if and only if $(i, j) \in E$. The actions change, but the actors and those acted upon remain the same. Denote the degree of agent $i \in N$ in G by $d(i)$. This allows for various topologies, including heterogeneous ones.¹ We need to distinguish between the actions of agent i on agent j and the other way around. To be able to mention directed edges, we shall treat this graph as a symmetric directed one, meaning that for every $(i, j) \in E$, we have $(j, i) \in E$. Time is modeled by a set of discrete moments $t \in T \triangleq \{0, 1, 2, \dots\}$, defining a time slot whenever at least one agent acts. Agent i acts at times $T_i \triangleq \{t_{i,0} = 0, t_{i,1}, t_{i,2}, \dots\} \subseteq T$, and $\cup_{i \in N} T_i = T$. We assume that all agents act at $t = 0$, since otherwise we could not sometimes consider the last action of another agent, which would force us to complicate the model and render it even harder for theoretical analysis. If all agents always act at the same times ($T_1 = T_2 = \dots = T_n = T$), we say they act *synchronously*. Throughout the whole dissertation, relatively to the synchronous case, analyzing the general model does not add much game theoretic interest. The general, not necessarily synchronous, model is practically important, on the expense of adding quite some complexity, so we do analyze it. However, a purely theoretically inclined reader may safely assume synchronicity, skipping the excessive discussion necessary only for the general case, such as, for instance, Lemma 5 and Section 4.

For the sake of asymptotic analysis, we assume that each agent gets to act an infinite number of times; that is, T_i is infinite for every $i \in N$. Any real application will, of course, realize only a finite part of it, and infinity models the unboundedness of the process in time.

When (i, j) is in E , we denote the weight of an action by agent $i \in N$ on another agent $j \in N$ at moment $t \in T_i$ by $x_{i,j}: T_i \rightarrow \mathbb{R}$. To extend $x_{i,j}$ to the whole T , we first define *the last action time* $s_i(t): T \rightarrow T_i$ of agent i as the largest $t' \in T_i$ that is at most t . Since $0 \in T_i$, this is well defined. Now, we extend $x_{i,j}$ to T by $x_{i,j}(t) \triangleq x_{i,j}(s_i(t))$, and we have defined $x_{i,j}(t): T \rightarrow \mathbb{R}$. So, $x_{i,j}$ is the last action of agent i on (another) agent j . For example, when interacting by file sharing, the actions of sending a valid piece of a file, nothing,

¹I.e., with various degrees.

or a piece with a virus are decreasing in weight. Since only the weight of an action is relevant, we usually write “action” while referring to its weight.

We denote the total received contribution from all the neighbors $N(i)$ at their last action times not later than t by $\text{got}_i(t): T \rightarrow \mathbb{R}$; formally, $\text{got}_i(t) \triangleq \sum_{j \in N(i)} x_{j,i}(t)$.

2.2. RECIPROCATION

We now define two reciprocation attitudes, which define how an agent reciprocates. In order to model the inner inclination of an agent to act and the readiness to react, we need the following notions. The kindness of agent i is denoted by $k_i \in \mathbb{R}$; w.l.o.g., $k_n \geq \dots \geq k_2 \geq k_1$ throughout the thesis. Kindness models inherent inclination to help others; in particular, it determines the first action of an agent, before others have acted. We model agent i 's inclination to mimic a neighboring agent's action and the actions of the whole neighborhood in G by reciprocation coefficients $r_i \in [0, 1]$ and $r'_i \in [0, 1]$ respectively, such that $r_i + r'_i \leq 1$. Here, r_i is the fraction of $x_{i,j}(t)$ that is determined by the last action of j upon i , and r'_i is the fraction that is determined by $\frac{1}{|N(i)|}$ th of the total contribution to i from all the neighbors at the last time. Conceptually, reacting to last actions, one reacts to the actor, since “who you are is what you do” [28].

Intuitively, with the *fixed* attitude, actions depend on the agent's kindness at every time, while the *floating* attitude is loose, moving freely in the reciprocation process, and kindness directly influences such behavior only at $t = 0$. In both cases $x_{i,j}(0) \triangleq k_i$.

Definition 12. For the fixed reciprocation attitude, agent i 's action on another agent j is determined by the other agent's action weighted by r_i , by the total action of the neighbors weighted by r'_i and divided by the number of the neighbors, and by the agent's kindness weighted by $1 - r_i - r'_i$. That is, for $t \in T_i$,

$$x_{i,j}(t) \triangleq (1 - r_i - r'_i) \cdot k_i + r_i \cdot x_{j,i}(t-1) + r'_i \cdot \frac{\text{got}_i(t-1)}{|N(i)|}.$$

Definition 13. In the floating reciprocation attitude, agent i 's action is a weighted average of that of the other agent j , of the total action of the neighbors divided by the number of the neighbors, and of her own last action. To be precise, for $t \in T_i$,

$$x_{i,j}(t) \triangleq (1 - r_i - r'_i) \cdot x_{i,j}(t-1) + r_i \cdot x_{j,i}(t-1) + r'_i \cdot \frac{\text{got}_i(t-1)}{|N(i)|}.$$

The relations are (usually inhomogeneous) linear recurrences with constant coefficients, but many variables. We could express the dependence of $x_{i,j}(t)$ only on $x_{i,j}(t')$ with $t' < t$, but then the coefficients would not be constant, besides the case of two *fixed* agents. We are not aware of a method to use the general recurrence theory to improve our results.

The notation is summarized in Table 4.1.

2.3. CONTEXT AND EXAMPLES

We emphasize that compared to the other reciprocation models, our model takes reciprocal actions as given and looks at the process, while other models either consider how

Term:	Meaning:
T_i	The time moments when agent i acts.
Synchronous	$T_1 = T_2 = \dots = T_n$.
$s_i(t): T \rightarrow T_i$	$\max\{t' \in T_i \mid t' \leq t\}$.
$x_{i,j}(t): T \rightarrow \mathbb{R}$	The action of i on another agent j at time $s_i(t)$.
$\text{got}_i(t): T \rightarrow \mathbb{R}$	$\sum_{j \in N(i)} x_{j,i}(t)$.
k_i	The kindness of agent i .
$r_i, r'_i \in [0, 1], r_i + r'_i \leq 1$	The reciprocation coefficients of agent i .
Agent i has the <i>fixed</i> reciprocation attitude, j is another agent	At moment $t \in T_i$, $x_{i,j}(t) \triangleq \begin{cases} (1 - r_i - r'_i) \cdot k_i + r_i \cdot x_{j,i}(t-1) & t > t_{i,0} \\ + r'_i \cdot \frac{\text{got}_i(t-1)}{ N(i) } & \\ k_i & t = t_{i,0} = 0. \end{cases}$
Agent i has the <i>floating</i> reciprocation attitude, j is another agent	At moment $t \in T_i$, $x_{i,j}(t) \triangleq \begin{cases} (1 - r_i - r'_i) \cdot x_{i,j}(t-1) & t > t_{i,0} \\ + r_i \cdot x_{j,i}(t-1) + r'_i \cdot \frac{\text{got}_i(t-1)}{ N(i) } & \\ k_i & t = t_{i,0} = 0. \end{cases}$

Table 4.1: The notation used throughout the thesis.

reciprocation originates, such as the evolutionary model of Axelrod [8], or take it as given and consider specific games, such as in [4, 17–19].

There are several reminiscent but different models. The *floating* model resembles opinions that converge to a consensus [29–32], while the *fixed* model resembles converging to a general equilibrium of opinions [33]. Of course, unlike the models of spreading opinions, we consider different actions on various neighbors, determined by direct reaction and a reaction to the whole neighborhood. Still, because of some technical reminiscence to our models, we do use [29, Theorem 2] to prove Proposition 6. Another similar model is that of monotonic concession in negotiation [34] and that of bargaining over dividing a pie between two agents [35]. The main difference is that in those models, the agents decide what to do, while in our case, they follow the reciprocation formula. This reciprocal reaction fits the large body of literature on reciprocation, and we concentrate on what will occur, if the agents decide to concede reciprocally.

We now illustrate the model on Example 5.

Example 5 (Continued). *Let (just here) $n = 3$ and the reciprocation coefficients be $r_1 = r_2 = 0.5, r'_1 = r'_2 = 0.3, r_3 = 0.8, r'_3 = 0.1$. Assume the kindness to be $k_1 = 0, k_2 = 0.5$ and $k_3 = 1$. Since this is a small group, all the colleagues may interact, so the graph is a clique.² At $t = 0$, every agent's action on every other agent is equal to her kindness value, so agent*

²A clique is a fully connected graph.

1 does nothing, agent 2 supports emotionally a lot, and 3 provides advice. If all agents act synchronously, meaning $T_1 = T_2 = T_3 = \{0, 1, \dots\}$, and all agents get carried away by the process, meaning that they forget the kindness in the sense of employing floating reciprocation, then, at $t = 1$ they act as follows: $x_{1,2}(1) = (1 - 0.5 - 0.3) \cdot 0 + 0.5 \cdot 0.5 + 0.3 \cdot \frac{0.5+1}{2} = 0.475$ (supports emotionally a lot), $x_{1,3}(1) = (1 - 0.5 - 0.3) \cdot 0 + 0.5 \cdot 1 + 0.3 \cdot \frac{0.5+1}{2} = 0.975$ (provides advice), $x_{2,1}(1) = (1 - 0.5 - 0.3) \cdot 0.5 + 0.5 \cdot 0 + 0.3 \cdot \frac{0+1}{2} = 0.25$ (supports emotionally a little), and so on. Please note that now, unlike at $t = 0$, agent 1 acts on different agents differently.

Consider modeling tit for tat [9]:

Example 6. In our model, a tit for tat agent with two options: cooperate or defect is easily modeled with $r_i = 1$, $k_i = 1$, meaning that the original action is cooperating (1) and the next action is the current action of the other agent. If one of two tit-for-tat agents makes a mistake and begins with defection ($k_2 = 0$), acting synchronously, then they will alternate.

If the agents are human, this example predicts an indefinitely long alternation, which seems unrealistic to us. Similarly, an agent that sticks to his actions regardless the other seems highly implausible. This provides evidence that extreme values of the reciprocation coefficients are uncommon in life.

3. PAIRWISE INTERACTION

We now consider an interaction of two agents, 1 and 2, since this assumption allows proving more than we can in the general case. We assume that the pair has no other neighbors. This allows to also assume that $r'_1 = r'_2 = 0$, since now, r_i and r'_i play the same role.

To predict behavior, this section studies reciprocation mainly by proving convergence and finding its limit, which represents the actions after stabilizing. This section considers three sets of reciprocation attitudes: both agents are *fixed*, both are *floating*, and one is *fixed* while the other one is *floating*. For each such setting, we study convergence, properties of the sequence of an agent's actions and the relationship between the actions of the two agents.

If T_1 contains exactly all the even numbered slots and T_2 contains zero and all the odd ones, we say that the agents *alternate*. Since agent 1 can only act on agent 2 and vice versa, we write $x(t)$ for $x_{1,2}(t)$ and $y(t)$ for $x_{2,1}(t)$.

To formally discuss the actions after the interaction has stabilized, we consider the limits (if exist)³ $\lim_{p \rightarrow \infty} x(t_{1,p})$, and $\lim_{t \rightarrow \infty} x(t)$, for agent 1, and $\lim_{p \rightarrow \infty} y(t_{2,p})$ and $\lim_{t \rightarrow \infty} y(t)$ for agent 2. Since the sequence $\{x(t)\}$ is $\{x(t_{1,p})\}$ with finite repetitions, the limit $\lim_{t \rightarrow \infty} x(t)$ exists if and only if $\lim_{p \rightarrow \infty} x(t_{1,p})$ does. If they exist, they are equal; the same holds for $\lim_{t \rightarrow \infty} y(t)$ and $\lim_{p \rightarrow \infty} y(t_{2,p})$. Denote $L_x \triangleq \lim_{t \rightarrow \infty} x(t)$ and $L_y \triangleq \lim_{t \rightarrow \infty} y(t)$.

3.1. Fixed RECIPROCATION

In this section we prove that if both agents are *fixed* and at least one of them does not just mimic the other, but considers own kindness, then their action sequences converge;

³Agent i acts at the times in $T_i = \{t_{i,0} = 0, t_{i,1}, t_{i,2}, \dots\}$.

we also find the limits. We also show each sequence of actions alternates and that the relative positions between the two sequences of the two agents can be that one is always greater than the other or not.

In order to prove this theorem, we first show that it is sufficient to analyze the synchronous case, i.e., $T_1 = T_2 = T$.

Lemma 5. *Consider a pair of interacting agents. Denote the action sequences that would result in case both agents acted synchronously, (i.e., if held $T_1 = T_2 = T$), by $\{x'(t)\}_{t \in T}$ and $\{y'(t)\}_{t \in T}$, respectively. Then, the action sequences⁴ $\{x(t_{1,p})\}_{p \in \mathbb{N}}$ and $\{y(t_{2,p})\}_{p \in \mathbb{N}}$ are subsequences of $\{x'(t)\}_{t \in T}$ and $\{y'(t)\}_{t \in T}$, respectively, besides that a repetition streak in $\{x(t_{1,p})\}_{p \in \mathbb{N}}$ and $\{y(t_{2,p})\}_{p \in \mathbb{N}}$ may be represented by a single element in $\{x'(t)\}_{t \in T}$ and $\{y'(t)\}_{t \in T}$.*

The proof follows from Definition 12 by induction on time steps.

Proof. We prove by induction that for each $p > 0$, the sequence $x(t_{1,0}), x(t_{1,1}), \dots, x(t_{1,p})$ is a subsequence of the $x'(0), x'(1), \dots, x'(t_{1,p})$ and the sequence $y(t_{2,0}), y(t_{2,1}), \dots, y(t_{2,p})$ is a subsequence of the $y'(0), y'(1), \dots, y'(t_{2,p})$, perhaps, with removing some repetitions.

For $p = 0$, this is immediate, since $t_{i,0} = 0$.

For the induction step, assume that the lemma holds for $p - 1$ and prove it for $p > 0$. By definition, $x(t_{1,p}) = (1 - r_1) \cdot k_1 + r_1 \cdot y(t_{1,p-1})$, and since by the induction hypothesis, $y(t_{1,p-1})$ is an element in the sequence $\{y'(t)\}$, we conclude that $x(t_{1,p})$ is an element in the sequence $\{x'(t)\}$. Moreover, in $\{x'(t)\}$ this element comes after (or coincides with) $x(t_{1,p-1})$, because either $x(t_{1,p-1}) = x(0) = x'(0)$ or $x(t_{1,p-1}) = (1 - r_1) \cdot k_1 + r_1 \cdot y(t_{1,p-1} - 1)$ and $y(t_{1,p-1} - 1)$ precedes (or coincides with) $y(t_{1,p-1})$ in $\{y'(t)\}$ by the induction hypothesis. This proves the induction step for agent 1, and it is proven by analogy for 2. \square

We now assume the synchronous case and prove that the action sequences oscillate. Oscillation is an interesting property of the convergence process on its own right and is also helpful to prove convergence.

Lemma 6. *In the synchronous case, for every $t > 0$: $x(2t - 1) \geq x(2t + 1)$, and for every $t \geq 0$: $x(2t) \leq x(2t + 2) \leq x(2t + 1)$. By analogy, $\forall t > 0$: $y(2t - 1) \leq y(2t + 1)$, and $\forall t \geq 0$: $y(2t) \geq y(2t + 2) \geq y(2t + 1)$. All the inequations are strict if and only if $0 < r_1, r_2 < 1, k_2 > k_1$.⁵*

Since $x(2t) \leq x(2t + 2) \leq x(2t + 1) \Rightarrow x(2t) \leq x(2t + 1)$, we obtain for $t > 0$: $x(2t - 1) \geq x(2t + 1) \geq x(2t)$, and for every $t \geq 0$: $x(2t) \leq x(2t + 2) \leq x(2t + 1)$. By analogy, $\forall t > 0$: $y(2t - 1) \leq y(2t + 1) \leq y(2t)$, and $\forall t \geq 0$: $y(2t) \geq y(2t + 2) \geq y(2t + 1)$. Intuitively, this means that the sequence $\{x(t)\}$ is alternating while its amplitude is getting smaller, and the same holds for the sequence $\{y(t)\}$, with another alternation direction. The intuitive reasons are that first, agent 1 increases her action, while 2 decreases it. Then, since 2 has decreased her action, so does 1, while since 1 has increased hers, so does 2. This

⁴Agent i acts at the times in $T_i = \{t_{i,0} = 0, t_{i,1}, t_{i,2}, \dots\}$.

⁵Recall that we always assume that $k_2 \geq k_1$.

alternating process is going on, and the amplitude subsides, since the changing part is convexly combined with the constant one.

We now prove the lemma by induction on time.

Proof. For $t = 0$, we need to show that $x(0) \leq x(2) \leq x(1)$ and $y(0) \geq y(2) \geq y(1)$. We know that $x(0) = k_1$, $x(1) = (1 - r_1) \cdot k_1 + r_1 \cdot k_2$, and $y(0) = k_2$, $y(1) = (1 - r_2) \cdot k_2 + r_2 \cdot k_1$. Since $y(1) \leq k_2$, we have $x(2) = (1 - r_1) \cdot k_1 + r_1 \cdot y(1) \leq x(1)$. Since $y(1) \geq k_1$, we also have $x(2) = (1 - r_1) \cdot k_1 + r_1 \cdot y(1) \geq x(0)$. The proof that $y(0) \geq y(2) \geq y(1)$ is analogous.

For the induction step, for any $t > 0$, assume that the lemma holds for $t - 1$, which means $x(2t - 3) \geq x(2t - 1)$ (for $t > 1$), $x(2t - 2) \leq x(2t) \leq x(2t - 1)$, and $y(2t - 3) \leq y(2t - 1)$ (for $t > 1$), $y(2t - 2) \geq y(2t) \geq y(2t - 1)$.

We now prove the lemma for t . By Definition 12, $x(2t - 1) = (1 - r_1)k_1 + r_1y(2t - 2)$ and $x(2t + 1) = (1 - r_1)k_1 + r_1y(2t)$. Since we assume $y(2t - 2) \geq y(2t)$, we have $x(2t - 1) \geq x(2t + 1)$. By analogy, we can prove that $y(2t - 1) \leq y(2t + 1)$.

By definition, $x(2t) = (1 - r_1)k_1 + r_1y(2t - 1)$ and $x(2t + 2) = (1 - r_1)k_1 + r_1y(2t + 1)$. Since $y(2t - 1) \leq y(2t + 1)$, we have $x(2t) \leq x(2t + 2)$. By definition, $x(2t + 1) = (1 - r_1)k_1 + r_1y(2t)$. Since $y(2t) \geq y(2t - 1)$, we conclude that $x(2t + 1) \geq x(2t)$. By analogy, we prove that $y(2t + 1) \leq y(2t)$. From this and from the recursive definitions, we conclude that $x(2t + 2) \leq x(2t + 1)$, and we have shown that $x(2t) \leq x(2t + 2) \leq x(2t + 1)$. By analogy, we prove that $y(2t) \geq y(2t + 2) \geq y(2t + 1)$.

The equivalence of strictness in all the inequations to $0 < r_1, r_2 < 1, k_2 > k_1$ is proven by repeating the proof with strict inequalities, and by noticing that not having one of the conditions $0 < r_1, r_2 < 1, k_2 > k_1$ implies equality in at least one of the statements of the lemma, while if all these conditions hold, then the strictness holds. \square

With these results we now prove the following central theorem regarding convergence.

Theorem 15. *If the reciprocation coefficients are not both 1, which means $r_1 r_2 < 1$, then we have, for $i \in N$: $\lim_{p \rightarrow \infty} x_{i,j}(t_{i,p}) = \frac{(1-r_i)k_i + r_i(1-r_j)k_j}{1-r_i r_j}$.*

The idea is to prove monotonicity and boundedness, drawing convergence. The limits are found by substituting the limits to Definition 12.

Proof. Using Lemma 5, we assume the synchronous case. We first prove convergence, and then find its limit. For each agent, Lemma 6 implies that the even actions form a monotone sequence, and so do the odd ones. Both sequences are bounded, which can be easily proven by induction, and therefore each one converges. The whole sequence converges if and only if both limits are the same. We now show that they are indeed the same for the sequences $\{x(2t - 1)\}$ and $\{x(2t)\}$; the proof for $\{y(2t - 1)\}$ and $\{y(2t)\}$ is

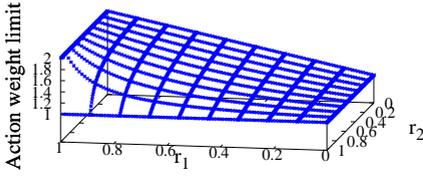


Figure 4.1: The limit of the actions of agent 1 as a function of the reciprocity coefficients, for a *Fixed-fixed* reciprocation, $k_1 = 1, k_2 = 2$. Given r_1 , agent 2 receives most when $r_2 = 0$.

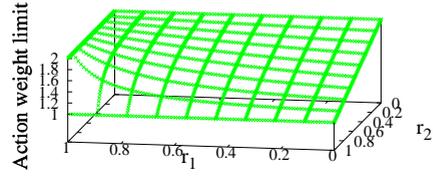


Figure 4.2: The limit of the actions of agent 2 as a function of the reciprocity coefficients, for a *Fixed-fixed* reciprocation, $k_1 = 1, k_2 = 2$. Given r_2 , agent 1 receives most when $r_1 = 1$.

analogous.

$$\begin{aligned}
 x(t+1) - x(t) &= (1 - r_1)k_1 + r_1y(t) - (1 - r_1)k_1 - r_1y(t-1) \\
 &= r_1(y(t) - y(t-1)) = r_1r_2(x(t-1) - x(t-2)) = \dots \\
 &= (r_1r_2)^{\lfloor t/2 \rfloor} \begin{cases} x(1) - x(0) & t = 2s, s \in \mathbb{N} \\ x(2) - x(1) & t = 2s + 1, s \in \mathbb{N}. \end{cases}
 \end{aligned}$$

As $r_1r_2 < 1$, this difference goes to 0 as t goes to ∞ . Thus, $x(t)$ converges (and so does $y(t)$). To find the limits $L_x = \lim_{t \rightarrow \infty} x(t)$ and $L_y = \lim_{t \rightarrow \infty} y(t)$, notice that in the limit we have $(1 - r_1)k_1 + r_1L_y = L_x$ and $(1 - r_2)k_2 + r_2L_x = L_y$ with the unique solution: $L_x = \frac{(1-r_1)k_1+r_1(1-r_2)k_2}{1-r_1r_2}$ and $L_y = \frac{(1-r_2)k_2+r_2(1-r_1)k_1}{1-r_1r_2}$. \square

Remark 2. *If, unlike the theorem assumes, $r_1r_2 = 1$, then since $r_1r_2 = 1$ if and only if $r_1 = r_2 = 1$, in the synchronous case, each agent just repeats what the other one did last time, thereby interchangeably playing k_1 and k_2 . In particular, unless $k_1 = k_2$, then no convergence takes place. If the synchronicity breaks by agent i acting alone, then both agents will act k_i from this time on.*

The theorem's assumption that not both reciprocation coefficients are 1 and the similar assumptions in the following theorems (such as $0 < r_i < 1$) mean that the agent neither ignores the other's action, nor does it copy the other's action. These are to be expected in real life, as we also mention in Example 6. The limits of actions as functions of the reciprocation coefficients are shown in Figures 4.1 and 4.2. We see that given r_1 , agent 2 receives most when $r_2 = 0$, and given r_2 , agent 1 receives most when $r_1 = 1$. This fits the intuition, that the kinder agent 2 increases interaction by not reacting to the less kind agent 1, while the opposite holds for agent 1. We now reconsider Example 5 for two agents.

Example 5 (Continued). *if agents 1 and 2 employ fixed reciprocation, $r_1 = r_2 = 0.5, r'_1 = r'_2 = 0.0$ and $k_1 = 0, k_2 = 0.5$, then we obtain $L_x = \frac{0.5 \cdot (1-0.5) \cdot 0.5}{1-0.5 \cdot 0.5} = 1/6$ and $L_y = \frac{(1-0.5) \cdot 0.5}{1-0.5 \cdot 0.5} = 1/3$.*

We see that $L_x \leq L_y$, which is intuitive, since the agents are always considering their kindness, so the kinder one acts with a bigger weight also in the limit. When this limit

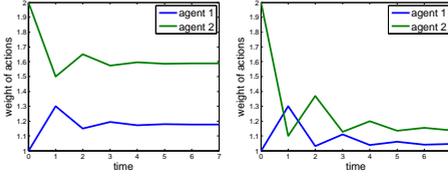


Figure 4.3: Simulation of actions for the synchronous case, with $r_1 + r_2 < 1$, $r_2 = 0.5$ on the left, and $r_1 + r_2 > 1$, $r_2 = 0.9$ on the right. This is a *fixed-fixed* reciprocation, with $k_1 = 1$, $k_2 = 2$, $r_1 = 0.3$. Each agent oscillates, while converging to her limit.

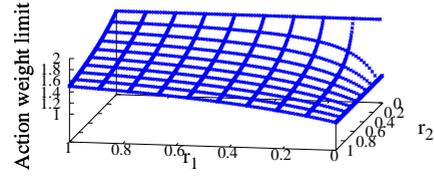


Figure 4.4: The common limit of the actions as a function of the reciprocity coefficients, for a *floating-floating* reciprocation, $k_1 = 1$, $k_2 = 2$. Given r_2 , agent 1 receives most when $r_1 = 1$, and given r_1 , agent 2 receives most when $r_2 = 0$.

inequality is strict, we have $x(t) < y(t)$ for all $t \geq t_0$ for some t_0 . To find when it is strict, consider the following:

$$\begin{aligned} \frac{(1-r_1)k_1 + r_1(1-r_2)k_2}{1-r_1r_2} &= \frac{(1-r_2)k_2 + r_2(1-r_1)k_1}{1-r_1r_2} \\ \iff (1-r_1)(1-r_2)k_1 &= (1-r_1)(1-r_2)k_2 \\ \iff r_1 = 1 \vee r_2 = 1 \vee k_1 &= k_2, \end{aligned}$$

and thus, it is strict if and only if $r_1 < 1 \wedge r_2 < 1 \wedge k_1 < k_2$. Even when the limit inequality is strict, $x(t) < y(t)$ may hold only from some time on, and not all the way. For illustration, in the simulation of the actions over time in Figure 4.3, on the left, $y(t)$ is always larger than $x(t)$, and on the right, they alternate several times before $y(t)$ becomes larger. We observe oscillations, predicted by Lemma 6, and rapid convergence, which will ensue from Theorem 20.

3.2. Floating RECIPROCATION

We first prove convergence, and later we analyze the relative position between the two sequences of actions and the monotonicity of each of these sequences.

If both agents have the *floating* reciprocation attitude, their action sequences converge to a common limit, as the following important theorem states.

Theorem 16. *If the reciprocation coefficients are neither both 0 nor both 1, which means $0 < r_1 + r_2 < 2$, then, as $t \rightarrow \infty$, $x(t)$ and $y(t)$ converge to a common limit. In the synchronous case ($T_1 = T_2 = T$), they both approach*

$$\frac{1}{2} \left(k_1 + k_2 + (k_2 - k_1) \frac{r_1 - r_2}{r_1 + r_2} \right) = \frac{r_2}{r_1 + r_2} k_1 + \frac{r_1}{r_1 + r_2} k_2.$$

The common limit of the actions is shown in Figure 4.4. As in the *fixed, fixed* case, we observe that to boost cooperation, the kindest should be stable, while the less kind should mimic the kindest.

In Example 5, we have the following.

Example 5 (Continued). *if agents 1 and 2 employ floating reciprocation, $r_1 = r_2 = 0.5$, $r'_1 = r'_2 = 0.0$ and $k_1 = 0$, $k_2 = 0.5$, then we obtain $L_x = L_y = (1/2) \cdot 0 + (1/2) \cdot 0.5 = 0.25$.*

The idea of the proof is to show that $\{[\min\{x(t), y(t)\}, \max\{x(t), y(t)\}]\}_{t=1}^{\infty}$ is a nested sequence of segments, which lengths approach zero, and therefore, $\{x(t)\}$ and $\{y(t)\}$ converge to a common limit. Finding this limit stems from finding $\lim_{t \rightarrow \infty} (x(t) + y(t))$.

Throughout the section, whenever we need concrete T_1, T_2 , we consider the synchronous case. The alternative case is fully handled in Section 4.

Proof. We first prove that the convergence takes place.

If both agents act at time $t > 0$, then $y(t) - x(t)$

$$\begin{aligned} &= x(t-1)(r_2 - 1 + r_1) + y(t-1)(1 - r_2 - r_1) \\ &= y(t-1)(1 - r_1 - r_2) - x(t-1)(1 - r_1 - r_2) \\ &= (1 - r_1 - r_2)(y(t-1) - x(t-1)). \end{aligned} \quad (4.1)$$

Since $0 < r_1 + r_2 < 2$, we have $|(1 - r_1 - r_2)| < 1$.

If only agent 1 acts at time $t > 0$, then $y(t) - x(t)$

$$\begin{aligned} &= y(t-1)(1 - r_1) - x(t-1)(1 - r_1) \\ &= (1 - r_1)(y(t-1) - x(t-1)). \end{aligned} \quad (4.2)$$

If $r_1 > 0$, then $|1 - r_1| < 1$. Similarly, if only agent 2 acts, then

$$y(t) - x(t) = (1 - r_2)(y(t-1) - x(t-1)). \quad (4.3)$$

Since $r_1 + r_2 > 0$, either r_1 or r_2 is greater than 0, and since each agent acts an infinite number of times, we obtain $\lim_{t \rightarrow \infty} |y(t) - x(t)| = 0$. Since $\forall t > 0 : x(t), y(t) \in [\min\{x(t-1), y(t-1)\}, \max\{x(t-1), y(t-1)\}]$, we have a nested sequence of segments, which lengths approach zero, thus $x(t)$ and $y(t)$ both converge, and to a common limit.

Assuming $T_1 = T_2 = T$ now, we find the common limit. For all $t > 0$,

$$\begin{aligned} x(t) + y(t) &= x(t-1)(1 - r_1 + r_2) + y(t-1)(r_1 + 1 - r_2) \\ &= x(t-1) + y(t-1) + (r_1 - r_2)(y(t-1) - x(t-1)) \\ &\Rightarrow \lim_{t \rightarrow \infty} x(t) + y(t) = k_1 + k_2 + \sum_{t=0}^{\infty} (r_1 - r_2)(y(t) - x(t)) \\ &\stackrel{(4.1)}{=} k_1 + k_2 + (r_1 - r_2) \sum_{t=0}^{\infty} (1 - r_1 - r_2)^t (k_2 - k_1) \\ &\underbrace{\quad}_{\text{geom. series}} \xrightarrow{t \rightarrow \infty} k_1 + k_2 + (r_1 - r_2) \frac{k_2 - k_1}{r_1 + r_2} = k_1 + k_2 + (k_2 - k_1) \frac{r_1 - r_2}{r_1 + r_2}. \end{aligned}$$

Since we have shown that both limits exist and are equal, each is equal to half of $k_1 + k_2 + (k_2 - k_1) \frac{r_1 - r_2}{r_1 + r_2}$. \square

Remark 3. *If, unlike the theorem assumes, $r_1 + r_2 = 0$, then since $r_1 + r_2 = 0$ if and only if $r_1 = r_2 = 0$, each agent keeps doing the same thing all the time: agent 1 does k_1 and 2 does k_2 . If, unlike the theorem assumes, $r_1 + r_2 = 2$, then since $r_1 + r_2 = 2$ if and only if $r_1 = r_2 = 1$, in the synchronous case, each agent just repeats what the other one did last time, thereby interchangeably playing k_1 and k_2 . In particular, unless $k_1 = k_2$, then no convergence takes place. If the synchronicity breaks by agent i acting alone, then both agents will act k_i from this time on.*

The following gives the relation between $x(t)$ s and $y(t)$ s.

Proposition 1. *If $r_1 + r_2 \leq 1$, then, for every $t \geq 0$: $y(t) \geq x(t)$.*

If $r_1 + r_2 \geq 1$, then the following holds. $y(0) \geq x(0)$. For every $t > 0$, $t \in T_1 \cap T_2$, we have $y(t-1) \geq x(t-1) \Rightarrow y(t) \leq x(t)$, and $y(t-1) \leq x(t-1) \Rightarrow y(t) \geq x(t)$. For any other t , we have $y(t-1) \geq x(t-1) \Rightarrow y(t) \geq x(t)$, and $y(t-1) \leq x(t-1) \Rightarrow y(t) \leq x(t)$. In words, $x(t)$ s and $y(t)$ s alter their relative positions if and only if both act.

Proof. Consider the case $r_1 + r_2 \leq 1$ first. We employ induction. The basis is $t = 0$, where $y(0) = k_2 \geq k_1 = x(0)$.

For the induction step, we assume the proposition for all the times smaller than $t > 0$ and prove it for t . If only 1 acts at t , then $y(t) = y(t-1)$ and $x(t) = (1-r_1)x(t-1) + r_1y(t-1)$. Therefore, $y(t) \geq x(t)$ if and only if $y(t-1) \geq (1-r_1)x(t-1) + r_1y(t-1)$, which is equivalent to $(1-r_1)y(t-1) \geq (1-r_1)x(t-1)$, which holds by the induction hypothesis. If only agent 2 acts at t , then $x(t) = x(t-1)$ and $y(t) = (1-r_2)y(t-1) + r_2x(t-1)$. Therefore,

$$y(t) \geq x(t) \iff (1-r_2)y(t-1) + r_2x(t-1) \geq x(t-1) \iff (1-r_2)y(t-1) \geq (1-r_2)x(t-1),$$

which is true by the induction hypothesis.

If both agents act at t , then $x(t) = (1-r_1)x(t-1) + r_1y(t-1)$ and $y(t) = (1-r_2)y(t-1) + r_2x(t-1)$. Therefore, $y(t) \geq x(t)$ if and only if $(1-r_2)y(t-1) + r_2x(t-1) \geq (1-r_1)x(t-1) + r_1y(t-1)$ if and only if $(1-r_1-r_2)y(t-1) \geq (1-r_1-r_2)x(t-1)$, which is true by the induction hypothesis and using the assumption $r_1 + r_2 \leq 1$.

Consider the case $r_1 + r_2 \geq 1$ now. We employ induction again. The basis is $t = 0$, where $y(0) = k_2 \geq k_1 = x(0)$.

For the induction step, assume the proposition for all values smaller than $t > 0$ and prove it for t . The cases where only agent 1 acts at t and where only 2 acts at t are shown analogously to how they are shown for the case $r_1 + r_2 \leq 1$. If both agents act at t , then we have shown that $y(t) \geq x(t)$ if and only if $(1-r_1-r_2)y(t-1) \geq (1-r_1-r_2)x(t-1)$, which means that $y(t-1) \geq x(t-1) \Rightarrow y(t) \leq x(t)$ and $y(t-1) \leq x(t-1) \Rightarrow y(t) \geq x(t)$, assuming $r_1 + r_2 \geq 1$. \square

We now discuss the question of monotonicity of action sequences.

Proposition 2. *Assume $r_1 < 1, r_2 < 1$. Then, if $r_1 + r_2 \leq 1$, then $\{x(t)\}$ do not decrease and $\{y(t)\}$ do not increase. On the other hand, if $r_1 + r_2 > 1$, then both $\{x(t)\}$ and $\{y(t)\}$ are not monotonic, unless $T_1 \cap T_2 = \{0\}$, in which case they are monotonic. In the action times alternate, then each agent's actions alternate in positions.*

Proof. The proposition implies that if $r_1 + r_2 \leq 1$, then $\{x(t)\}$ do not decrease and $\{y(t)\}$ do not increase, since the next $x(t)$ (or $y(t)$) is either the same or a combination of the last one with a higher value (lower value, for $y(t)$).

For $r_1 + r_2 > 1$, both $\{x(t)\}$ and $\{y(t)\}$ are not monotonic, unless $T_1 \cap T_2 = \{0\}$, in which case they are monotonic, for the reason above (in this case we always have $y(t) \geq x(t)$). For $T_1 \cap T_2 \neq \{0\}$, take any positive t in $T_1 \cap T_2$. Then the larger value at $t-1$ becomes the smaller one at t , thereby getting smaller, and the smaller value gets larger analogously. In the future, the new smaller will only grow and the new larger will decrease, thereby behaving non-monotonically. In particular, in the alternating case, each agent's actions alternate. \square

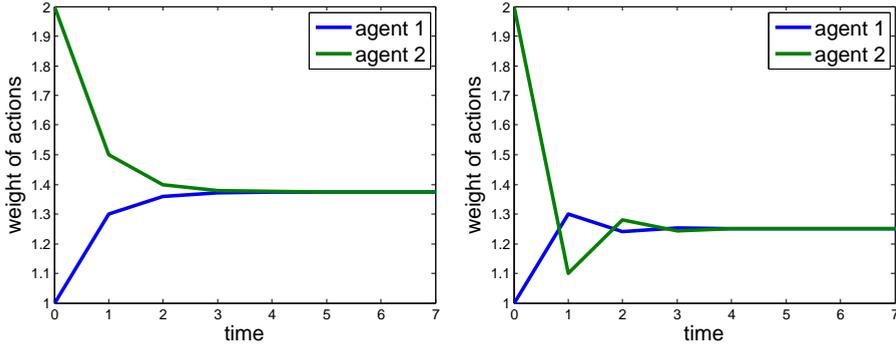


Figure 4.5: Simulation results for the synchronous case, $r_1 + r_2 < 1$ (left) and $r_1 + r_2 > 1$ (right). We assume a *floating-floating* reciprocation, $k_1 = 1, k_2 = 2, r_1 = 0.3$. In the left graph, $r_2 = 0.5$, while in the right one, $r_2 = 0.9$. In the left graph, agent 1's actions are smaller than those of 2; agent 1's actions increase, while those of agent 2 decrease. In the right graph, the actions of the agents alter their relative positions at each time step, and each agent's actions go up and down.

Some examples are simulated in Figure 4.5, fitting Theorem 16 and showing monotonicity and the actions of agent 2 being larger in case $r_1 + r_2 \leq 1$.

3.3. Fixed AND Floating RECIPROCATION

We prove monotonicity, and subsequently use it to establish convergence. Finally, we analyze the relative position between the two sequences of actions.

Assume that agent 1 employs the *fixed* reciprocation attitude, while 2 acts by the *floating* reciprocation. We will show the convergence Theorem 17 using the following lemma, proving monotonicity of the action sequences from some time on.

Lemma 7. *If $r_2 > 0$ and $r_1 + r_2 \leq 1$, then, for every $t \geq t_{1,1} : x(t+1) \leq x(t)$, and for every $t \geq 0 : y(t+1) \leq y(t)$.*

This lemma assumes the agents do not react too strongly, since $r_1 + r_2$ is at most 1. We now prove this lemma by induction on t , using the definition of reciprocation.

Proof. We employ induction. The basis consists of the following subcases: $t = 0, 0 < t < t_{1,1}$ and $t = t_{1,1}$. For $t = 0$, we have either $y(1) = y(0)$ or $y(1) = (1 - r_2)k_2 + r_2k_1 \leq k_2 = y(0)$.

For any $0 < t < t_{1,1}$, we have either $y(t+1) = y(t)$ or $y(t+1) = (1 - r_2)y(t) + r_2k_1 \stackrel{k_1 \leq y(t)}{\leq} (1 - r_2)y(t) + r_2y(t) = y(t)$.

For $t = t_{1,1}$, we either have $x(t_{1,1} + 1) = x(t_{1,1})$ or $x(t_{1,1} + 1) = (1 - r_1)k_1 + r_1y(t_{1,1})$, and anyway $x(t_{1,1}) = (1 - r_1)k_1 + r_1y(t_{1,1} - 1)$ by the definition of $t_{1,1}$. Since $y(t_{1,1}) \leq y(t_{1,1} - 1)$ by the induction hypothesis, we have $x(t_{1,1} + 1) \leq x(t_{1,1})$. As to $y(t)$ s, we have either $y(t_{1,1} + 1) = y(t_{1,1})$ or

$$\begin{aligned} y(t_{1,1} + 1) &= (1 - r_2)y(t_{1,1}) + r_2x(t_{1,1}) \\ \Rightarrow y(t_{1,1} + 1) &\leq y(t_{1,1}) \stackrel{r_2 > 0}{\iff} x(t_{1,1}) \leq y(t_{1,1}). \end{aligned}$$

Either $y(t_{1,1}) = y(t_{1,1} - 1)$ or $y(t_{1,1}) = (1 - r_2)y(t_{1,1} - 1) + r_2k_1$. In the first case,

$$\begin{aligned} x(t_{1,1}) \leq y(t_{1,1}) &\iff (1 - r_1)k_1 + r_1y(t_{1,1} - 1) \leq y(t_{1,1} - 1) \\ &\iff (1 - r_1)k_1 \leq (1 - r_1)y(t_{1,1} - 1), \end{aligned}$$

which always holds. In the second case,

$$\begin{aligned} x(t_{1,1}) \leq y(t_{1,1}) &\iff (1 - r_1)k_1 + r_1y(t_{1,1} - 1) \leq (1 - r_2)y(t_{1,1} - 1) + r_2k_1 \\ &\iff (1 - r_1 - r_2)k_1 \leq (1 - r_1 - r_2)y(t_{1,1} - 1), \end{aligned}$$

which is true, since $k_1 \leq y(t_{1,1} - 1)$ and $(r_1 + r_2) \leq 1$. Thus, the basis is proven.

For the induction step, for any $t > t_{1,1}$, assume that $x(t) \leq x(t - 1) \leq \dots \leq x(t_{1,1})$, and $y(t) \leq y(t - 1) \leq \dots \leq y(0)$.

We now prove the lemma for t . If $x(t + 1) = x(t)$, then trivially $x(t + 1) \leq x(t)$. Otherwise, $x(t + 1) = (1 - r_1)k_1 + r_1y(t) \leq (1 - r_1)k_1 + r_1y(s_1(t) - 1) = x(t)$, where the inequality stems from the induction hypothesis. For $y(t)$ s, if $y(t + 1) = y(t)$, then trivially $y(t + 1) \leq y(t)$. Otherwise, $y(t + 1) = (1 - r_2)y(t) + r_2x(t) \leq (1 - r_2)y(s_2(t) - 1) + r_2x(s_2(t) - 1) = y(t)$. The above inequality stems from the induction hypothesis, if $s_2(t) - 1 \geq t_{1,1}$, so that the induction hypothesis for $x(t)$ s holds as well as the one for $y(t)$ s. Otherwise ($s_2(t) \leq t_{1,1}$), the above inequality is proven as follows:

$$\begin{aligned} (1 - r_2)y(t) + r_2x(t) &\leq (1 - r_2)y(s_2(t) - 1) + r_2x(s_2(t) - 1) \\ &\iff \\ (1 - r_2)(y(t) - y(s_2(t) - 1)) &\leq r_2(x(s_2(t) - 1) - x(t)) = r_2(k_1 - x(t)). \end{aligned}$$

Notice that $(y(t) - y(s_2(t) - 1)) = y(s_2(t)) - y(s_2(t) - 1) \stackrel{s_2(t) \leq t_{1,1}}{=} (1 - r_2)y(s_2(t) - 1) + r_2k_1 - y(s_2(t) - 1) = r_2(k_1 - y(s_2(t) - 1))$. Thus, continuing the above chain of equivalences,

$$\begin{aligned} (1 - r_2)(y(t) - y(s_2(t) - 1)) &\leq r_2(k_1 - x(t)) \\ \iff (1 - r_2)r_2(k_1 - y(s_2(t) - 1)) &\leq r_2(k_1 - x(t)) \\ \stackrel{r_2 > 0}{\iff} (1 - r_2)(k_1 - y(s_2(t) - 1)) &\leq (k_1 - x(t)) \iff x(t) - r_2k_1 \leq (1 - r_2)y(s_2(t) - 1). \quad (4.4) \end{aligned}$$

To show this, notice that $s_2(t) \leq t_{1,1} < t \Rightarrow s_2(t) + 1 \leq t$. In addition, agent 2 acts at time slot $t_{1,1} - 1$ (since 1 does not), and therefore $s_2(t) \geq t_{1,1} - 1$. Therefore, using the induction hypothesis for $x(t)$ s we obtain

$$x(t) \stackrel{t \geq s_2(t) + 1 \geq t_{1,1}}{\leq} x(s_2(t) + 1) = (1 - r_1)k_1 + r_1y(s_2(t)),$$

where the equality stems from the fact that if $s_2(t) + 1 \notin T_1$, then it is in T_2 , and therefore, $t = s_2(t)$, a contradiction. Therefore,

$$\begin{aligned} x(t) - r_2k_1 &\leq (1 - r_1 - r_2)k_1 + r_1y(s_2(t)) \leq (1 - r_1 - r_2)y(s_2(t)) + r_1y(s_2(t)) \\ &= (1 - r_2)y(s_2(t)) \leq (1 - r_2)y(s_2(t) - 1), \end{aligned}$$

and (4.4) has been proven. The chain of equivalences that ends with (4.4) begins with $(1 - r_2)y(t) + r_2x(t) \leq (1 - r_2)y(s_2(t) - 1) + r_2x(s_2(t) - 1)$, and we have $y(t + 1) = (1 - r_2)y(t) + r_2x(t) \leq (1 - r_2)y(s_2(t) - 1) + r_2x(s_2(t) - 1) = y(t)$, completing the induction step. \square

This lemma provides monotonicity from some point on, enabling us to prove the convergence of the process when the agents react mildly.

Theorem 17. *If $r_2 > 0$ and $r_1 + r_2 \leq 1$, then, $\lim_{t \rightarrow \infty} x(y) = \lim_{t \rightarrow \infty} y(t) = k_1$.*

Proof. We first prove that the convergence takes place, and then find its limit. For each agent, Lemma 7 implies that her actions are monotonically non-increasing. Since the actions are bounded below by k_1 , which can be easily proven by induction, both action sequences converge.

To find the limits, notice that in the limit we have

$$(1 - r_1)k_1 + r_1L_y = L_x \quad (4.5)$$

$$(1 - r_2)L_y + r_2L_x = L_y. \quad (4.6)$$

From (4.6), we conclude that $L_x = L_y$, since $r_2 > 0$. Substituting this to (4.5) gives us $L_x = L_y = k_1$, since $r_2 > 0$ and $r_1 + r_2 \leq 1$ together imply $r_1 < 1$. \square

Remark 4. *If, unlike the theorem assumes, $r_2 = 0$, then agent 2 keeps doing the same thing all the time: k_2 , and agent 1 keeps doing $(1 - r_1)k_1 + r_1k_2$ all the time when $t > 0$. If, unlike the theorem assumes, $r_1 + r_2 > 1$, but the rest holds, then it is still open what happens.*

The relation between the sequences of $x(t)$ s and $y(t)$ s is given by the following proposition (also covering the case $r_1 + r_2 \geq 1$, when agents react actively too each other).

Proposition 3. *If $r_1 + r_2 \leq 1$, then for every $t \geq 0$: $y(t) \geq x(t)$.*

If $r_1 + r_2 \geq 1$, then $y(0) \geq x(0)$. For every $t > 0$ such that $t \in T_1 \cap T_2$, we have $y(t - 1) \leq x(t - 1) \Rightarrow y(t) \geq x(t)$. For any $t \in T_1 \setminus T_2$, we have $y(t) \geq x(t)$, and for any $t \in T_2 \setminus T_1$, we have $y(t - 1) \geq x(t - 1) \Rightarrow y(t) \geq x(t)$, and $y(t - 1) \leq x(t - 1) \Rightarrow y(t) \leq x(t)$.

Proof. Consider the case $r_1 + r_2 \leq 1$ first. We employ induction. The basis is $t = 0$, where $y(0) = k_2 \geq k_1 = x(0)$.

For the induction step, we assume the proposition for all values smaller than $t > 0$ and prove the proposition for t . If only agent 1 acts at t , then $y(t) = y(t - 1)$ and $x(t) = (1 - r_1)k_1 + r_1y(t - 1)$. Therefore,

$$\begin{aligned} y(t) \geq x(t) &\iff y(t - 1) \geq (1 - r_1)k_1 + r_1y(t - 1) \\ &\iff (1 - r_1)y(t - 1) \geq (1 - r_1)k_1, \end{aligned}$$

which is true.

If only agent 2 acts at t , then $x(t) = x(t - 1)$ and $y(t) = (1 - r_2)y(t - 1) + r_2x(t - 1)$. Therefore,

$$\begin{aligned} y(t) \geq x(t) &\iff (1 - r_2)y(t - 1) + r_2x(t - 1) \geq x(t - 1) \\ &\iff (1 - r_2)y(t - 1) \geq (1 - r_2)x(t - 1), \end{aligned}$$

which is true by the induction hypothesis.

If both agents act at t , then $x(t) = (1 - r_1)k_1 + r_1y(t - 1)$ and $y(t) = (1 - r_2)y(t - 1) + r_2x(t - 1)$. Therefore,

$$\begin{aligned} y(t) \geq x(t) &\iff (1 - r_2)y(t - 1) + r_2x(t - 1) \geq (1 - r_1)k_1 + r_1y(t - 1) \\ &\iff (1 - r_1 - r_2)y(t - 1) \geq (1 - r_1)k_1 - r_2x(t - 1). \end{aligned}$$

Since $x(t - 1) \geq k_1$, it is enough to show that $(1 - r_1 - r_2)y(t - 1) \geq (1 - r_1)x(t - 1) - r_2x(t - 1) = (1 - r_1 - r_2)x(t - 1)$, which is true by the induction hypothesis and using the assumption $r_1 + r_2 \leq 1$. Thus, the case $r_1 + r_2 \leq 1$ has been proven.

Consider the case $r_1 + r_2 \geq 1$ now. We employ induction. The basis is $t = 0$, where $y(0) = k_2 \geq k_1 = x(0)$.

For the induction step, we assume the proposition for all values smaller than $t > 0$ and prove the proposition for t . The cases where only agent 1 acts at t and where only 2 acts at t are shown by analogy to how they are shown for the case $r_1 + r_2 \geq 1$. If both agents act at t , then we have shown that

$$y(t) \geq x(t) \iff (1 - r_1 - r_2)y(t - 1) \geq (1 - r_1)k_1 - r_2x(t - 1). \quad (4.7)$$

Now, if $y(t - 1) \leq x(t - 1)$, then $(1 - r_1 - r_2)y(t - 1) \geq (1 - r_1 - r_2)x(t - 1) \geq (1 - r_1)k_1 - r_2x(t - 1)$, and from (4.7) we have $y(t) \geq x(t)$. \square

We have not seen yet whether Theorem 17 holds for strong reactions ($r_1 + r_2 > 1$). Nonetheless, we do know that monotonicity from some time on (Lemma 7) does not hold. We also know that $y(t)$ being always at least as large as $x(t)$ or the other way around fails to hold in this case. As a counterexample for both of them, consider the case of $r_2 = 1, 0 < r_1 < 1, k_2 > k_1$. One can readily prove by induction that for all t we have $x(2t + 1) > x(2t) = x(2t + 2)$ and $y(2t) > y(2t - 1) = y(2t + 1)$, and thus both sequences are not monotonic. In addition, one can inductively prove that $x(2t + 1) > y(2t + 1), x(2t) < y(2t)$, and therefore no sequence is always larger than the other one.

Figure 4.6 shows how the actions evolve over time, fitting Theorem 17, Lemma 7 and Proposition 3 for $r_1 + r_2 \leq 1$. The actions seem to converge also in the unproven case $r_1 + r_2 > 1$.

In the case of the mirroring assumption that agent 1 acts according to the *floating* reciprocation attitude, while 2 acts according to the *fixed* reciprocation, we can obtain the following similar results by analogy.

Theorem 18. *If $r_1 > 0$ and $r_1 + r_2 \leq 1$, then, $\lim_{t \rightarrow \infty} x(t) = \lim_{t \rightarrow \infty} y(t) = k_2$.*

The proof is analogous, with the lemma being

Lemma 8. *If $r_1 > 0$ and $r_1 + r_2 \leq 1$, then, for every $t \geq t_{2,1} : y(t + 1) \geq y(t)$, and for every $t \geq 0 : x(t + 1) \geq x(t)$.*

Remark 5. *If, unlike the theorem assumes, $r_1 = 0$, then agent 1 keeps doing the same thing all the time: k_1 , and agent 2 keeps doing $(1 - r_2)k_2 + r_2k_1$ all the time when $t > 0$. If, unlike the theorem assumes, $r_1 + r_2 > 1$, but the rest holds, then it is still open what happens.*

Regarding the relation between $x(t)$ s and $y(t)$ s, we prove the following, by analogy to how Proposition 3 is proven:

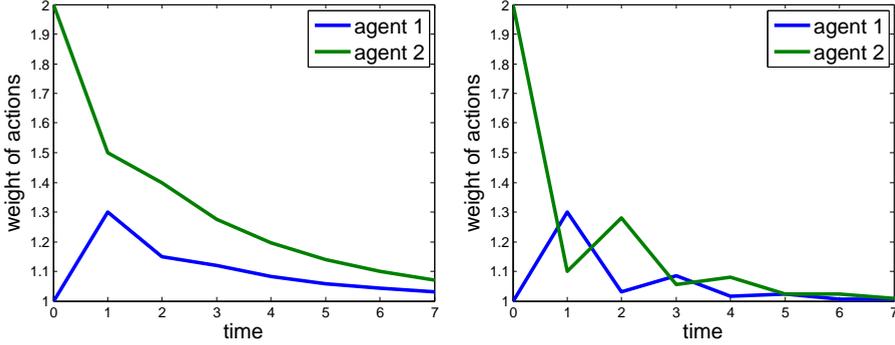


Figure 4.6: Simulation of actions for the synchronous case, with $r_1 + r_2 < 1$, $r_2 = 0.5$ on the left, and $r_1 + r_2 > 1$, $r_2 = 0.9$ on the right. This is a *fixed-floating* reciprocation, with $k_1 = 1, k_2 = 2, r_1 = 0.3$. In the left graph, agent 1's actions are smaller than those of 2; agent 1's actions decrease after $t = 1$, while those of agent 2 decrease all the time. The value of the common limit fits the theorem's prediction.

Proposition 4. *If $r_1 + r_2 \leq 1$, then for every $t \geq 0$: $y(t) \geq x(t)$. If $r_1 + r_2 \geq 1$, then, $y(0) \geq x(0)$. For every $t > 0$, such that $t \in T_1 \cap T_2$, we have $y(t - 1) \leq x(t - 1) \Rightarrow y(t) \geq x(t)$. For any $t \in T_2 \setminus T_1$, we have $y(t) \geq x(t)$, and for any $t \in T_1 \setminus T_2$, we have $y(t - 1) \geq x(t - 1) \Rightarrow y(t) \geq x(t)$ and $y(t - 1) \leq x(t - 1) \Rightarrow y(t) \leq x(t)$.*

To solve the case of $r_1 + r_2 > 1$, we will prove the following crucial corollary from Proposition 7. This corollary is used further, rather than the limited Theorem 17.

Corollary 3. *Consider pairwise interaction, where one agent i employs fixed reciprocation and the other agent j employs the floating one, and every agent acts at least once every q times. Assume that $r_i < 1$ and $r_j > 0$. Then, both limits exist and are equal to k_i . The convergence is geometrically fast.*

We omit the proof at this stage.

For all the considered cases, we have the following

Proposition 5. *If both L_x and L_y exist, then $L_x \leq L_y$.*

We now continue to study alternating pairwise interaction.

4. PAIRWISE INTERACTION: ALTERNATING CASE

In the previous section, we sometimes show more when assuming that the interaction is synchronous. However, Theorem 16 can be extended for the alternating case (T_1 contains precisely all the even times and T_2 contains zero and all the odd ones) as follows:

Theorem 19. *In the case where agents act alternately, which is when T_1 contains precisely all the even times and T_2 contains zero and all the odd ones, they both approach*

$$\frac{1}{2} \left(k_1 + k_2 + \frac{(r_1 - r_2 - r_1 r_2)}{r_1 + r_2 - r_1 r_2} (k_2 - k_1) \right) = \frac{r_2}{r_1 + r_2 - r_1 r_2} k_1 + \frac{r_1 - r_1 r_2}{r_1 + r_2 - r_1 r_2} k_2.$$

The idea of the proof is proving that $x(t) + y(t)$ approach $k_1 + k_2 + \frac{(r_1 - r_2 - r_1 r_2)}{r_1 + r_2 - r_1 r_2} (k_2 - k_1)$.

Proof. We now assume the alternating case. Consider the behavior of $x(t) + y(t)$. For an even $t > 0$, only agent 1 acts and we have

$$x(t) + y(t) = x(t-1) + y(t-1) + (r_1)(y(t-1) - x(t-1)).$$

For an odd t , only 2 acts and we have

$$x(t) + y(t) = x(t-1) + y(t-1) + (-r_2)(y(t-1) - x(t-1)).$$

And therefore, we have

$$\begin{aligned} \Rightarrow \lim_{t \rightarrow \infty} x(y) + y(t) &= k_1 + k_2 + \sum_{t=0}^{\infty} (-r_2)(y(2t) - x(2t)) + \sum_{t=0}^{\infty} (r_1)(y(2t+1) - x(2t+1)) \\ &\stackrel{(4.2),(4.3)}{=} k_1 + k_2 - (r_2) \cdot \sum_{t=0}^{\infty} ((1-r_1)^t (1-r_2)^t) (k_2 - k_1) + (r_1) \cdot \sum_{t=0}^{\infty} ((1-r_1)^t (1-r_2)^{t+1}) (k_2 - k_1) \\ &= k_1 + k_2 + \sum_{t=0}^{\infty} (1-r_1)^t (1-r_2)^t (r_1 - r_2 - r_1 r_2) (k_2 - k_1) \\ &\quad \underbrace{\text{geom. series}}_{t \rightarrow \infty} k_1 + k_2 + (r_1 - r_2 - r_1 r_2) (k_2 - k_1) \frac{1}{1 - (1-r_1)(1-r_2)} \\ &= k_1 + k_2 + \frac{(r_1 - r_2 - r_1 r_2)}{r_1 + r_2 - r_1 r_2} (k_2 - k_1). \end{aligned}$$

Since we have shown that both limits exist and are equal, each equals to half of $k_1 + k_2 + \frac{(r_1 - r_2 - r_1 r_2)}{r_1 + r_2 - r_1 r_2} (k_2 - k_1)$. \square

Remark 6. *If, unlike the theorem assumes, $r_1 + r_2 = 2$, then since $r_1 + r_2 = 2 \iff r_1 = r_2 = 1$, in the alternating case, agent 2 plays at time 1 the strategy of agent 1 at time 0, which is k_1 , and since then, each player plays it.*

Having dealt with the pairwise interaction, we move to the general case.

5. MULTI-AGENT INTERACTION

Continuing predicting reciprocal interaction to facilitate decision support, we now analyze the general reciprocal interaction, when agents interact with many agents. To formally discuss the actions after the interaction has settled down, we consider the limits (if exist)⁶ $\lim_{p \rightarrow \infty} x_{i,j}(t_{1,p})$, and $\lim_{t \rightarrow \infty} x_{i,j}(t)$, for agents i and j . Since the sequence $\{x_{i,j}(t)\}$ is $\{x_{i,j}(t_{1,p})\}$ with finite repetitions, the limit $\lim_{p \rightarrow \infty} x_{i,j}(t_{1,p})$ exists if and only if $\lim_{t \rightarrow \infty} x_{i,j}(t)$ does. If they exist, they are equal. Denote $L_{i,j} \triangleq \lim_{t \rightarrow \infty} x_{i,j}(t)$.

We first provide general convergence results, and then we find the common limit for the case when at most one agent is *fixed* and synchronous in Theorem 20. This allows concluding about how an agent can maximize the common value by picking her reciprocation coefficient. Finally, we prove convergence for the general case of any acting dynamics described by a contraction (see Definition 15). In this section, the ambivalent case of $r_i + r'_i = 1$, which can be taken as either *fixed* or *floating*, is taken to be *floating*.

First, we have convergence for the case of *floating* agents.

⁶Agent i acts at the times in $T_i = \{t_{i,0} = 0, t_{i,1}, t_{i,2}, \dots\}$.

Proposition 6. Consider a connected interaction graph, where all the agents are floating and for every agent i , $r_i + r'_i < 1$. Then, for all pairs of agents $i \neq j$ such that $(i, j) \in E$, the limit $L_{i,j}$ exists; all these limits are equal to each other.

Proof. Follows directly from [29, Theorem 2]. This article and similar articles on multi-agent coordination [30, 31] prove convergence when all the agents are *floating*. \square

We now show convergence, when some agents are *fixed*.

Proposition 7. Consider a connected interaction graph, where for all the agents i , $r'_i > 0$. Assume that at least one agent employs the fixed attitude and every agent acts at least once every q times, for a natural $q > 0$. Then, for all pairs of agents $i \neq j$ such that $(i, j) \in E$, the limit $L_{i,j}$ exists. The convergence is geometrically fast.

The proof expresses the dependency of actions on the previous actions as a matrix multiplication and directly proves convergence.

Proof. We express how each action depends on the actions in the previous time in matrix $A(t) \in \mathbb{R}_+^{|\mathcal{E}| \times |\mathcal{E}|}$, which, in the synchronous case, is defined as follows:

$$A(t)((i, j), (k, l)) \triangleq \begin{cases} (1 - r_i - r'_i) & \text{if } k = i, l = j; \\ r_i + r'_i \frac{1}{|\mathcal{N}^+(i)|} & \text{if } k = j, l = i; \\ r'_i \frac{1}{|\mathcal{N}^+(i)|} & \text{if } k \neq j, l = i; \\ 0 & \text{otherwise,} \end{cases} \quad (4.8)$$

where the first line is missing for the *fixed* agents, since for them, own behavior does not matter. If, for each time $t \in T$, the column vector $\vec{p}(t) \in \mathbb{R}_+^{|\mathcal{E}|}$ describes the actions at time t , in the sense that its (i, j) th coordinate contains $x_{i,j}(t)$ (for $(i, j) \in E$), then we have $\vec{p}(t+1) = A(t+1)\vec{p}(t) + \vec{k}'$, where \vec{k}' is the relevant kindness vector, formally defined as

$$k'(t)((i, j)) \triangleq \begin{cases} (1 - r_i - r'_i)k_i & \text{if } i \text{ is fixed;} \\ 0 & \text{otherwise.} \end{cases}$$

In a not necessarily synchronous case, only a subset of agents act at a given time t . For an acting agent i , every $A(t)((i, j), (k, l))$ is defined as in the synchronous case. For a non-acting agent i , we define

$$A(t)((i, j), (k, l)) \triangleq \begin{cases} 1 & \text{if } k = i, l = j; \\ 0 & \text{otherwise.} \end{cases} \quad (4.9)$$

The kindness vector is defined as

$$k'(t)((i, j)) \triangleq \begin{cases} (1 - r_i - r'_i)k_i & \text{if } i \text{ is fixed and acting;} \\ 0 & \text{otherwise.} \end{cases}$$

By induction, we obtain $\vec{p}(t) = \prod_{t'=1}^t A(t')\vec{p}(0) + \sum_{\vec{k}' \in K} \left\{ \left(\sum_{l \in S_{\vec{k}'}(t)} \prod_{t'=l}^t A(t') \right) \vec{k}' \right\}$, where K is the set of all possible kindness vectors and $S_{\vec{k}'}(t)$ is a set of the appearance times of \vec{k}' , which are at most t .

We aim to show that $\vec{p}(t)$ converges. First, defining $r_i(M)$ to be the sum of the i th row of M , note [36, Eq. 3], namely

$$r_i(AB) = \sum_{k=1}^n \sum_{j=1}^n a_{i,j} b_{j,k} = \sum_{j=1}^n \sum_{k=1}^n a_{i,j} b_{j,k} = \sum_{j=1}^n a_{i,j} r_j(B). \quad (4.10)$$

Since the sum of every row in any $A(t)$ is at most 1, we conclude that if $B \leq \beta C$, $C_{i,j} \equiv 1$, then also $A(t)B \leq \beta C$.

We now prove that an upper bound of the form βC on the entries of $\prod_{t=p}^s A(t)$ converges to zero geometrically. We have just shown that this bound never increases. First, $A(p) \leq C$, yielding the bound in the beginning. Now, let i be a *fixed* agent, and assume he acts at time t . Thus, each row in $A(t)$ which relates to the edges entering i sums to less than 1, and from (4.10) we gather that the upper bound on the appropriate rows in $A(t)B$ decreases relatively to the bound on B by some constant ratio. Since the graph is connected, for all agents i , $r'_i > 0$, and every agent acts every q times, we will have, after enough multiplications, that the bound on all the entries will have decreased by a constant ratio.

Every agent acts at least once every q times, so we gather that for some $q' > 0$, every q' times, the product of matrices becomes at most a given fraction, say $\alpha \in (0, 1)$, of the product q' times before. This implies a geometric convergence of $\prod_{t'=1}^t A(t')$. As for $\sum_{l \in S_{\vec{k}}(t)} \prod_{t'=l}^t A(t')$, we use the just proven exponential upper bound on the product of $A(t)$ s, and obtain $\sum_{l \in S_{\vec{k}}(t)} \prod_{t'=l}^t A(t') \leq \sum_{l \in S_{\vec{k}}(t)} \alpha^{\lceil \frac{t-l+1}{q'} \rceil} C \leq \sum_{l \in S_{\vec{k}}(t)} \alpha^{\frac{t-l}{q'} - 1} C = \alpha^{\frac{t}{q'} - 1} \left(\sum_{l \in S_{\vec{k}}(t)} \alpha^{\frac{l}{q'}} \right) C \text{ geom. seq. } \frac{\alpha^{\frac{t}{q'} - 1}}{1 - \alpha^{\frac{1}{q'}}} C$, proving a geometric convergence of the series $\sum_{l \in S_{\vec{k}}(t)} \prod_{t'=l}^t A(t')$. Therefore, $\vec{p}(t)$ converges, and it does so geometrically fast. \square

As an immediate conclusion of this proposition, we can finally prove Corollary 3, generalizing Theorem 17 for the case $r_1 + r_2 > 1$.

Corollary 3. *Consider pairwise interaction, where one agent i employs fixed reciprocation and the other agent j employs the floating one, and every agent acts at least once every q times. Assume that $r_i < 1$ and $r_j > 0$. Then, both limits exist and are equal to k_i . The convergence is geometrically fast.*

Proof. If $r_i > 0$, then Proposition 7 implies geometrically fast convergence. We find the limits as in the proof of Theorem 17.

On the other hand, if $r_i = 0$, then agent i constantly acts k_i , and j acts, at any $t \in T_j$, $x_{j,i}(t) = (1 - r_j)x_{j,i}(t-1) + r_j k_i$. We have $x_{j,i}(t) - k_i = (1 - r_j)x_{j,i}(t-1) + r_j k_i - k_i = (1 - r_j)(x_{j,i}(t-1) - k_i)$. Since $r_j > 0$, this converges exponentially to zero, and therefore, $x_{j,i}(t)$ converges exponentially to k_i . \square

Remark 7. *If, unlike the corollary assumes, $r_i = 1$, then, if $r_j < 1$, then we know from the float-float case that both action sequences approach $\frac{r_j}{1+r_j} k_i + \frac{1}{1+r_j} k_j$. If $r_i = r_j = 1$, then in the synchronous case, each agent just repeats what the other one did last time, thereby interchangeably playing k_1 and k_2 . In particular, unless $k_1 = k_2$, then no convergence*

takes place. If the synchronicity breaks by agent i acting alone, then both agents will act k_i from this time on. If, unlike the corollary assumes, $r_j = 0$, then, agent i constantly acts $(1 - r_i)k_i + r_i k_j$ and j constantly acts k_j .

We now turn to finding the limit. We manage to do this in the synchronous case, when all the agents are *floating* or all the *fixed* agents have the same kindness. For all reciprocation attitudes, the following central theorem also provides an alternative proof of convergence in the synchronous case.

Theorem 20. *Given a connected interaction graph, consider the synchronous case where for all agents i , $r'_i > 0$. If there exists a cycle of an odd length in the graph (or at least one agent i employs floating reciprocation and has $r_i + r'_i < 1$), then, for all pairs of agents $i \neq j$ such that $(i, j) \in E$, the limit $L_{i,j}$ exists, and the convergence is geometrically fast. Moreover, if all the agents employ floating reciprocation, then all these limits are equal to each other and it is a convex combination of the kindness values, namely⁷*

$$L = \frac{\sum_{i \in N} \left(\frac{d(i)}{r_i + r'_i} \cdot k_i \right)}{\sum_{i \in N} \left(\frac{d(i)}{r_i + r'_i} \right)}. \quad (4.11)$$

If at least one agent is fixed, then each $L_{i,j}$ is a positive combination of all the kindness values of the agents who are fixed. Moreover, if, all the fixed agents have the same kindness k , then all these limits are equal to k .

In any case, when not all the agents are floating, then changing only the kindness of the floating agents does not change the limits (also follows from the limits being positive combinations of all the kindness values of the agents who are fixed).

Let us say several words about the assumptions. If all the agents are *fixed*, we can prove that the actions are subsequences of the actions in the synchronous case (a straightforward generalization of Lemma 5.) Thus, the synchronous case represents all the cases in the limit, when all the agents are *fixed*. The assumption of a cycle of an odd length virtually always holds, since three people influencing each other form such a cycle.

Equation (4.11) means that the limit is a weighted average of the kindness values, the weight of k_i being the number of i 's neighbors divided by the sum of her reciprocation coefficients. Intuitively, the influence of an agent's kindness is proportional to her influence and inversely proportional to her responsiveness. An obvious conclusion of the theorem is that the *fixed* agents are, intuitively spoken, more important than the *floating* ones, at least their kindness is.

We finally prove Theorem 20. The idea is to express how each action depends on the actions in the previous time in a dynamics matrix A , and prove the theorem by applying the famous Perron–Frobenius theorem [37, Theorem 1.1, 1.2] to this matrix.

Proof. We first prove the case where all the agents use *floating* reciprocation. We now

⁷ $d(i)$ denotes the degree of $i \in N$ in the interaction graph.

define the dynamics matrix $A \in \mathbb{R}_+^{|E| \times |E|}$:

$$A((i, j), (k, l)) \triangleq \begin{cases} (1 - r_i - r'_i) & \text{if } k = i, l = j; \\ r_i + r'_i \frac{1}{|N^+(i)|} & \text{if } k = j, l = i; \\ r'_i \frac{1}{|N^+(i)|} & \text{if } k \neq j, l = i; \\ 0 & \text{otherwise.} \end{cases} \quad (4.12)$$

According to the definition of *floating* reciprocation, if for each time $t \in T$ the column vector $\vec{p}(t) \in \mathbb{R}_+^{|E|}$ describes the actions at time t , in the sense that its (i, j) th coordinate contains $x_{i,j}(t)$ (for $(i, j) \in E$), then $\vec{p}(t+1) = A\vec{p}(t)$. We then call $\vec{p}(t)$ an action vector. Initially, $\vec{p}(i,j)(0) = k_i$.

Further, we shall need to use the Perron–Frobenius theorem for primitive matrices. We now prepare to use it, and first we show that A is primitive. First, A is irreducible since we can move from any $(i, j) \in E$ to any $(k, l) \in E$ as follows. We can move from an action to its reverse, since if $k = j, l = i$, then $A((i, j), (k, l)) = r_i + r'_i \frac{1}{|N^+(i)|} > 0$. We can also move from an action to another action by the same agent, since we can move to any action on the same agent and then to its reverse. To move to an action on the same agent, notice that if $l = i$, then $A((i, j), (k, l)) \geq r'_i \frac{1}{|N^+(i)|} > 0$. Now, we can move from any action (i, j) to any other action (k, l) by moving to the reverse action (j, i) (if $k = j, l = i$, we are done). Then, follow a path from j to k in graph G by moving to the appropriate action by an agent and then to the reverse, as many times as needed till we are at the action (k, j) and finally to the action (k, l) . Thus, A is irreducible.

By definition, A is non-negative. A is aperiodic, since either at least one agent i has $r_i + r'_i < 1$, and thus the diagonal contains non-zero elements, or there exists a cycle of an odd length in the interaction graph G . In the latter case, let the cycle be i_1, i_2, \dots, i_p for an odd p . Consider the following cycles on the index set of the matrix: $(i, j), (j, i), (i, j)$ for any $(i, j) \in E$ and $(i_2, i_1), (i_3, i_2), \dots, (i_p, i_{p-1}), (i_1, i_p), (i_2, i_1)$. Their lengths are 2 and p , respectively, which greatest common divisor is 1, implying aperiodicity. Being irreducible and aperiodic, A is primitive by [37, Theorem 1.4]. Since the sum of every row is 1, the spectral radius of A is 1.

According to the Perron–Frobenius theorem for primitive matrices [37, Theorem 1.1], the absolute values of all the eigenvalues except one eigenvalue of 1 are strictly less than 1. The eigenvalue 1 has unique right and left eigenvectors, up to a constant factor. Both these eigenvectors are strictly positive. Therefore, [37, Theorem 1.2] implies that $\lim_{t \rightarrow \infty} A^t = \vec{1} \vec{v}'$, where \vec{v}' is the left eigenvector of the value 1, normalized such that $\vec{v}' \vec{1} = 1$, and the approach rate is geometric. Therefore, we obtain $\lim_{t \rightarrow \infty} \vec{p}(t) = \lim_{t \rightarrow \infty} A^t \vec{p}(0) = \vec{1} \vec{v}' \vec{p}(0) = \vec{1} \sum_{(i,j) \in E} v'((i, j)) k_i$. Thus, actions converge to $\vec{1}$ times $\sum_{(i,j) \in E} v'((i, j)) k_i$.

To find this limit, consider the vector v' defined by $v'((i, j)) = \frac{1}{r_i + r'_i}$. Substitution shows it is a left eigenvector of A . To normalize it such that $\vec{v}' \vec{1} = 1$, divide this vector by the sum of its coordinates, which is $\sum_{i \in N} \frac{d(i)}{r_i + r'_i}$, obtaining $v'((i, j)) = \frac{1}{\sum_{i \in N} \frac{d(i)}{r_i + r'_i}} \cdot \frac{1}{r_i + r'_i}$.

Therefore, the common limit is $\frac{\sum_{i \in N} \left(\frac{d(i)}{r_i + r'_i} \cdot k_i \right)}{\sum_{i \in N} \left(\frac{d(i)}{r_i + r'_i} \right)}$.

We now prove the case where at least one agent employs *fixed* reciprocation. We define the dynamics matrix A analogously to the previous case, besides that the first line from (4.12) is missing for the *fixed* agents, since for them, own behavior does not matter. In this case, we have $\vec{p}(t+1) = A\vec{p}(t) + \vec{k}'$, where \vec{k}' is the relevant kindness vector, formally defined as

$$k'((i, j)) \triangleq \begin{cases} (1 - r_i - r'_i)k_i & \text{if } i \text{ is } \textit{fixed}; \\ 0 & \text{otherwise.} \end{cases}$$

By induction, we obtain $\vec{p}(t) = A^t\vec{p}(0) + (\sum_{l=0}^{t-1} A^l)\vec{k}'$.

Analogously to the previous case, A is irreducible, non-negative and aperiodic. Therefore A is primitive. Since at least one agent employs *fixed* reciprocation, at least one row of A sums to less than 1, and therefore the spectral radius of A is strictly less than 1.

Now, the Perron–Frobenius implies that all the eigenvalues are strictly smaller than 1. Since we have $\lim_{t \rightarrow \infty} \vec{p}(t) = \lim_{t \rightarrow \infty} A^t\vec{p}(0) + (\lim_{t \rightarrow \infty} \sum_{l=0}^{t-1} A^l)\vec{k}'$, [37, Theorem 1.2] implies that this limit exists (the first part converges to zero, while the second one is a series of geometrically decreasing elements.) Since A is primitive, $(\lim_{t \rightarrow \infty} \sum_{l=0}^{t-1} A^l) > 0$.

When all the fixed agents have the same kindness k , we now find the limits. Taking the limits in the equality $\vec{p}(t+1) = A\vec{p}(t) + \vec{k}'$ yields $(I - A)\lim_{t \rightarrow \infty} \vec{p}(t) = \vec{k}'$. [37, Lemma B.1] implies that $I - A$ is invertible and therefore, if we guess a vector \vec{x} that fulfills $(I - A)\vec{x} = \vec{k}'$, it will be the limit. Since the vector with all actions equal to k satisfies this equation, we conclude that all the limits are equal to k . In any case, when there exists at least one *fixed* agent, changing only the kindness of the *floating* agents will not change the (unique) solution of $(I - A)\vec{x} = \vec{k}'$, and, therefore, will not change the limits. \square

This theorem in particular states that when all the agents are *floating*, the limit does not depend on who acts on whom, but only on how many one acts. This implies that when a manager thinks about how a group of co-workers will work together, she should only care about the numbers of connections, ignoring who is connected to whom.

Remark 8. *The proof of the case of at least one fixed agent easily extends to the combined reciprocation attitude, defined as follows.*

Definition 14. *In the combined reciprocation attitude, agent i is characterized by the parameters k_i, r_i, r'_i , and r''_i , such that $r_i + r'_i + r''_i \leq 1$. Agent i 's action on another agent j is determined by the other agent's action weighted by r_i , by the total action of the neighbors weighted by r'_i and divided by the number of the neighbors, by the agent's own last action weighted by r''_i , and by her kindness, weighted by $1 - r_i - r'_i - r''_i$. That is, for $t \in T_i$,*

$$x_{i,j}(t) \triangleq (1 - r_i - r'_i - r''_i) \cdot k_i + r''_i \cdot x_{i,j}(t-1) + r_i \cdot x_{j,i}(t-1) + r'_i \cdot \frac{\text{got}_i(t-1)}{|\mathbf{N}(i)|}.$$

The extension of the convergence proof is straight-forward.

In order to understand the implications of the limit of actions in practically important cases, let us consider several examples of (4.11).

Example 7. If the interaction graph is regular, meaning that all the degrees are equal to each other, we have $L = \frac{\sum_{i \in N} \left(\frac{k_i}{r_i + r'_i} \right)}{\sum_{i \in N} \left(\frac{1}{r_i + r'_i} \right)}$. This holds for cliques, modeling small human collectives or groups of countries, and for cycles, modeling circular computer networks. Here, only reciprocation coefficients determine, who influences the limit more. The number of connections does not change the limit, as long as this number is the same for everyone.

Example 8. For star networks, modeling networks of a supervisor of several people or entities, assume w.l.o.g. that agent 1 is the center, and we obtain $L = \frac{\frac{n-1}{r_1+r'_1} \cdot k_1 + \sum_{i \in N \setminus \{1\}} \left(\frac{k_i}{r_i+r'_i} \right)}{\frac{n-1}{r_1+r'_1} + \sum_{i \in N \setminus \{1\}} \left(\frac{1}{r_i+r'_i} \right)}$.

In determining the limit, the central agent has $n - 1$ times more weight than any other agent does.

We now conclude about the optimal reciprocation, which goes back to providing decision support. This topic is discussed much more in the following chapter.

Proposition 8. If (4.11) holds, then agent i who wants to maximize the common value L , and who can choose either r_i or r'_i , in certain limits $[a, b]$, for $a > 0$, should choose either the smallest possible or the largest possible coefficient, as follows. We assume we choose r_i , but the same holds for r'_i with the obvious adjustments. She should set r_i to b , if $\sum_{j \in N \setminus \{i\}} \left(\frac{d(j)}{r_j+r'_j} \cdot k_j \right) - k_i \left(\sum_{j \in N \setminus \{i\}} \left(\frac{d(j)}{r_j+r'_j} \right) \right)$ is positive, to a , if that is negative, and to an arbitrary value, if zero. When the discriminating expression is not zero, only these choices are optimal.

Proof. Consider the derivative:

$$\begin{aligned} \frac{\partial L}{\partial(r_i)} &= \frac{\frac{-d(i)k_i}{(r_i+r'_i)^2} \left(\sum_{j \in N} \left(\frac{d(j)}{r_j+r'_j} \right) \right) + \sum_{j \in N} \left(\frac{d(j)}{r_j+r'_j} \cdot k_j \right) \frac{d(i)}{(r_i+r'_i)^2}}{\left(\sum_{j \in N} \left(\frac{d(j)}{r_j+r'_j} \right) \right)^2} \\ &= \frac{\frac{d(i)}{(r_i+r'_i)^2} \left(\sum_{j \in N} \left(\frac{d(j)}{r_j+r'_j} \cdot k_j \right) - k_i \left(\sum_{j \in N} \left(\frac{d(j)}{r_j+r'_j} \right) \right) \right)}{\left(\sum_{j \in N} \left(\frac{d(j)}{r_j+r'_j} \right) \right)^2} \\ &= \frac{\frac{d(i)}{(r_i+r'_i)^2} \left(\sum_{j \in N \setminus \{i\}} \left(\frac{d(j)}{r_j+r'_j} \cdot k_j \right) - k_i \left(\sum_{j \in N \setminus \{i\}} \left(\frac{d(j)}{r_j+r'_j} \right) \right) \right)}{\left(\sum_{j \in N} \left(\frac{d(j)}{r_j+r'_j} \right) \right)^2}. \end{aligned}$$

Therefore, the derivative is zero either for all r_i or for none. In any case, the maximum is attained at an endpoint. To avoid complicated substitution, we consider the derivative sign instead:

$$\frac{\partial u_i}{\partial r_i} \geq 0 \iff \sum_{j \in N \setminus \{i\}} \left(\frac{d(j)}{r_j+r'_j} \cdot k_j \right) - k_i \left(\sum_{j \in N \setminus \{i\}} \left(\frac{d(j)}{r_j+r'_j} \right) \right) \geq 0,$$

and so when $\sum_{j \in N \setminus \{i\}} \left(\frac{d(j)}{r_j + r'_j} \cdot k_j \right) - k_i \left(\sum_{j \in N \setminus \{i\}} \left(\frac{d(j)}{r_j + r'_j} \right) \right)$ is nonnegative, i should choose the largest r_i , which is b , and she should choose $r_i = a$ otherwise. When the derivative is not zero, these choices are the only optimal ones. \square

We can also prove a general convergence result, allowing agents to act in a more general way than modeled above. We need the following definition:

Definition 15. Given a metric space (X, d) , function $f: X \rightarrow X$ is called contraction, if for any $x_1, x_2 \in X$, we have $d(f(x_1), f(x_2)) \leq qd(x_1, x_2)$, for some $q \in (0, 1)$.

Theorem 21. Given an interaction graph, assume the synchronous case, where every agent acts in the following way. Let $S \subseteq \mathbb{R}$ be a compact set. Similarly to the proof of Theorem 20, assume that for each time $t \in T$, the column vector $\vec{p}(t) \in S^{|E|}$ describes the actions at time t , in the sense that its (i, j) th⁸ coordinate contains $x_{i,j}(t)$, and that there exists a contraction $f: S^{|E|} \rightarrow S^{|E|}$ with respect to the Euclidean metric, such that $\vec{p}(t+1) = f(\vec{p}(t))$. Initially, $\vec{p}_{(i,j)}(0) = k_i$. Then, for all pairs of agents $i \neq j$ such that $(i, j) \in E$, the limit $L_{i,j}$ exists. The convergence is geometrically fast.

This theorem is not a generalization of Theorem 20, since matrix in A in the proof of Theorem 20 needs not be a contraction.

Proof. By definition of action, $\vec{p}(t) = f^t(\vec{p}(0))$, and using Banach's fixed point theorem [38, Exercise 6.88], we know that $f^t(\vec{p}(0))$ converges to the unique fixed point of f in $S^{|E|}$, with a geometrical speed, thereby proving the theorem. \square

This theorem means that absolutely any interaction that brings the actions closer becomes stable exponentially quickly.

We have studied some properties of interaction theoretically, and we now turn to the not yet understood properties and analyze them using simulations.

6. SIMULATIONS

Since we have managed to find limits for interaction of n agents only if at most one agent is *fixed*, we now analyze the model with simulations. We start by corroborating the theoretical results from the previous sections. Then, we analyze the dependency of the limits of actions on the reciprocation coefficients and test whether Proposition 5 generalizes for any number n of agents. We employ MatLab simulations, running at least 100 synchronous rounds, which allows achieving practical convergence.

We first study the case of three agents who can influence each other, meaning that the interaction graph is a clique. We begin by corroborating the already proven result that when at least one *fixed* agents exists, then the kindness of the *floating* agents does not influence the actions in the limit. This means that changing the kindness leaves the limits without change, and it is indeed what we have observed. Another proven thing we corroborate is that when exactly one *fixed* agent exists, then all the actions approach her kindness as time approaches infinity. Indeed, when the actions are plotted as functions of time, we obtain graphs such as those in Figure 4.7, demonstrating convergence to the

⁸For $(i, j) \in E$.

kindness of the only *fixed* agent. The left graph on that figure demonstrates, that exponential convergence may be quite slow, and this is a new observation we did not know from the theory. We also corroborate that the limiting values of the actions depend linearly on the kindness values of all the fixed agents, the proportionality coefficients being independent of the other kindness values. This means that the limits of actions as functions of kindness look like the graphs in Figure 4.8. In order to reasonably cover the sampling space, all the above-mentioned regularities have also been automatically checked for the combinations of kindness values of 1, 2, 3, 4, 5, over r_i and r'_i values of 0.1, 0.3, 0.5, 0.7, 0.9, such that $r_i + r + i' \leq 1$, and over all the relevant reciprocation attitudes. The checks were up to the absolute precision of 0.01.

We do not know the exact limits when there exist two or more *fixed* agents with distinct kindness values. We do know that the dependencies on the kindness values are linear, but we lack theoretical knowledge about the dependencies of the limits of actions on the reciprocation coefficients, and this is relevant to maximizing the limits of the actions. This find this dependency, we simulate the interaction for various reciprocation coefficients, obtaining monotonic graphs like those in Figure 4.9 and in Figure 4.10. Note that we can have both increasing and decreasing graphs in the same scenario, and also convex and concave graphs. The observed monotonicity has been automatically verified for all the above-mentioned combinations of parameters. This monotonicity means that if an agent wants to maximize the limit of the actions of some agent on some other agent, she can do this by choosing an extreme value of r_i or r'_i . This is in line with the more concrete maximization rules from Proposition 8.

A natural question is whether Proposition 5 can be extended for more than two agents. Since the kindness of the *floating* agents does not effect the limits, this sort of monotonicity with respect to kindness would require all the limits of the actions to be the same. We therefore ask whether the monotonicity holds at least for the actions of the *fixed* agents. The answer is negative, as Figure 4.12 shows.

The next thing we study is regularity in how the degree of a fourth agent influences the limits of the actions. We consider the limits of the actions as functions of the fourth agent's degree, obtaining graphs like those on in Figure 4.11. We find no regularity in these graphs; in particular, no monotonicity holds in the general case.

To summarize, we have studied reciprocation theoretically, completing the gaps in theory with simulations. In the next section, we look back at the model and consider a possible generalization.

7. OPINIONS

When defining the reciprocal reaction, we used the last action of the other agent to model the opinion about the other agent. In this section, we consider a more general modeling of opinions, to obtain a better understanding of our model and to see what extensions are possible to our model.

We can explicitly define the *opinion* of agent i about another agent j at time t , $\text{opin}_{i,j}: \mathbb{R}^{t+1} \rightarrow \mathbb{R}$, as $\text{opin}_{i,j}(t) \triangleq x_{j,i}(t)$, which is the last action of j upon i . Then, we obtain that in the *fixed* reciprocation attitude, for $t > 0$, $x_{i,j}(t) \triangleq (1 - r_i - r'_i) \cdot k_i + r_i \cdot \text{opin}_{i,j}(t-1) + r'_i \cdot \frac{\sum_{j \in N(i)} \text{opin}_{i,j}(t-1)}{|N(i)|}$, and in the *floating* reciprocation attitude, for $t > 0$,

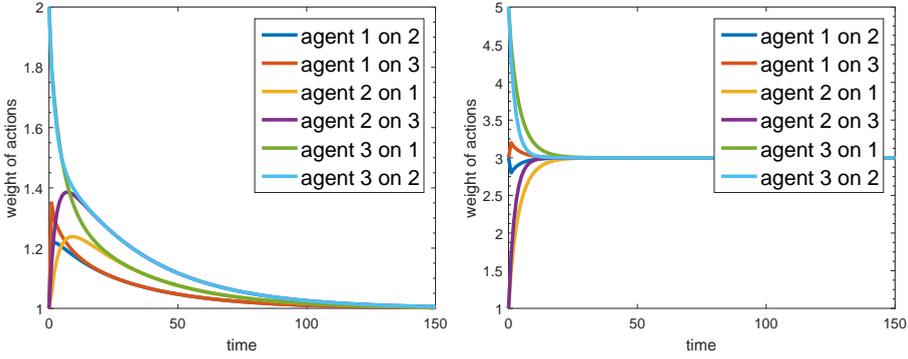


Figure 4.7: Simulation results for the synchronous case, with one *fixed* and two *floating* agents, for $r_1 = 0.1, r_2 = 0.1, r_3 = 0.1, r'_1 = 0.5, r'_2 = 0.1, r'_3 = 0.1$. In the left graph, $k_1 = 1, k_2 = 1, k_3 = 2$, while in the right one, $k_1 = 3, k_2 = 1, k_3 = 5$. The common limits, which are equal to the kindness of agent 1, fit the prediction of Theorem 20.

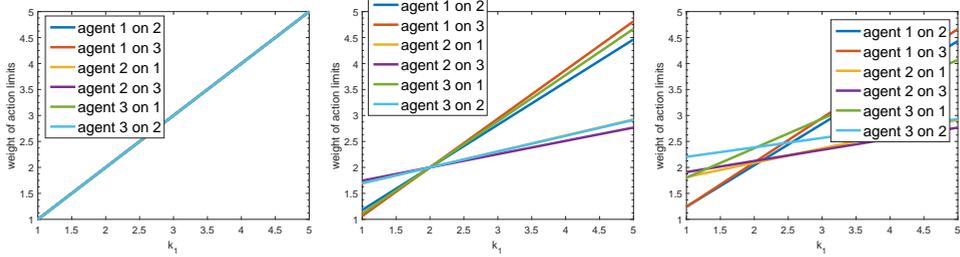


Figure 4.8: Simulation results for the synchronous case, where the limits of actions are plotted as functions of the kindness of a *fixed* agent 1, for $r_1 = 0.2, r_2 = 0.1, r_3 = 0.6, r'_1 = 0.1, r'_2 = 0.4, r'_3 = 0.1, k_2 = 2, k_3 = 3$. In the left graph, agent 1 is the only *fixed* agent, in the middle graph, both 1 and 2 are *fixed*, while in the right one, all the agents are *fixed*. The linear dependencies on the kindness of agent 1 fit the prediction of Theorem 20.

$$x_{i,j}(t) \triangleq (1 - r_i - r'_i) \cdot x_{i,j}(t-1) + r_i \cdot \text{opin}_{i,j}(t-1) + r'_i \cdot \frac{\sum_{j \in N(i)} \text{opin}_{i,j}(t-1)}{|N(i)|}.$$

Naturally, a more general definition of opinion is possible. To this end, we define the temporal distance in T_i , for an $i \in N$, which designates how many times agent i acted between two given times in T_i . Formally,

Definition 16. For an $i \in N$ and two times $t_{i,l}, t_{i,m} \in T_i$, we define $d_{T_i}: T_i^2 \rightarrow \mathbb{R}_+$ by $d_{T_i}(t_{i,l}, t_{i,m}) \triangleq |l - m|$.

Now, we define opinion.

Definition 17. Define the cumulative opinion of i about j at time t to be $\text{opin}_{i,j}(t) \triangleq \sum_{t' \in T_j, t' \leq t} \delta_i(d_{T_j}(t', s_j(t)) + 1) \cdot x_{j,i}(t')$, where $\delta_i(p): \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is the discount function, expressing how much the passed time influences the importance of an action.

The current model fits in this definition as follows.

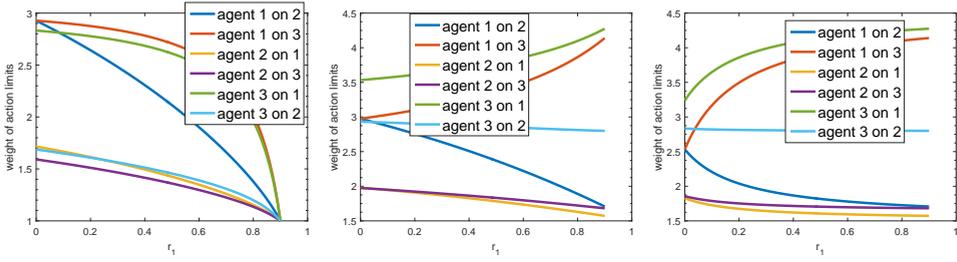


Figure 4.9: Simulation results for the synchronous case, where the limits of actions are plotted as functions of r_1 , for $r_2 = 0.1, r_3 = 0.6, r'_1 = 0.1, r'_2 = 0.4, r'_3 = 0.1, k_1 = 3, k_2 = 1, k_3 = 5$. In the left graph, agent 3 is the only *floating* agent, in the middle graph, all the agents are *fixed*, while in the right one, 1 is the only *floating* agent. All the graphs exhibit monotonicity.

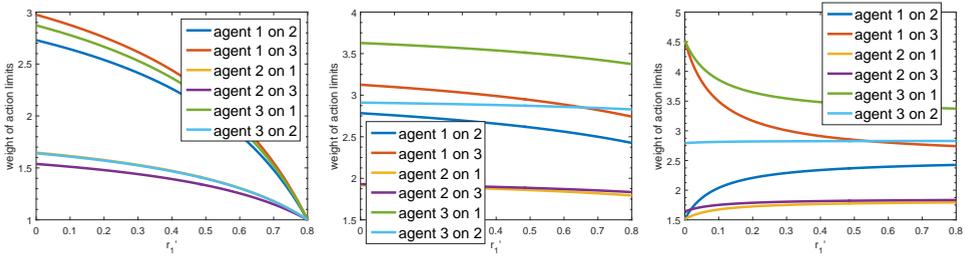


Figure 4.10: Simulation results for the synchronous case, where the limits of actions are plotted as functions of r'_1 , for $r_1 = 0.2, r_2 = 0.1, r_3 = 0.6, r'_2 = 0.4, r'_3 = 0.1, k_1 = 3, k_2 = 1, k_3 = 5$. In the left graph, agent 3 is the only *floating* agent, in the middle graph, all the agents are *fixed*, while in the right one, 1 is the only *floating* agent. All the graphs exhibit monotonicity.

Remark 9. Our definition of opinion as $\text{opin}_{i,j}(t) = x_{j,i}(t)$ is a particular case of this model, where the discount function is $\delta_i(p) = \begin{cases} 1 & p = 1, \\ 0 & \text{otherwise.} \end{cases}$

We now put our work in the context of the related literature.

8. RELATED WORK

We now further motivate our model and place it among works of others. We now describe works that motivate and inspire our model, beyond the main motivation, presented in Section 1 and Section 2.3. Then, we describe other works that use convex combination to model behavior.

In Section 1, we present the main reasons for predicting reciprocation and for modeling it as we do. In addition to the direct motivation for our model, presented in Section 1, we were inspired by Trivers [13] (a psychologist), who describes a balance between an inner quality (immutable kindness) and costs/benefits when determining an action. This idea of balancing the inner and the outer appears also in our model. Trivers also discusses a naturally selected complicated balance between altruistic and cheating tendencies, which we model as kindness, which represents the inherent inclination

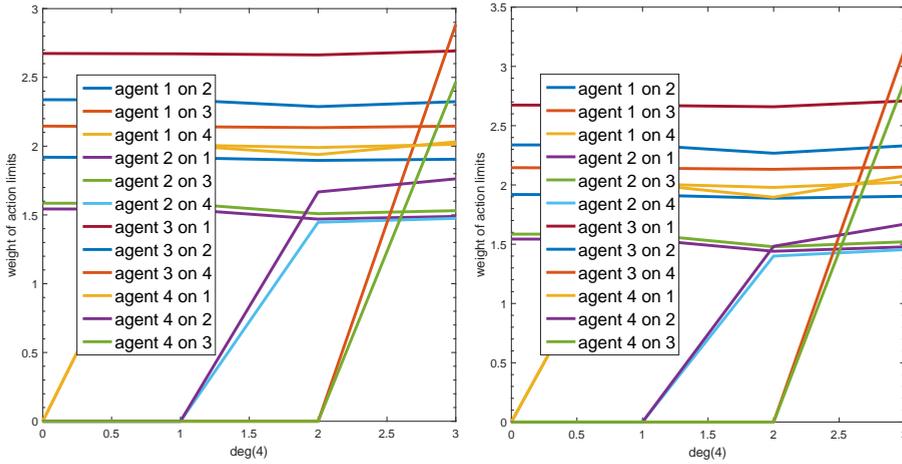


Figure 4.11: Simulation results for the synchronous case, where the limits of actions are plotted as functions of the degree of agent 4, for $r_1 = 0.2, r_2 = 0.1, r_3 = 0.6, r_4 = 0.5, r'_1 = 0.1, r'_2 = 0.4, r'_3 = 0.1, r'_4 = 0.2, k_1 = 2, k_2 = 1, k_3 = 4, k_4 = 2$. In the left graph, all the agents are *fixed*, while in the right one, 4 is the only *floating* agent. No regularity observed.

to contribute. The balance between complying and not complying is mentioned in the conclusion of [27], motivating the convex combination between own kindness or action and others' actions.

Additional motivation stems from the bargaining and negotiation realm, where Pruitt [39] mentions that in negotiation, cooperation often takes place in the form of reciprocation and that personal traits influence the way of cooperation, which corresponds in our model to the personal kindness and reciprocation coefficients.

The idea of humans behaving according to a convex combination resembles another model, that of the altruistic extension, like [40–42], and Chapter iii.2 in [43]. In these papers, utility is often assumed being a convex combination, while we consider a mechanism of an action being a convex combination.

We have presented and fully analyzed the model, and it remains to conclude.

9. CONCLUSIONS AND FUTURE WORK

This section summarizes the chapter, draw conclusions and presents the possibilities to aim to in the future.

We consider networks of interacting agents. This models ubiquitous situations, such as file sharing networks and interacting colleagues. In order to facilitate behavioral decisions for people and agents owned by people, we need to predict what interaction a given setting will engender. To this end, we model two reciprocation attitudes in this chapter, where a reaction is a weighted combination of the action of the other player, the total action of the neighborhood and either one's own kindness or one's own last action. This combination's weights are defined by the reciprocation coefficients. For a pairwise interaction, we show that actions converge, find the exact limits, and show that if you consider

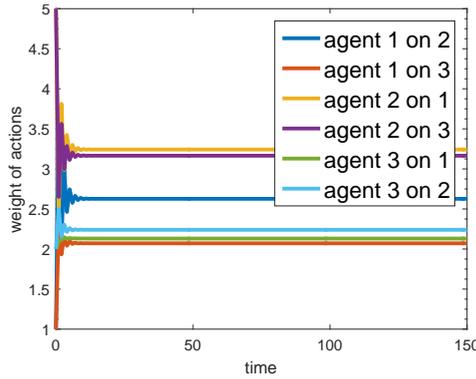


Figure 4.12: Simulation results for the synchronous case, with three *fixed* agents, for $r_1 = 0.5, r_2 = 0.2, r_3 = 0.1, r'_1 = 0.3, r'_2 = 0.5, r'_3 = 0.2, k_1 = 1, k_2 = 5, k_3 = 2$. We observe that $k_1 < k_3$, but $\lim_{t \rightarrow \infty} x_{1,2}(t) > \lim_{t \rightarrow \infty} x_{3,2}(t)$.

your kindness while reciprocating (*fixed*), then, asymptotically, your actions values get closer to your kindness than if you consider it only at the outset. For a general network, we prove convergence and find the common limit if all agents act synchronously and consider their last own action (*floating*), besides at most one agent. Dealing with the case when multiple agents consider their kindness (*fixed*) is mathematically hard, so we use simulations.

To illustrate the implication of our results, consider Example 5.

Example 5 (Continued). *Like in Section 2, assume that all the agents employ floating reciprocation, $n = 3$, and everyone may act on everyone else. Let the kindness values be $k_1 = 0, k_2 = 0.5$ and $k_3 = 1$, and let the reciprocation coefficients be $r_1 = r_2 = 0.5, r'_1 = r'_2 = 0.3, r_3 = 0.8, r'_3 = 0.1$. Then, (4.11) implies that all the actions approach $25/52$ in the limit, meaning that all the colleagues support each other emotionally a lot.*

In addition to predicting the development of reciprocal interactions, our results explain why persistent agents have more influence on the interaction. An expression of the converged behavior is that while growing up, people acquire their own style of reciprocating with acquaintances [24]. In organizations, many styles are often very similar from person to person, forming organizational cultures [25].

We saw in theory and we know from everyday life that the reciprocation process may seem confusing, but the exponential convergence promises the confusion to be short. Actually, we can have a moderately rapid exponential convergence, such as observed in the left graph in Figure 4.7, but mostly, the process converges quickly. Another important conclusion is that employing *floating* reciprocation makes us achieve equality. In the synchronous case, to achieve a common limit it is also enough for all the *fixed* agents to have the same kindness. We also show that if all agents employ *floating* reciprocation and act synchronously, then the influence of an agent is proportional to her number of neighbors and inversely proportional to her tendency to reciprocate, that is, the stability. Therefore, facilitating kinder agents to act on more agents and impeding the less kind

ones from acting on others will increase cooperation. We prove that in the synchronous case, the limit is either a linear combination of the kindness values of all the *fixed* agents or, if all the agents are *floating*, a linear combination of the kindness values of all the agents. Thus, an agent's kindness influences nothing, or it is a linear factor, thereby enabling a very eager agent to influence the limits arbitrarily, by having the *fixed* attitude and the appropriate kindness.

As we see in examples, real situations may require more complex modeling, motivating further research. For instance, the following directions are interesting.

1. Modeling interactions with a known finite time horizon.
2. Merging the two reciprocation attitudes by allowing an agent's action depend on both own kindness and previous own action. It is easy to extend the convergence from Theorem 20 to this case, as we mention in remark 8, but finding the limit is still open.
3. Since people may change while reciprocating, modeling changes in the reciprocity coefficients and/or reciprocation attitude is important. In addition, groups of colleagues and nations get and lose people, motivating modeling a dynamically changing set of reciprocating agents. Even with the same set of agents, the interaction graph may change as people move around.
4. Though it seems extremely hard, it would be nice to consider our model in the light of a game theoretic model of an extensive form game, such as [18].
5. An agent could have different kindness values towards different agents, to represent her prejudgement. Another extension would be allowing the same action be perceived differently by various agents.
6. We used others' research, based on real data, as a basis for the model; actually evaluating the model on relevant data, like the arms race actions, may be enlightening.

We study interaction processes where agents reciprocate with some given parameters, and show that maximizing L would require extreme values of reciprocation coefficients. In the next chapter, we aim to predict real situations better and to provide constructive advice about what parameters and attitudes of the agents are most efficient. To this end, we define utility functions to the agents and consider the game where agents choose their own parameters before the interaction commences. Changing habits is indeed hard, but people are able to change their behavior. The agents' strategic behavior may come at a cost with respect to the social welfare, so considering price of anarchy [44] and stability [45] of such a game is in order. For practical uses, considering how to influence agents to change their behavior is also relevant.

The chapter analyzes reciprocation process analytically and with simulations. This allows estimating whether an interaction will be profitable to a given agent and lays the foundation for further modeling and analysis of reciprocation, in order to anticipate and improve the individual utilities and the social welfare.

REFERENCES

- [1] G. Polevoy, M. de Weerd, and C. Jonker, *The convergence of reciprocation*, in *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '16 (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2016) pp. 1431–1432.
- [2] M. A. Raghunandan and C. A. Subramanian, *Sustaining cooperation on networks: An analytical study based on evolutionary game theory*, in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS'12 (2012) pp. 913–920.
- [3] A. Rubinstein, *Modeling Bounded Rationality*, 1st ed., Vol. 1 (The MIT Press, 1997).
- [4] A. Falk and U. Fischbacher, *A theory of reciprocity*, *Games and Economic Behavior* **54**, 293 (2006).
- [5] E. Fehr and S. Gächter, *Fairness and retaliation: The economics of reciprocity*, *Journal of Economic Perspectives* **14**, 159 (2000).
- [6] W. Güth, R. Schmittberger, and B. Schwarze, *An experimental analysis of ultimatum bargaining*, *Journal of Economic Behavior & Organization* **3**, 367 (1982).
- [7] J. Sobel, *Interdependent preferences and reciprocity*, *Journal of Economic Literature* **43**, 392 (2005).
- [8] R. Axelrod, *The emergence of cooperation among egoists*, *American Political Science Review* **75**, 306 (1981).
- [9] R. Axelrod, *The evolution of cooperation*, Basic books (Basic Books, 1984).
- [10] R. Axelrod and W. Hamilton, *The evolution of cooperation*, *Science* **211**, 1390 (1981), <http://www.sciencemag.org/content/211/4489/1390.full.pdf>.
- [11] J. A. Fletcher and M. Zwick, *Unifying the theories of inclusive fitness and reciprocal altruism*, *The American Naturalist* **168**, 252 (2006).
- [12] S. Van Segbroeck, J. M. Pacheco, T. Lenaerts, and F. C. Santos, *Emergence of fairness in repeated group interactions*, *Phys. Rev. Lett.* **108**, 158104 (2012).
- [13] R. L. Trivers, *The evolution of reciprocal altruism*, *The Quarterly Review of Biology* **46**, 35 (1971).
- [14] H. Gintis, *Game Theory Evolving: A Problem-centered Introduction to Modeling Strategic Behavior*, Economics / Princeton University Press (Princeton University Press, 2000).
- [15] E. Fehr, U. Fischbacher, and S. Gächter, *Strong reciprocity, human cooperation, and the enforcement of social norms*, *Human Nature* **13**, 1 (2002).
- [16] U. Segal and J. Sobel, *Tit for tat: Foundations of preferences for reciprocity in strategic settings*, *Journal of Economic Theory* **136**, 197 (2007).

- [17] J. C. Cox, D. Friedman, and S. Gjerstad, *A tractable model of reciprocity and fairness*, *Games and Economic Behavior* **59**, 17 (2007).
- [18] M. Dufwenberg and G. Kirchsteiger, *A theory of sequential reciprocity*, *Games and Economic Behavior* **47**, 268 (2004).
- [19] M. Rabin, *Incorporating fairness into game theory and economics*, *The American Economic Review* **83**, pp. 1281 (1993).
- [20] M. A. Nowak and K. Sigmund, *Evolution of indirect reciprocity*, *Nature* **437**, 1291 (2005).
- [21] W. J. Dixon, *Reciprocity in United States-Soviet relations: Multiple symmetry or issue linkage?* *American Journal of Political Science* **30**, pp. 421 (1986).
- [22] M. D. Ward, *Modeling the USA-USSR arms race*, *Transactions of The Society for Modeling and Simulation International* **43**, 196 (1984).
- [23] J. Gottman, C. Swanson, and J. Murray, *The mathematics of marital conflict: Dynamic mathematical nonlinear modeling of newlywed marital interaction*, *Journal of Family Psychology* **13**, 3 (1999).
- [24] B. W. Roberts, K. E. Walton, and W. Viechtbauer, *Patterns of mean-level change in personality traits across the life course: a meta-analysis of longitudinal studies*. *Psychological bulletin* **132**, 1 (2006).
- [25] G. Hofstede, *Culture and organizations*, *International Studies of Management & Organization* **10**, 15 (1980).
- [26] P. Breitman and C. Hatch, *How to Say No Without Feeling Guilty: And Say Yes to More Time, More Joy, and what Matters Most to You* (Broadway Books, 2000).
- [27] W. Ury, *The Power of a Positive No: How to Say No and Still Get to Yes* (Random House Publishing Group, 2007).
- [28] H. McAllister, *Who You Are Is What You Do: Making Choices About Life After School*, illustrated ed. (Wilkins Farago Pty Ltd, 2013).
- [29] V. Blondel, J. Hendrickx, A. Olshevsky, and J. Tsitsiklis, *Convergence in multiagent coordination, consensus, and flocking*, in *Decision and Control, 2005 and 2005 European Control Conference. CDC-ECC '05. 44th IEEE Conference on* (2005) pp. 2996–3000.
- [30] L. Moreau, *Stability of multiagent systems with time-dependent communication links*, *Automatic Control, IEEE Transactions on* **50**, 169 (2005).
- [31] J. Tsitsiklis, D. Bertsekas, and M. Athans, *Distributed asynchronous deterministic and stochastic gradient optimization algorithms*, *Automatic Control, IEEE Transactions on* **31**, 803 (1986).

- [32] M. H. DeGroot, *Reaching a consensus*, *Journal of the American Statistical Association* **69**, 118 (1974).
- [33] D. Bindel, J. Kleinberg, and S. Oren, *How bad is forming your own opinion?* in *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on* (2011) pp. 57–66.
- [34] J. S. Rosenschein and G. Zlotkin, *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers* (MIT Press, Cambridge, MA, USA, 1994).
- [35] A. Rubinstein, *Perfect equilibrium in a bargaining model*, *Econometrica* **50**, 97 (1982).
- [36] B. K. Butler and P. H. Siegel, *Sharp bounds on the spectral radius of nonnegative matrices and digraphs*, *Linear Algebra and its Applications* **439**, 1468 (2013).
- [37] E. Seneta, *Non-negative Matrices and Markov Chains*, Springer Series in Statistics (Springer, 2006).
- [38] E. Hewitt and K. Stromberg, *Real and Abstract Analysis: A Modern Treatment of the Theory of Functions of a Real Variable*, Graduate Texts in Mathematics (Springer New York, 1975).
- [39] D. Pruitt, *Negotiation Behavior*, Library and Information Science (Academic Press, 1981).
- [40] P.-A. Chen, B. de Keijzer, D. Kempe, and G. Schäfer, *The robust price of anarchy of altruistic games*, in *Internet and Network Economics*, Lecture Notes in Computer Science, Vol. 7090, edited by N. Chen, E. Elkind, and E. Koutsoupias (Springer Berlin Heidelberg, 2011) pp. 383–390.
- [41] M. Hoefer and A. Skopalik, *Altruism in atomic congestion games*, *ACM Trans. Econ. Comput.* **1**, 21:1 (2013).
- [42] M. Rahn and G. Schäfer, *Bounding the inefficiency of altruism through social contribution games*, in *Web and Internet Economics*, Lecture Notes in Computer Science, Vol. 8289, edited by Y. Chen and N. Immorlica (Springer Berlin Heidelberg, 2013) pp. 391–404.
- [43] J. O. Ledyard, *Public Goods: A Survey of Experimental Research*, Public Economics 9405003 (EconWPA, 1994).
- [44] E. Koutsoupias and C. Papadimitriou, *Worst-case equilibria*, in *16th Annual Symposium on Theoretical Aspects of Computer Science* (Trier, Germany, 1999) pp. 404–413.
- [45] E. Anshelevich, A. DasGupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden, *The price of stability for network design with fair cost allocation*, in *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on* (2004) pp. 295–304.

5

THE GAME OF RECIPROCATION HABITS

An eye for an eye will make the whole world blind.

Graham , 1914

People often act on reciprocal habits, almost automatically responding to others' actions. A robot who interacts with humans may also reciprocate, in order to come across natural and to be predictable. We aim to facilitate a decision support system that advises on utility-efficient habits in these ubiquitous interactions. To this end, given a model for reciprocation behavior with parameters that represent habits, we define a game that describes what habit one should adopt to increase the utility of the process. The used model specifies an agent's action as a weighted combination of the others' previous actions (reacting) and either i) her innate kindness, or ii) her own previous action (inertia). We analyze reciprocation attitude change only for a pairwise interaction, and the coefficient change for any number of agents. For the case of two agents, to analyze what happens when everyone reciprocates rationally, we define a game where an agent may choose her habit, which is either her reciprocation attitude (i or ii), or both her reciprocation attitude and weight. For a general connected network, when all agents have attitude ii), we define a game where an agent chooses her weights. We characterize the Nash equilibria of these games and consider their efficiency. We find that the less kind agents should adjust to the kinder agents to improve both their own utility as well as the social welfare. This constitutes advice on improving cooperation and explains real life phenomena in human interaction, such as the societal benefits from adopting the behavior of the kindest person, or becoming more polite as one grows up.

This chapter is an extended version of paper [1].

1. INTRODUCTION

In the previous chapter, we started studying reciprocal interactions. We modeled it as a predefined process and predicted its development. This chapter models strategic choices an agent can take regarding her *own habits* before such an interaction commences. By habits we mean parameters of agent's behavior. Since people tend to reciprocate on habits, a way to optimize one's behavior is choosing one's own habits.

We now review the previous knowledge of reciprocation and then present the new work. Interaction is central in human behavior, e.g., at school, in file sharing over networks, and in business cooperation. While interacting, people tend to reciprocate, i.e., react on the past actions of others [2–4]. Imagine software agents owned by individuals repeatedly competing with the same people online. People expect reciprocal behavior and tend to behave so themselves. Consider virtual assistants. They need to be reciprocal in order to be credible. Think of countries at an arms race or about arguing friends. They also tend to be nicer if the other side is nicer [5–7]. In these and other cases of repeated interaction, we can help people and artificial agents obtain more from the interaction by providing decision support. The decision is how to reciprocate. Reciprocating efficiently includes defining to one's software agent or other artificial agents how to reciprocate with humans. In order to help people strategically choose efficient approaches for reciprocating, and to predict that strategic choice of how to reciprocate, a model is needed that is theoretically approachable and has enough predictive power.

We now elaborate on some examples, starting with an arms race.

Example 9. Consider n countries $1, 2, \dots, n$; each country can put a certain arsenal of weapons at the border with its neighbors, or point some amount of missiles at her potential enemies. What a country approximately does with respect to another country at a given year is what was done in the previous year, adjusted to react to what the other countries did. If they armed themselves against us, we also will, and if the others aimed at us less, so shall we. This process is often reciprocal with linear reactions [5, 6]. Perhaps, one reason for that is that politicians can explain a reciprocal action as a proper reaction to the nation. A crucial question is how to make this process efficient, so that one's country, and, preferably, everyone incurs the least possible cost.

Till now in this example, an action had a negative influence on the other country. We can also consider a positive influence on the other side in this context; for instance, a concession.

Software agents can reciprocate automatically.

Example 10. Consider software agents that run on computers in a cloud, and they need to agree on how much resources each is allocated. Since their owners may want to be nice to others reciprocally, it is reasonable to make them reciprocate. Naturally, everyone wants her agent to reciprocate as efficiently as possible, and also the society can save much money by efficient reciprocation.

Companies can reciprocate while achieving mutual gain.

Example 11. Reciprocation is useful in business life [8]. Reciprocating means helping the other, for example, by redirecting potential clients to another company. It is definitely economically important to make this reciprocation efficient.

The existent studies of reciprocation (sometimes repeated) either attempt to explain why reciprocation is there in the first place [9–12], or, given that reciprocation exists, they analyze what happens in a short interaction where being reciprocal pays off [2, 13, 14]. We, on the other hand, consider a lengthy interaction, that is (naturally) bound to be reciprocal, but changing the approach of reciprocation is possible, in order to receive more and do less.

To study such interactions, we employ the model from the previous chapter, which formally modeled and analyzed repeated intrinsic reciprocation, to understand how reciprocity makes interaction evolve with time. We briefly summarize the model. Actions, which are influences of an agent on another one, are represented by *weight*, where a higher value means a more desirable contribution to its recipient. That model was mainly inspired by arms race models [5, 6] and a model of spouses arguments [7]. Given the model, the previous chapter analyzes the interaction it engenders. This model consists of two reciprocation attitudes, where the action of an agent is a convex combination¹ between i) one's own kindness or ii) one's own last action (mental inertia), and the other's last action (reaction). The combination is determined by the agent's reciprocation coefficient. Since the last own action is, recursively, a product of previous actions, it represents the agent at a given time, including her history. Attitude i), which is connected to kindness, is called *fixed*, and ii) depending on one's own last action is called *floating*. We concentrate, for technical reasons, on two agents, like two rival superpowers.

Such a reciprocation process converges, and in many cases, the actions in the limit are known from the previous chapter; the required previous results are summarized in the next section. A natural question to ask next in order to provide decision support and predict the strategic reciprocation is in what way the agents can strategically influence the reciprocation process for their own good, and what will the social welfare become when every individual behaves strategically. Setting one's way of reciprocating resembles Mastenbroek's [15, Chapter 14] recommendation to know one's own negotiating style and adjust it. Assuming that people strategically choose each action is unrealistic, since people usually act on habits [16], and a strategic choice consists of choosing a habit for the reciprocal interaction. Here, the habit, chosen after deliberation, can be the balance between reacting and being faithful to oneself, as defined in the model. It is also easy to prescribe a "habit" to a robot.

Choosing habits resembles bounded rationality, especially that of procedures of choice [17, Chapter 2]. Indeed, our agent follows the procedure of rationally choosing among the possible habits. The difference is that choosing a habit does include a rational step, and is, therefore, amenable to a standard game-theoretic analysis, like NE and price of anarchy and stability. Choosing habits resembles metagames as well, when an agent chooses a representative to play the underlying game for her. For instance, Rubinstein [17, Chapter 8] and [18, Chapter 9] define a *machine game*, where an agent wants a well-paying strategy that is simple to implement. This tradeoff is modeled by choosing a finite deterministic automaton to play the repeated game, where the agent's utility increases in the utility of the underlying game and decreases in the number of the states of the chosen automaton. The game's equilibria are studied. The equilibria in this game are found for the case of the utility of the repeated game being defined as the limit-

¹A combination is convex if it has nonnegative weights that sum up to 1.

of-means or with discounting in [19]. Various combinations between the utility of the repeated game and the complexity are studied; for instance, lexicographic utilities are studied in [20] and [21]. A player in a machine game chooses a finite automaton, while our player chooses a habit. Choosing an automaton, however, considers the bounding effect of finiteness and attempt to minimize the automaton's state space, while we simply consider a best possible habit, all habits being equally simple. Therefore, our model neither generalizes theirs nor is our model generalized by theirs. Additionally, no finite automaton is able to model reciprocation, though any degree of approximation is, of course, possible by extending the state space of such an automaton.

To model strategically setting one's habits, we first define the utility of an agent as the value an agent receives for the action by the other agent minus the cost of the action contributed by herself. Then, we consider the one-shot game of setting one's own reciprocation attitude or coefficient, each of which represents a habit.

All the agents choose their reciprocation habits and then the reciprocation process plays itself. Our contributions include a characterization of this game's Nash equilibria (NE) and a discussion of their efficiencies. We consider only pure NE in this chapter. Analyzing this game provides an insight into how people and machines could change their behavior to achieve a more desirable behavior in the limit of the interaction process. This desirability can be to themselves or to the society. In addition to predicting the strategically reciprocal behavior and advising on what to do, the analysis explains the following known phenomena. First, in reciprocation, we often notice that when the example of the kindest person is followed by others, it makes the group more successful [22]. We also notice that people generally tend to become more polite as they grow up [23], which is yet another example of the utility of learning from the behavior of the kindest.

To make this chapter reasonably self-contained, Section 2 provides other necessary background about the previous chapter. Please skip this section if you have read that chapter. Using our definition of utility from Section 3, we analyze setting one's own reciprocation attitude or coefficient to maximize own utility from the reciprocation. We analyze changing reciprocation attitude for a pairwise interaction. Pairwise interactions still allow for many agents provided assuming that the agents do not mix one relationship with the other ones. Therefore, we first consider the case of pairwise interaction, and the game of choosing the reciprocation attitude in sections 3, 4 and 5, proving the central Theorems 22 and 23. We also model in Section 6 what happens if an agent can choose both own attitude and reciprocation coefficient. The answers are given in the key Theorems 24 and 25.

Then, we consider the general case of agents interacting with many agents. Here, we are able to analyze the game of choosing the balance between being faithful to her inner self and reacting. We consider the best reaction of an agent to a given choice of the other agents in Section 8. In order to analyze the general situation, when everyone reciprocates optimally, we consider the game of simultaneously choosing the balance and the Nash equilibria of this game in Section 10, summarizing our findings in Theorem 26. To study whether the various Nash equilibria are good or bad for the society, we calculate their efficiency relatively to the optimum possible situations, which requires finding the optimum total utility in Section 9.

2. BACKGROUND

For the self-containment of this chapter, we restate some definitions and results from the previous chapter. We begin by presenting the model. Then, we mention the convergence results for pairwise interaction, and finally, provide a theorem for n agents. The numeration of the statements copies the original one from Chapter 4.

Recall that an agent $i \in N$ acts at the times in $T_i \subseteq \{0, 1, 2, \dots\}$. Agent i has kindness k_i , defining her inner inclination to act on others, and reciprocation coefficients $r_i, r'_i \in [0, 1]$, defining what fraction of her action is reciprocal. Therefore, $r_i + r'_i \leq 1$. W.l.o.g., $k_n \geq \dots \geq k_2 \geq k_1$. We define the total received action by $\text{got}_i(t) \triangleq \sum_{j \in N(i)} x_{j,i}(t)$. There are two reciprocation attitudes, as follows. In both cases $x_{i,j}(0) \triangleq k_i$.

Definition 12. For the fixed reciprocation attitude, agent i 's reaction on the other agent j and on the neighborhood is determined by the agent's kindness weighted by $1 - r_i - r'_i$, by the other agent's action weighted by r_i and by the total action of the neighbors weighted by r'_i and divided by the number of the neighbors: That is, for $t \in T_i$,

$$x_{i,j}(t) \triangleq (1 - r_i - r'_i) \cdot k_i + r_i \cdot x_{j,i}(t-1) + r'_i \cdot \frac{\text{got}_i(t-1)}{|N(i)|}.$$

Definition 13. In the floating reciprocation attitude, agent i 's action is a weighted average of her own last action, of that of the other agent j and of the total action of the neighbors divided by the number of the neighbors: To be precise, for $t \in T_i$,

$$x_{i,j}(t) \triangleq (1 - r_i - r'_i) \cdot x_{i,j}(t-1) + r_i \cdot x_{j,i}(t-1) + r'_i \cdot \frac{\text{got}_i(t-1)}{|N(i)|}.$$

We prove the following convergence theorems, representing what takes place once the actions have stabilized. For two agents, we assume, w.l.o.g., that $r'_i = 0$.

For two *fixed* agents, we prove:

Theorem 15. If the reciprocation coefficients are not both 1, which means $r_1 r_2 < 1$, then we have, for $i \in N$: $\lim_{p \rightarrow \infty} x_{i,j}(t_{i,p}) = \frac{(1-r_i)k_i + r_i(1-r_j)k_j}{1-r_i r_j}$.

For two agents, in the *floating* case, we show:

Theorem 16. If the reciprocation coefficients are neither both 0 and nor both 1, which means $0 < r_1 + r_2 < 2$, then, as $t \rightarrow \infty$, $x(t)$ and $y(t)$ converge to a common limit. In the synchronous case ($T_1 = T_2 = T$), they both approach

$$\frac{1}{2} \left(k_1 + k_2 + (k_2 - k_1) \frac{r_1 - r_2}{r_1 + r_2} \right) = \frac{r_2}{r_1 + r_2} k_1 + \frac{r_1}{r_1 + r_2} k_2.$$

For a *fixed* and a *floating* agent, the following corollary shows convergence:

Corollary 3. Consider pairwise interaction, where one agent i employs fixed reciprocation and the other agent j employs the floating one, and every agent acts at least once every q times. Assume that $r_i < 1$ and $r_j > 0$. Then, both limits exist and are equal to k_i . The convergence is geometrically fast.

For multiple agents, we prove the following:

Theorem 20. *Given a connected interaction graph, consider the synchronous case where for all agents i , $r'_i > 0$. If there exists a cycle of an odd length in the graph (or at least one agent i employs floating reciprocation and has $r_i + r'_i < 1$), then, for all pairs of agents $i \neq j$ such that $(i, j) \in E$, the limit $L_{i,j}$ exists, and the convergence is geometrically fast. Moreover, if all the agents employ floating reciprocation, then all these limits are equal to each other and it is a convex combination of the kindness values, namely²*

$$L = \frac{\sum_{i \in N} \left(\frac{d(i)}{r_i + r'_i} \cdot k_i \right)}{\sum_{i \in N} \left(\frac{d(i)}{r_i + r'_i} \right)}.$$

If at least one agent is fixed, then each $L_{i,j}$ is a positive combination of all the kindness values of the agents who are fixed. Moreover, if, all the fixed agents have the same kindness k , then all these limits are equal to k .

In any case, when not all the agents are floating, then changing only the kindness of the floating agents does not change the limits (also follows from the limits being positive combinations of all the kindness values of the agents who are fixed).

We next begin by analyzing strategic behavior in pairwise interaction.

TWO AGENTS

We begin with considering the case of a pairwise interaction with agents setting their reciprocation attitude.

3. TWO AGENTS: UTILITY MAXIMIZATION

As a first step to analyzing strategic choices, we now define the utility of an agent and consider how an agent can maximize her utility by choosing either her reciprocation coefficient or reciprocation attitude, before the interaction begins. This can be expected from a rational agent, who reciprocates, but chooses her reciprocation habits. Since in reality the behavioral parameters of others are unknown, choosing an optimal behavior will probably be harder, through trial and error, and the theory predicts the trend of these choices.

3.1. UTILITY DEFINITION FOR n AGENTS

An agent's utility at a given time moment is the action one receives minus the effort incurred by the action one performs. Colloquially, this is what the agent gets minus what she gives. Modeling the effect of own actions on the actor's utility, besides the incurred effort, is an interesting direction to model. This classical way of defining utility is expressed, for instance, in the quasilinear preferences of auction theory [24, Chapter 9.3]. Formally,

² $d(i)$ denotes the degree of $i \in N$ in the interaction graph.

Definition 18. The utility of agent i at moment t , $u_{i,t}: \mathbb{R}^{\deg(i)} \times \mathbb{R}^{\deg(i)} \rightarrow \mathbb{R}$, is defined as

$$u_{i,t}(x_{i,j}(t), x_{j,i}(t)) \triangleq \sum_{j \in N(i)} x_{j,i}(t) - \beta_i \sum_{j \in N(i)} x_{i,j}(t),$$

where β_i is the relative importance of the performed actions for i 's utility. The personal price of acting is higher, equal or lower than of receiving an action, if β is bigger, equal or smaller than 1, respectively.

In particular, for two agents, we denote $x(t) \triangleq x_{1,2}(t)$ and $y(t) \triangleq x_{2,1}(t)$. Thus, agent 1's utility at time t is $y(t) - \beta_1 x(t)$ and 2's utility at time t is $x(t) - \beta_2 y(t)$.

Remark 10 (Signs). We take acting with a minus sign, to account for the effort it takes (unless β_i is negative, which would mean that the agent enjoys making effort). According to this formula, when $\beta_i > 0$, a negative action would suddenly contribute to the utility; we needed to take the absolute value. Instead, we will assume that actions are always non-negative, which is equivalent to all kindness values being non-negative. We still can have negative influence, we have simply mathematically transformed all the original kindness values by adding a sufficiently large number so that they all have become nonnegative.

To model the utility in the long run, we give the following

Definition 19. Define the asymptotic utility, or just the utility, of agent i , $u_i: (\mathbb{R}^{\deg(i)})^\infty \times (\mathbb{R}^{\deg(i)})^\infty \rightarrow \mathbb{R}$, as $u_i(\bigcup_{t=0}^\infty \{x_{i,j}(t), x_{j,i}(t)\}) \triangleq \lim_{t \rightarrow \infty} u_{i,t}(x_{i,j}(t), x_{j,i}(t))$. When the parameters in the parentheses are clear from the context, we may omit them.

Remark 11 (Other definitions). This is the utility we consider in the paper by default. The utility might be defined otherwise, like a discounted sum, though since we have an exponential convergence, it is possible to simplify it to looking at the limit, assuming that the discounting is not extremely quick. A more elaborate discussion of this question appears at Section 11.1, after Proposition 21.

3.2. CHOOSING RECIPROCATION BEHAVIOR FOR TWO AGENTS

In order to analyze strategic choice of reciprocation habits, we consider how an agent can maximize its utility by choosing how to reciprocate. In Example 5 on page 79, this models a colleague changing her behavior to improve her own well-being as a result of a psychologist's advice. In the case of Example 9, this models a country setting a smart foreign policy with respect to arming. First, suppose that the only available option of agent i to modify the process is by setting its reciprocation coefficient r_i . We therefore analyze how i 's utility depends on r_i . Choosing the reciprocation attitude is studied afterwards. In the results of this section, the asymmetry of the agents stems from $k_2 \geq k_1$.

For the *fixed* reciprocation attitude, we prove:

Proposition 9. In the fixed reciprocation attitude, the following holds: If $r_2 < 1$ and agent 1 wants to maximize his utility by choosing his reciprocation coefficient r_1 , then he should

$$\text{set } r_1 \text{ to be } \begin{cases} 1 & \text{if } r_2 > \beta_1, \\ \text{anything} & r_2 = \beta_1, \\ 0 & r_2 < \beta_1. \end{cases}$$

If $r_1 < 1$ and agent 2 wants to maximize his utility by choosing his reciprocation coefficient r_2 , then he should set r_2 to be

$$\begin{cases} 0 & \text{if } r_1 > \beta_2, \\ \text{anything} & r_1 = \beta_2, \\ 1 & r_1 < \beta_2. \end{cases}$$

These choices are the only utility maximizing ones.

The idea of the proof is to express the utility of an agent and differentiate it by her reciprocation coefficient, to find candidates for the extrema. We obtain that the maximum value is attained at an endpoint, and try both endpoints to find a maximum.

Proof. Let us prove for agent 1 choosing r_1 . We first express 1's utility and then maximize it. Since $r_2 < 1$, we have $r_1 r_2 < 1$, and from Theorem 15 on page 85,

$$\begin{aligned} \lim_{t \rightarrow \infty} x(t) &= \frac{(1-r_1)k_1 + r_1(1-r_2)k_2}{1-r_1 r_2}, \\ \lim_{t \rightarrow \infty} y(t) &= \frac{(1-r_2)k_2 + r_2(1-r_1)k_1}{1-r_1 r_2} \\ \Rightarrow u_1 &= \frac{(1-r_2)k_2 + r_2(1-r_1)k_1}{1-r_1 r_2} - \beta_1 \frac{(1-r_1)k_1 + r_1(1-r_2)k_2}{1-r_1 r_2}. \end{aligned}$$

To find a maximum point of this utility as a function of r_1 , we differentiate:

$$\frac{\partial(u_1)}{\partial(r_1)} = \dots = \frac{(r_2 - \beta_1)(1-r_2)}{(1-r_1 r_2)^2} (k_2 - k_1).$$

Therefore, if $r_2 = \beta_1$, then the derivative is zero, and the utility is constant. Otherwise, the maximum is attained only at a single endpoint: at the right endpoint, if the $r_2 > \beta_1$, and at the left endpoint if $r_2 < \beta_1$.

The case of agent 2 choosing r_2 is proven by analogy. \square

For the *floating* reciprocation attitude, we prove:

Proposition 10. *In the floating reciprocation attitude in a synchronous³ reciprocation, the following holds: If $r_2 < 1$ and agent 1 wants to maximize his utility by choosing his reciprocation coefficient r_1 , then he should set r_1 to be*

$$\begin{cases} 1 & \text{if } r_2 > 0 \text{ and } \beta_1 < 1, \\ 0 & \text{if } r_2 > 0 \text{ and } \beta_1 > 1, \\ \text{anything} & \text{if } r_2 > 0 \text{ and } \beta_1 = 1, \\ 0 & \text{if } r_2 = 0 \text{ and } \beta_1 > 0, \\ \text{anything positive} & \text{if } r_2 = 0 \text{ and } \beta_1 < 0, \\ \text{anything} & \text{if } r_2 = 0 \text{ and } \beta_1 = 0. \end{cases}$$

³That is, $T_1 = T_2 = T$, i.e. both agents act at all times.

If $r_1 < 1$ and agent 2 wants to maximize his utility by choosing his reciprocation coefficient r_2 , then he should set r_2 to be

$$\left\{ \begin{array}{ll} 0 & \text{if } r_1 > 0 \text{ and } \beta_2 < 1, \\ 1 & \text{if } r_1 > 0 \text{ and } \beta_2 > 1, \\ \text{anything} & \text{if } r_1 > 0 \text{ and } \beta_2 = 1, \\ \text{anything positive} & \text{if } r_1 = 0 \text{ and } \beta_2 > 0, \\ 0 & \text{if } r_1 = 0 \text{ and } \beta_2 < 0, \\ \text{anything} & \text{if } r_1 = 0 \text{ and } \beta_2 = 0. \end{array} \right.$$

These choices are the only utility maximizing ones.

The idea of the proof is as in the previous proof.

Proof. Let us prove for agent 1. We first express 1's utility and then maximize it. Assume first that $0 < r_2 < 1$, and therefore $0 < r_1 + r_2 < 2$. Therefore, from Theorem 16 on page 87,

$$\begin{aligned} \lim_{t \rightarrow \infty} x(t) &= \lim_{t \rightarrow \infty} y(t) = \frac{1}{2} \left(k_1 + k_2 + (k_2 - k_1) \frac{r_1 - r_2}{r_1 + r_2} \right) \\ \Rightarrow u_1 &= (1 - \beta_1) \frac{1}{2} \left(k_1 + k_2 + (k_2 - k_1) \frac{r_1 - r_2}{r_1 + r_2} \right). \end{aligned}$$

To find a maximum point of this utility as a function of r_1 , we differentiate:

$$\frac{\partial(u_1)}{\partial(r_1)} = \dots = \frac{(1 - \beta_1)(k_2 - k_1)r_2}{(r_1 + r_2)^2}$$

Therefore, the derivative is zero for $1 = \beta_1$, positive for $1 > \beta_1$, and negative for $1 < \beta_1$, and we have proven the proposition for the choice of agent 1, when $r_2 > 0$.

In the case of $r_2 = 0$, notice that $r_1 = r_2 = 0$ results in $u_1 = k_2 - \beta_1 k_1$, while for $r_1 > 0$, Theorem 16 implies that $u_1 = k_2 - \beta_1 k_2$. Consequently, if $\beta_1 > 0$, only $r_1 = 0$ is optimal for agent 1's utility, if β_1 is zero, $u_1 \equiv k_2$ regardless r_1 , and $\beta_1 < 0$ implies that any positive r_1 is optimal.

The case of agent 2 choosing r_2 is proven by analogy. \square

Remark 12 (An intuitive case). For $\beta_1 = \beta_2 = 0$, which is the case when both agents want only to receive more action weight, both results are very intuitive, since the agent with the smaller original kindness should choose to be very reciprocating, while the other agent should choose to be completely non-reciprocating, thereby remaining kind and pulling the other agent to act more.

For synchronous interaction with *fixed* and *floating* attitudes, we prove:

Proposition 11. In the case of synchronous interaction, if agent i employs the fixed and j the floating attitude, the following holds: Given any $r_j < 1$, agent i can maximize her

utility by setting r_i to be

$$\left\{ \begin{array}{ll} 1 & r_j > 0 \text{ and } (1 - \beta_i)k_i \leq (1 - \beta_i)k_j, \\ \text{Any value } < 1 & r_j > 0 \text{ and } (1 - \beta_i)k_i \geq (1 - \beta_i)k_j, \\ 0 & r_j = 0 \text{ and } k_j \geq \beta_i k_i, \\ 1 & r_j = 0 \text{ and } k_j \leq \beta_i k_i. \end{array} \right.$$

If $r_i < 1$, and agent j wants to maximize her utility, she should set r_j to be

$$\begin{cases} 0 & (r_i - \beta_j)(k_j - k_i) \geq 0, \\ \text{Anything positive} & \text{otherwise.} \end{cases}$$

These choices are the only utility maximizing ones.

The idea is to use Corollary 3 from page 94 to find the possible utilities and to compare them.

Proof. If, in addition to being smaller than 1, we also have that $r_j > 0$, then Corollary 3 implies that any $r_i < 1$ makes both limits be k_i . Since $r_i = 1$ allows us assume both agents are *floating*, Theorem 16 on page 87 gives the common limit of $\frac{r_j}{1+r_j} k_i + \frac{1}{1+r_j} k_j$. This is at least as good to i as k_i is if and only if

$$\begin{aligned} (1 - \beta_i)k_i \leq (1 - \beta_i)\left(\frac{r_j}{1+r_j} k_i + \frac{1}{1+r_j} k_j\right) &\iff (1+r_j)(1 - \beta_i)k_i \leq (1 - \beta_i)(r_j k_i + k_j) \\ &\iff (1 - \beta_i)k_i \leq (1 - \beta_i)k_j. \end{aligned}$$

On the other hand, if $r_j = 0$, then i acts $(1 - r_i)k_i + r_i k_j$ and j acts k_j , and we have $u_i = k_j - \beta_i((1 - r_i)k_i + r_i k_j) = (1 - r_i)(k_j - \beta_i k_i)$. Therefore, the sign of $k_j - \beta_i k_i$ determines which r_i is optimal.

Assume $r_i < 1$ now. Then, if $r_j > 0$, both action sequences converge to k_i , by Corollary 3, thus $u_j = (1 - \beta_j)k_i$. If $r_j = 0$, then i constantly acts $(1 - r_i)k_i + r_i k_j$ and j constantly acts k_j , and so $u_j = (1 - r_i)k_i + r_i k_j - \beta_j k_j$. Comparing the utilities, we obtain that the first one is greater than the second one if and only if $(r_i - \beta_j)(k_i - k_j) > 0$. \square

If the kindness values and reciprocation coefficient are set, and an agent may only choose between *fixed* or *floating* reciprocation, we prove:

Proposition 12. *In a synchronous reciprocation, if $0 < r_1, r_2 < 1$, then, if agent 1 wants to maximize her utility, and she may only choose whether to employ fixed or floating reciprocation, then she should choose*

$$\begin{cases} \text{fixed} & \text{if (agent 2 plays fixed} \wedge \{\beta_1 \geq r_2\}) \vee (\text{agent 2 plays floating} \wedge \{\beta_1 \geq 1\}), \\ \text{floating} & \text{if (agent 2 plays fixed} \wedge \{\beta_1 \leq r_2\}) \vee (\text{agent 2 plays floating} \wedge \{\beta_1 \leq 1\}). \end{cases}$$

If agent 2 wants to maximize his utility by choosing fixed or floating reciprocation, then he should choose

$$\begin{cases} \text{floating} & \text{if (agent 1 plays fixed} \wedge \{\beta_2 \geq r_1\}) \vee (\text{agent 1 plays floating} \wedge \{\beta_2 \geq 1\}), \\ \text{fixed} & \text{if (agent 1 plays fixed} \wedge \{\beta_2 \leq r_1\}) \vee (\text{agent 1 plays floating} \wedge \{\beta_2 \leq 1\}). \end{cases}$$

Supposing $k_1 < k_2$, an attitude choice given in this proposition is the only best one if and only if the relevant inequality on the right-hand side of the conditions holds strictly.

The idea of the proof is to compare the possibilities, to see when which option is best.

Proof. First, consider the choice of agent 1. We compare all the reciprocation attitudes of 1, to see which one yields her the larger utility. Assume first that 2 plays *fixed*. Then, 1 playing *fixed* yields $u_1 = \frac{(1-r_2)k_2+r_2(1-r_1)k_1}{1-r_1r_2} - \beta_1 \frac{(1-r_1)k_1+r_1(1-r_2)k_2}{1-r_1r_2}$, and 1 playing *floating* yields $u_1 = (1-\beta_1)k_2$. We compare these utilities now.

$$\begin{aligned} \frac{(1-r_2)k_2+r_2(1-r_1)k_1}{1-r_1r_2} - \beta_1 \frac{(1-r_1)k_1+r_1(1-r_2)k_2}{1-r_1r_2} &\geq (1-\beta_1)k_2 \iff \dots \\ &\iff (1-r_1)(\beta_1-r_2)k_2 \geq (1-r_1)(\beta_1-r_2)k_1. \end{aligned}$$

Since $r_1 < 1 \Rightarrow (1-r_1) > 0$, we conclude that *fixed* is preferable if and only if $\beta_1 - r_2 \geq 0$. The only remaining thing to show about the choice of 1 the case when 2 plays *floating*.

When agent 2 plays *floating*, 1 playing *fixed* yields $u_1 = (1-\beta_1)k_1$, and 1 playing *floating* yields $u_1 = (1-\beta_1)(\alpha k_1 + \beta k_2)$, for some $\alpha, \beta \geq 0, \alpha + \beta = 1$, because the common limit is in $[k_1, k_2]$. Since $r_1 + r_2 \leq 1$, the *floating* reciprocation process is monotonic and since $\min\{r_1, r_2\} > 0$, the agents move from their original acts in the process, so that $\alpha, \beta > 0$. We compare these values now. Since $(\alpha k_1 + \beta k_2) > k_1$, *floating* is preferable if and only if $1 - \beta_1 \geq 0$. The choice of agent 1 has been fully shown.

We prove now for the choice of agent 2. Assume first that agent 1 plays *fixed*. Then, 2 playing *fixed* yields $u_2 = \frac{(1-r_1)k_1+r_1(1-r_2)k_2}{1-r_1r_2} - \beta_2 \frac{(1-r_2)k_2+r_2(1-r_1)k_1}{1-r_1r_2}$, and 2 playing *floating* yields $u_2 = (1-\beta_2)k_1$. We compare these values now.

$$\begin{aligned} \frac{(1-r_1)k_1+r_1(1-r_2)k_2}{1-r_1r_2} - \beta_2 \frac{(1-r_2)k_2+r_2(1-r_1)k_1}{1-r_1r_2} &\geq (1-\beta_2)k_1 \iff \dots \\ &\iff (1-r_2)(r_1-\beta_2)k_2 \geq (1-r_2)(r_1-\beta_2)k_1. \end{aligned}$$

Because $(1-r_2) > 0$, *fixed* is preferable if and only if $(r_1 - \beta_2) \geq 0$. The only remaining thing to show about the choice of agent 2 is the case when 1 plays *floating*.

When 1 plays *floating*, 2 playing *fixed* yields $u_2 = (1-\beta_2)k_2$, and 2 playing *floating* yields $u_2 = (1-\beta_2)(\alpha k_1 + \beta k_2)$, for some $\alpha, \beta \geq 0, \alpha + \beta = 1$. Actually, since $r_1 + r_2 \leq 1$, the *floating* reciprocation process is monotonic and since $\min\{r_1, r_2\} > 0$, we get that the agents move from their original acts, so that $\alpha, \beta > 0$. We compare these values now. Since $(\alpha k_1 + \beta k_2) < k_2$, *fixed* is preferable if and only if $1 - \beta_2 \geq 0$. The choice of agent 2 has been fully shown. \square

Remark 13 (An intuitive case). *For $\beta_1 = \beta_2 = 0$, which is the case when both agents want only to receive more, this result is intuitive, since a less kind agent aligns to the kinder one, and the kinder one lets the other agent align to himself.*

We exemplify the results using Example 9.

Example 9 (Continued). *If countries 1 and 2 have $r_1 = r_2 = 0.5, r_1' = r_2' = 0, \beta_1 = 0, \beta_2 = 0.2$ (acting is cheap), then, whatever attitude 2 employs, 1 should employ floating, to maximize its utility.*

We have prepared the analysis of the game of choosing reciprocation habits. To prepare the ground for analyzing the efficiency of NE, our next step will be finding how the social welfare can be maximized.

4. TWO AGENTS: MAXIMIZING SOCIAL WELFARE

Maximizing the social welfare is relevant for analyzing the whole interaction of agents maximizing their own utilities as a game, to see how good equilibria are for the society relatively to the best possible social welfare. Regardless of the game, the manager (say, the boss of a group of interacting workers) wants to maximize the social welfare by influencing agents' behavior through propaganda or an incentive mechanism.

Before maximizing the social welfare, we formally define it, beginning with the case of n agents.

Definition 20. *The social welfare at time t ($SW_t: \mathbb{R}^{|E|} \rightarrow \mathbb{R}$) is defined as the sum of utilities at time t , i.e.,*

$$SW_t \triangleq \sum_{i \in N} u_{i,t} = \sum_{i \in N, j \in N(i) \setminus \{i\}} (1 - \beta_i) x_{i,j}(t). \quad (5.1)$$

In particular, for two agents,

$$SW_t \triangleq u_{1,t} + u_{2,t} = (1 - \beta_1)x(t) + (1 - \beta_2)y(t). \quad (5.2)$$

For the whole process,

Definition 21. *We define the (asymptotic) social welfare, $SW: (\mathbb{R}^{|E|})^\infty \rightarrow \mathbb{R}$, as $SW \triangleq \lim_{t \rightarrow \infty} SW_t$.*

Let us reconsider Example 5 on page 79 and Example 9.

Example 5 (Continued). *Changing behavioral parameters to increase the social welfare models the boss trying to spread the good practices among the colleagues, by conducting psychological seminars and personal talks.*

Example 9 (Continued). *Changing the behavioral parameters to increase the social welfare models the United Nations trying to spread good practices among countries.*

We now consider optimizing the social welfare for two agents. We first suppose that the only available option to influence the interaction network is through choosing the reciprocation coefficients of the agents, and ask what is the most efficient setup of the r_1, r_2 parameters. Afterwards, we consider setting the reciprocation attitudes. To this end, we now analyze how the asymptotic social welfare depends on these parameters. For given reciprocation attitudes (not necessarily the same attitudes for both agents), we prove

Proposition 13. *We can maximize the social welfare by setting r_1 and r_2 to*

$$\begin{cases} r_1 = 1, r_2 = 0 & \text{if } \max\{\beta_1, \beta_2\} \leq 1, \\ r_1 = 0, r_2 = 1 & \text{if } \min\{\beta_1, \beta_2\} \geq 1, \\ r_1 = r_2 = 0 & \text{if } \beta_1 \geq 1, \beta_2 \leq 1, \\ r_1 = 1, r_2 = 0 & \text{if } \beta_1 \leq 1, \beta_2 \geq 1, \beta_1 + \beta_2 \leq 2, \\ r_1 = 0, r_2 = 1 & \text{if } \beta_1 \leq 1, \beta_2 \geq 1, \beta_1 + \beta_2 \geq 2. \end{cases} \quad (5.3)$$

The idea of the proof is to consider, what limits should be maximized, to maximize the social welfare.

Proof. If $\max\{\beta_1, \beta_2\} \leq 1$, then if we maximize both $\lim_{t \rightarrow \infty} x(t)$ and $\lim_{t \rightarrow \infty} y(t)$, we maximize the social welfare. For $r_1 = 1, r_2 = 0$, we obtain⁴ $\lim_{t \rightarrow \infty} x(t) = k_2$ and $\lim_{t \rightarrow \infty} y(t) = k_2$, which are the maximum possible. Thus, $r_1 = 1, r_2 = 0$ maximizes the social welfare.

If $\min\{\beta_1, \beta_2\} \geq 1$, then if we minimize both $\lim_{t \rightarrow \infty} x(t)$ and $\lim_{t \rightarrow \infty} y(t)$, we maximize the social welfare. For $r_1 = 0, r_2 = 1$, we obtain⁴ $\lim_{t \rightarrow \infty} x(t) = k_1$ and $\lim_{t \rightarrow \infty} y(t) = k_1$, which are the minimum possible. Thus, $r_1 = 0, r_2 = 1$ maximizes the social welfare.

If $\beta_1 \geq 1, \beta_2 \leq 1$, then if we minimize $\lim_{t \rightarrow \infty} x(t)$ and maximize $\lim_{t \rightarrow \infty} y(t)$, we maximize the social welfare. For $r_1 = r_2 = 0$, we obtain⁴ $\lim_{t \rightarrow \infty} x(t) = k_1$ and $\lim_{t \rightarrow \infty} y(t) = k_2$; that is, $\lim_{t \rightarrow \infty} x(t)$ is the minimum possible and $\lim_{t \rightarrow \infty} y(t)$ is the maximum possible. Thus, $r_1 = r_2 = 0$ maximizes the social welfare.

If $\beta_1 \leq 1, \beta_2 \geq 1$, we first express the social welfare in a handier form, and subsequently show how we can maximize it. Denote $\delta \triangleq 1 - \beta_1 \Rightarrow \delta \geq 0$ and $\epsilon \triangleq 2 - \beta_1 - \beta_2$. Then, we have $1 - \beta_2 = -(\delta - \epsilon)$ and $\text{SW} = (1 - \beta_1) \lim_{t \rightarrow \infty} x(t) + (1 - \beta_2) \lim_{t \rightarrow \infty} y(t) = \delta \lim_{t \rightarrow \infty} x(t) - (\delta - \epsilon) \lim_{t \rightarrow \infty} y(t) = \delta(\lim_{t \rightarrow \infty} x(t) - \lim_{t \rightarrow \infty} y(t)) + \epsilon \lim_{t \rightarrow \infty} y(t)$.

Now, if $\beta_1 + \beta_2 \leq 2$, then $\epsilon \geq 0$ and thus, if we maximize $\lim_{t \rightarrow \infty} x(t) - \lim_{t \rightarrow \infty} y(t)$ and $\lim_{t \rightarrow \infty} y(t)$, we maximize the social welfare. For $r_1 = 1, r_2 = 0$, we obtain⁴ $\lim_{t \rightarrow \infty} x(t) = \lim_{t \rightarrow \infty} y(t) = k_2$, thus maximizing the first (since by Proposition 5 on page 94, $\lim_{t \rightarrow \infty} x_{i,j}(t) \leq \lim_{t \rightarrow \infty} x_{j,i}(t)$, the first is non-positive) and the second. Thus, $r_1 = 1, r_2 = 0$ maximizes the social welfare.

Now, if $\beta_1 + \beta_2 \geq 2$, then $\epsilon \leq 0$ and thus, if we maximize $\lim_{t \rightarrow \infty} x(t) - \lim_{t \rightarrow \infty} y(t)$ and minimize $\lim_{t \rightarrow \infty} y(t)$, we maximize the social welfare. For $r_1 = 0, r_2 = 1$, we obtain⁴ $\lim_{t \rightarrow \infty} x(t) = \lim_{t \rightarrow \infty} y(t) = k_1$, thus maximizing the first and minimizing the second. Thus, $r_1 = 0, r_2 = 1$ maximizes the social welfare. \square

Remark 14 (Usage domain). *Notice that this proposition holds also if we may influence both r_1, r_2 and the attitudes of the agents, since the proof maximizes and minimizes expressions for any possible attitudes.*

Consider an intuitive case.

Remark 15 (free actions). *For $\beta_1 = \beta_2 = 0$, this result is intuitive, since the agent with the smaller original kindness is set to be very reciprocating, while the other agent is set to be non-reciprocating, thereby pulling the other agent to act more.*

Suppose now that the reciprocation coefficients are set, and the manager only chooses whether the agents employ *fixed* or *floating* reciprocation.

Proposition 14. *If $0 < r_1, r_2 < 1$ and every agent acts at least once every q times⁵, then the*

⁴This is evident from the definition of *fixed* or *floating* reciprocation, without a convergence theorem.

⁵We do not require synchronicity, since Theorem 16 from page 87 is used without using the value of the common limit.

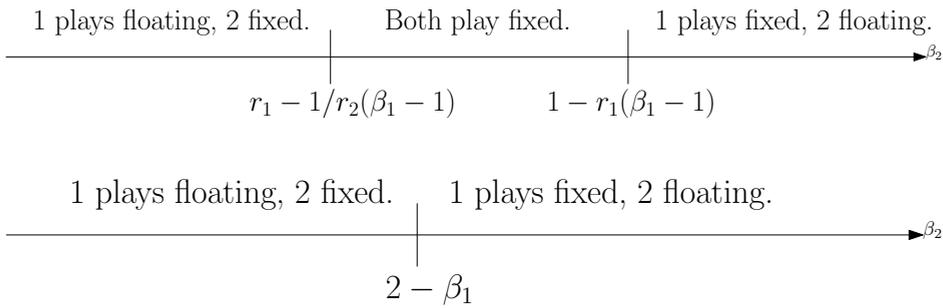


Figure 5.1: The upper figure is for $\beta_1 - 1 \geq 0$, and the lower figure is for $\beta_1 - 1 < 0$. The strategy profile written above denotes a profile to maximize the social welfare, based on where the value of β_2 resides.

social welfare is maximal by reciprocating as follows:

$$\begin{cases} 1 \text{ floating, } 2 \text{ fixed.} & \text{if } \beta_2 \leq 1 - \max\left\{\frac{1}{r_2}(\beta_1 - 1), \beta_1 - 1\right\}, \\ 1 \text{ fixed, } 2 \text{ fixed.} & \text{if } 1 - 1/r_2(\beta_1 - 1) \leq \beta_2 \leq 1 - r_1(\beta_1 - 1), \\ 1 \text{ fixed, } 2 \text{ floating.} & \text{if } \beta_2 \geq 1 - \min\{r_1(\beta_1 - 1), \beta_1 - 1\}. \end{cases}$$

The statement of the proposition can be expressed geometrically. We can maximize the social welfare depending on the real interval where β_2 is: Figure 5.1 shows a profile to maximize the social welfare, based on the segment where the value of β_2 belongs.

The idea of the proof is to express the social welfare in each case, characterize when the social welfare in one case is larger than in the other one, and conclude, how to maximize the social welfare.

Proof. The social welfare in each case is as follows:

$$\begin{aligned} & \text{SW}(1 \text{ plays fixed, } 2 \text{ fixed}) \\ \stackrel{\text{Theorem 15}}{=} & (1 - \beta_1) \frac{(1 - r_1)k_1 + r_1(1 - r_2)k_2}{1 - r_1 r_2} + (1 - \beta_2) \frac{(1 - r_2)k_2 + r_2(1 - r_1)k_1}{1 - r_1 r_2} \\ = \dots = & \frac{(1 - r_1)(1 + r_2 - \beta_1 - \beta_2 r_2)}{1 - r_1 r_2} k_1 + \frac{(1 - r_2)(1 + r_1 - \beta_1 r_1 - \beta_2)}{1 - r_1 r_2} k_2. \end{aligned}$$

$$\text{SW}(1 \text{ plays fixed, } 2 \text{ floating}) \stackrel{\text{Corollary 3}}{=} (1 - \beta_1)k_1 + (1 - \beta_2)k_1 = (2 - \beta_1 - \beta_2)k_1.$$

$$\text{SW}(1 \text{ plays floating, } 2 \text{ fixed}) \stackrel{\text{Corollary 3}}{=} (1 - \beta_1)k_2 + (1 - \beta_2)k_2 = (2 - \beta_1 - \beta_2)k_2.$$

$$\begin{aligned} & \text{SW}(1 \text{ plays floating, } 2 \text{ floating}) \\ \stackrel{\text{Theorem 16}}{=} & (1 - \beta_1)(\alpha k_1 + \beta k_2) + (1 - \beta_2)(\alpha k_1 + \beta k_2) = (2 - \beta_1 - \beta_2)(\alpha k_1 + \beta k_2), \\ & \text{for some } \alpha, \beta \geq 0, \alpha + \beta = 1. \end{aligned}$$

5

We first show that $\text{SW}(1 \text{ plays } \textit{floating}, 2 \textit{ floating})$ is not a candidate for the maximum social welfare. Since $\alpha k_1 + \beta k_2$ is a convex combination of k_1 and k_2 , we have the following: If $(2 - \beta_1 - \beta_2) \geq 0$, then $\text{SW}(1 \text{ plays } \textit{floating}, 2 \textit{ fixed}) \geq \text{SW}(1 \text{ plays } \textit{floating}, 2 \textit{ floating})$, and otherwise $\text{SW}(1 \text{ plays } \textit{fixed}, 2 \textit{ floating}) \geq \text{SW}(1 \text{ plays } \textit{floating}, 2 \textit{ floating})$.

Next, we characterize for each pair of the remaining three candidates, who is bigger in the pair and on what condition.

$$\begin{aligned} \text{SW}(1 \text{ plays } \textit{fixed}, 2 \textit{ fixed}) &\geq \text{SW}(1 \text{ plays } \textit{fixed}, 2 \textit{ floating}) \\ &\iff \dots \iff 1 - \beta_2 \geq r_1(\beta_1 - 1). \\ \text{SW}(1 \text{ plays } \textit{fixed}, 2 \textit{ fixed}) &\geq \text{SW}(1 \text{ plays } \textit{floating}, 2 \textit{ fixed}) \\ &\iff \dots \iff \frac{1}{r_2}(\beta_1 - 1) \geq 1 - \beta_2. \\ \text{SW}(1 \text{ plays } \textit{fixed}, 2 \textit{ floating}) &\geq \text{SW}(1 \text{ plays } \textit{floating}, 2 \textit{ fixed}) \\ &\iff \dots \iff \beta_1 - 1 \geq 1 - \beta_2. \end{aligned}$$

Therefore, if $\frac{1}{r_2}(\beta_1 - 1) \geq 1 - \beta_2 \geq r_1(\beta_1 - 1)$, then $\text{SW}(1 \text{ plays } \textit{fixed}, 2 \textit{ fixed})$ is the largest. If $1 - \beta_2 \leq \min\{r_1(\beta_1 - 1), \beta_1 - 1\}$, then $\text{SW}(1 \text{ plays } \textit{fixed}, 2 \textit{ floating})$ is the largest. If $1 - \beta_2 \geq \max\{\frac{1}{r_2}(\beta_1 - 1), \beta_1 - 1\}$, then $\text{SW}(1 \text{ plays } \textit{floating}, 2 \textit{ fixed})$ is the largest. These conclusions constitute the statement of the proposition. \square

A remarkable case follows.

Remark 16 (Free actions). *For $\beta_1 = \beta_2 = 0$, this result (agent 1 plays floating, 2 fixed) is intuitive, since the less kind agent aligns to the kinder one.*

By now, the preparation for analyzing the whole interaction as a game has been completed, so we proceed to define and to analyze the game.

5. TWO AGENTS: RECIPROCATION ATTITUDE GAME

We have considered an agent choosing her reciprocation coefficient or her *fixed* or *floating* reciprocation attitude, each choice yielding certain (asymptotic) utility to the agent. This situation is naturally modeled as a game where the strategies of each agent are the above choices and the respective utility is the asymptotic utility of the interaction. Recall that the utility of agent i is $\lim_{t \rightarrow \infty} \{x_{j,i}(t) - \beta_i x_{i,j}(t)\}$. This is a one-shot game, the attitude being chosen once, before the interaction commences. Analyzing this game allows predicting the situation, supplying some advice to an external party (such as the boss who wants to influence her employees) or to the agents themselves. As explained after Example 6 from page 83, human agents usually neither completely mimic the others' behavior, nor do they completely ignore it, which means $0 < r_1, r_2 < 1$. For simplicity, we also assume that all agents act synchronously. We call this game the *reciprocation attitude game* (RAG). The central Theorems 22 and 23 summarize our findings about RAG.

We first characterize the existence of pure NE in this game and subsequently look into their efficiency. We assume that $k_2 > k_1$ (strictly) in this section because if the kindness is equal, everyone just keeps acting with this equal value.

Theorem 22. *The Nash equilibria of RAG are characterized as follows:*

$$\begin{aligned}
 (\text{fixed}, \text{fixed}) \text{ is an NE} &\iff \beta_1 \geq r_2 \text{ and } \beta_2 \leq r_1. \\
 (\text{float}, \text{fixed}) \text{ is an NE} &\iff \beta_1 \leq r_2 \text{ and } \beta_2 \leq 1. \\
 (\text{fixed}, \text{float}) \text{ is an NE} &\iff \beta_1 \geq 1 \text{ and } \beta_2 \geq r_1. \\
 (\text{float}, \text{float}) \text{ is an NE} &\iff \beta_1 \leq 1 \text{ and } \beta_2 \geq 1.
 \end{aligned}$$

The proof utilizes Proposition 12 about utility maximization to see when no deviation is profitable.

Proof. Assume that $\beta_1 \geq r_2$ and $\beta_2 \leq r_1$. If the strategy profile is *(fixed, fixed)*, then, according to Proposition 12, no agent will have an incentive to unilaterally deviate, meaning this strategy profile is indeed an NE.

Assume now that *(fixed, fixed)* is an NE. We prove that $\beta_1 \geq r_2$ and $\beta_2 \leq r_1$ by contradiction. If $\beta_1 < r_2$, then Proposition 12 would imply that agent 1 would like to deviate, contradictory to the profile being an NE. If $\beta_2 > r_1$, then Proposition 12 would imply that 2 would like to deviate, contradictory to the NE.

Assume next that $\beta_1 \leq r_2$ and $\beta_2 \leq 1$. If the strategy profile is *(float, fixed)*, then, according to Proposition 12, no agent will have an incentive to unilaterally deviate, meaning this strategy profile is indeed an NE.

Assume now that *(float, fixed)* is an NE. We prove that $\beta_1 \leq r_2$ and $\beta_2 \leq 1$ by contradiction. If $\beta_1 > r_2$, then Proposition 12 would imply that 1 would like to deviate, contradictory to the profile being an NE. If $\beta_2 > 1$, then Proposition 12 would imply that 2 would like to deviate, contradictory to the profile being an NE.

Assume next that $\beta_1 \geq 1$ and $\beta_2 \geq r_1$. If the strategy profile is *(fixed, float)*, then, according to Proposition 12, no agent will have an incentive to unilaterally deviate, meaning this strategy profile is indeed an NE.

Assume now that *(fixed, float)* is an NE. We prove that $\beta_1 \geq 1$ and $\beta_2 \geq r_1$ by contradiction. If $\beta_1 < 1$, then Proposition 12 would imply that 1 would like to deviate, contradictory to the profile being an NE. If $\beta_2 < r_1$, then Proposition 12 would imply that 2 would like to deviate, contradictory to the profile being an NE.

Assume next that $\beta_1 \leq 1$ and $\beta_2 \geq 1$. If the strategy profile is *(float, float)*, then, according to Proposition 12, no agent will have an incentive to unilaterally deviate, meaning this strategy profile is indeed an NE.

Assume now that *(float, float)* is an NE. We prove that $\beta_1 \leq 1$ and $\beta_2 \geq 1$ by contradiction. If $\beta_1 > 1$, then Proposition 12 would imply that 1 would like to deviate, contradictory to the profile being an NE. If $\beta_2 < 1$, then Proposition 12 would imply that 2 would like to deviate, contradictory to the profile being an NE. \square

Remark 17 (Existence of NE). *If no characterizing condition holds, then no NE exists. For example, no characterizing condition holds when $\beta_1 = 0.8, \beta_2 = 0.9, r_1 = 0.5, r_2 = 0.2$, so no pure NE exists in this case. Since the game is finite, a mixed NE always exists by the classical result by Nash [25].*

We now illustrate the theorem for certain parameter values.

Example 12. Let $\beta_1 = 0.3, \beta_2 = 0.6$. Then, Theorem 22 states that

$$\begin{aligned} (\text{fixed}, \text{fixed}) \text{ is an NE} &\iff 0.3 \geq r_2 \text{ and } 0.6 \leq r_1. \\ (\text{float}, \text{fixed}) \text{ is an NE} &\iff 0.3 \leq r_2. \\ (\text{fixed}, \text{float}) \text{ is never an NE.} & \\ (\text{float}, \text{float}) \text{ is never an NE.} & \end{aligned}$$

5.1. PoA AND PoS

The manager or the government may want to know how far the social welfare in an equilibrium is from the maximum possible social welfare. To this end, we consider the famous measures of the efficiency of an equilibrium, namely price of anarchy [26] (PoA) and price of stability [27] (PoS). PoA is the smallest ratio of a social welfare in an NE to the optimum social welfare, and PoS is the largest such ratio.

Using Theorem 22, we know for each given set of parameters what all the NE are. Using Proposition 14, we know for each given set of parameters what the maximum social welfare is. Calculating the social welfare at each of the Nash equilibria and finding its ratio to the optimum social welfare enables us to find the price of anarchy and stability in the following theorem.

Theorem 23. The efficiency of the equilibria is as follows:

Conditions:	PoA = PoS:
$\{1 + r_2 - r_2\beta_2 > \beta_1 > r_2\}$ and $\{\beta_2 < r_1\}$	$\frac{\sum_{i=1,2; j \neq 1} (1-\beta_i) \frac{(1-r_i)k_i + r_i(1-r_j)k_j}{1-r_i r_j}}{(2-\beta_1-\beta_2)k_2}$
$\{1 + 1/r_1 - \beta_2/r_1 > \beta_1 > 1 + r_2 - r_2\beta_2\}$ and $\{\beta_2 < r_1\}$	1
$\{\beta_1 > 1 + 1/r_1 - \beta_2/r_1\}$ and $\{\beta_2 < r_1\}$	$\frac{\sum_{i=1,2; j \neq 1} (1-\beta_i) \frac{(1-r_i)k_i + r_i(1-r_j)k_j}{1-r_i r_j}}{(2-\beta_1-\beta_2)k_1}$
$\{\beta_1 < r_2\}$ and $\{\beta_2 < 1\}$	1
$\{\beta_1 > 1\}$ and $\{1 + r_1 - \beta_1 r_1 > \beta_2 > \max\{1 + \frac{1}{r_2} - \beta_1, r_1\}\}$	$\frac{(2-\beta_1-\beta_2)k_1}{\sum_{i=1,2; j \neq 1} (1-\beta_i) \frac{(1-r_i)k_i + r_i(1-r_j)k_j}{1-r_i r_j}}$
$\{\beta_1 > 1\}$ and $\{\beta_2 > \max\{1 + r_1 - \beta_1 r_1, r_1\}\}$	1
$\{\beta_1 < 1\}$ and $\{2 - \beta_1 > \beta_2 > 1\}$	$\frac{r_2}{r_1+r_2} \frac{k_1}{k_2} + \frac{r_1}{r_1+r_2}$
$\{\beta_1 < 1\}$ and $\{\beta_2 > 2 - \beta_1\}$	$\frac{r_2}{r_1+r_2} + \frac{r_1}{r_1+r_2} \frac{k_2}{k_1}$

In the case of equality in the conditions, the highest entry from our conditions that border the equal value is the price of stability in this case, and the lowest entry is the price or anarchy.

Proof. First, if equality in the conditions holds, then several equilibria exist, and we therefore take the best one for the price of stability and the worst one for the price of anarchy. The proof goes over all the cases from Theorem 22, split further into subcases, if Proposition 14 requires so.

If $\beta_1 > r_2$ and $\beta_2 < r_1$, Theorem 22 implies that the only NE is *(fixed, fixed)*. To find the optimal social welfare, we now look at several cases that fit Proposition 14. If $\beta_1 \leq 1$, then $\beta_1 - 1 \leq 0$, and Proposition 14 splits into 2 cases: $\beta_2 \leq 2 - \beta_1$ and $\beta_2 \geq 2 - \beta_1$. Since $\beta_2 \leq r_1 < 1 \leq 2 - \beta_1$, Proposition 14 implies that *(floating, fixed)* is optimal for the social

welfare. If, on the other hand, $\beta_1 \geq 1$, then $\beta_1 - 1 \geq 0$, and Proposition 14 splits into 3 cases: $\beta_2 \leq 1 + 1/r_2 - (1/r_2)\beta_1$, $1 + r_1 - r_1\beta_1 \geq \beta_2 \geq 1 + 1/r_2 - (1/r_2)\beta_1$, and $1 + r_1 - r_1\beta_1 \leq \beta_2$. In the first case, (floating, fixed) is optimal for the social welfare, and together with the option $\beta_1 \leq 1$ above, this yields the first row in the table of the theorem's statement. If $1 + r_1 - r_1\beta_1 \geq \beta_2 \geq 1 + 1/r_2 - (1/r_2)\beta_1$, then (fixed, fixed) is optimal, yielding the second row. If, finally, $1 + r_1 - r_1\beta_1 \leq \beta_2$ holds, then (fixed, floating) is optimal, and this gives us the third row.

If $\beta_1 < r_2$ and $\beta_2 < 1$, then (float, fixed) is the only NE, according to Theorem 22. Since $\beta_1 - 1 < 0$ and $\beta_2 < 1 < 2 - \beta_1$, Proposition 14 implies that (float, fixed) maximizes the social welfare, and we have proven the fourth row.

If $\beta_1 > 1, \beta_2 > r_1$, then Theorem 22 states (fixed, float) is the only NE. Regarding the social welfare, we look at Proposition 14. We have $\beta_1 - 1 \geq 0$, and we consider 3 cases: $\beta_2 \leq 1 + 1/r_2 - \beta_1/r_2$, $1 + 1/r_2 - \beta_1/r_2 \leq \beta_2 \leq 1 + r_1 - \beta_1 r_1$, and $\beta_2 \geq 1 + r_1 - \beta_1 r_1$. The first case is empty, since $1 + 1/r_2 - \beta_1/r_2 < r_1$. In the second case, (fixed, fixed) is optimal for the social welfare, yielding the fifth row. Finally, the optimum for the third case is obtained at (fixed, floating), and this is also an NE, and we have the sixth row.

If $\beta_1 < 1, \beta_2 > 1$, then (float, float) is the only NE. As for the social welfare, since $\beta_1 - 1 < 1$, Proposition 14 gives us two cases: $\beta_2 \leq 2 - \beta_1$ and $\beta_2 \geq 2 - \beta_1$. In the first case, (floating, fixed) is optimal, and we obtain the penultimate row. On the other hand, if $\beta_2 \geq 2 - \beta_1$, then (fixed, floating) maximizes the social welfare, and we get the last row. \square

In particular, if $\beta_1 < r_2, \beta_2 < 1$, then $PoA = PoS = 1$. We now illustrate the efficiency ranges on Example 12.

Example 12 (Continued). Recall that $\beta_1 = 0.3, \beta_2 = 0.6$. For these values, Theorem 23 implies the following.

Conditions:	Price of anarchy and stability:
$\{0.3 > r_2\}$ and $\{0.6 < r_1\}$	$\frac{\sum_{i=1,2; j \neq i} (1-\beta_i) \frac{(1-r_i)^{k_i+r_j} (1-r_j)^{k_j}}{1-r_i r_j}}{1.1k_2}$
$\{0.3 < r_2\}$	1

Consider Example 5 on page 79.

Example 5 (Continued). If colleagues 1 and 2 have $r_1 = r_2 = 0.5, r'_1 = r'_2 = 0, \beta_1 = 0, \beta_2 = 0.2$ (acting is cheap), then, as just mentioned, $PoA = PoS = 1$ and the only NE is (float, fixed). This is intuitive, since colleague 1 will align to the kinder 2, thereby each colleague maximizes the total action and, since acting is cheap, also her own utility and the social welfare.

5.2. CONVERGING TO NE

To analyze the stability of a Nash equilibrium, we recall the famous best response dynamics [18, Section 2.2], where each agent best responds to the current profile of the others. We prove that given an NE and any profile, we can let each agent simultaneously choose her reciprocation attitude to maximize her utility, such that it ends up in this NE.

Proposition 15. Given a Nash equilibrium in a reciprocation attitude game, for any strategy profile there exists a succession of profiles starting from it and terminating at the given

NE, such that each profile in this succession is an optimum reaction of each agent to the other one's strategy in the previous profile.

The idea of the proof is to use Theorem 22 to obtain information about the parameters, assuming that a given profile is an NE. Then, Proposition 12 is used to leverage the obtained information and show how to begin at any profile and end up at the considered NE.

Proof. We prove the proposition by showing such a succession of profiles for each possible NE and each starting profile. We denote by arrows \rightarrow the transition from one profile to another one, where the next profile is an optimum reaction of each agent to the other one's strategy in the previous profile. When we write a profile, in which we do not mention the strategy of an agent, we write a question mark instead, such as in $(?, \text{fixed})$.

$(\text{fixed}, \text{fixed})$ is an NE $\iff \beta_1 \geq r_2$ and $\beta_2 \leq r_1$, according to Theorem 22. The successions to this NE from various profiles follow:

$$(\text{fixed}, \text{float}) \xrightarrow{\text{Proposition 12}} (?, \text{fixed}) \xrightarrow{\text{Proposition 12}} (\text{fixed}, \text{fixed}).$$

$$(\text{float}, \text{fixed}) \xrightarrow{\text{Proposition 12}} (\text{fixed}, \text{fixed}).$$

$$(\text{float}, \text{float}) \xrightarrow{\text{Proposition 12}} (?, \text{fixed}) \xrightarrow{\text{as above}} (\text{fixed}, \text{fixed}).$$

And we have proven for the case when $(\text{fixed}, \text{fixed})$ is an NE.

$(\text{float}, \text{fixed})$ is an NE $\iff \beta_1 \leq r_2$ and $\beta_2 \leq 1$, according to Theorem 22. The successions to this NE from various profiles follow:

$$(\text{float}, \text{float}) \xrightarrow{\text{Proposition 12}} (\text{float}, \text{fixed}).$$

$$(\text{fixed}, \text{fixed}) \xrightarrow{\text{Proposition 12}} (\text{float}, ?) \xrightarrow{\text{as above}} (\text{float}, \text{fixed}).$$

$$(\text{fixed}, \text{float}) \xrightarrow{\text{Proposition 12}} (\text{float}, ?) \xrightarrow{\text{as above}} (\text{float}, \text{fixed}).$$

And we have proven for the case when $(\text{float}, \text{fixed})$ is an NE.

$(\text{fixed}, \text{float})$ is an NE $\iff \beta_1 \geq 1$ and $\beta_2 \geq r_1$, according to Theorem 22. The successions to this NE from various profiles follow:

$$(\text{fixed}, \text{fixed}) \xrightarrow{\text{Proposition 12}} (\text{fixed}, \text{float}).$$

$$(\text{float}, \text{fixed}) \xrightarrow{\text{Proposition 12}} (\text{fixed}, ?) \xrightarrow{\text{as above}} (\text{fixed}, \text{float}).$$

$$(\text{float}, \text{float}) \xrightarrow{\text{Proposition 12}} (\text{fixed}, ?) \xrightarrow{\text{as above}} (\text{fixed}, \text{float}).$$

And we have proven for the case when $(\text{fixed}, \text{float})$ is an NE.

$(\text{float}, \text{float})$ is an NE $\iff \beta_1 \leq 1$ and $\beta_2 \geq 1$, according to Theorem 22. The successions to this NE from various profiles follow:

$$(\text{float}, \text{fixed}) \xrightarrow{\text{Proposition 12}} (?, \text{float}) \xrightarrow{\text{Proposition 12}} (\text{float}, \text{float}).$$

$$(\text{fixed}, \text{float}) \xrightarrow{\text{as above}} (\text{float}, \text{float}).$$

$$(\text{fixed}, \text{fixed}) \xrightarrow{\text{Proposition 12}} (?, \text{float}) \xrightarrow{\text{as above}} (\text{float}, \text{float}).$$

And we have proven for the case when $(\text{float}, \text{float})$ is an NE. \square

This completes the analysis of the agents setting their own reciprocation attitudes. The next section considers agents who set both their own reciprocation attitudes and coefficients.

6. TWO AGENTS: RECIPROCATION ATTITUDE AND COEFFICIENT GAME

In the previous section we looked at the game of choosing a reciprocation attitude. It is also natural to consider what happens when the other habit, namely, the reciprocation coefficient, is chosen as well. Analyzing this game allows predicting the situation of more choice than the situation analyzed in RAG; for instance, the participants have more willpower or are just more knowledgeable than in RAG. As before, this is a one-shot game, the attitude and reciprocation coefficient being chosen once, before the interaction commences. As we did for RAG, since people usually neither completely mimic the others' behavior, nor do they completely ignore it, we assume $0 < r_1, r_2 < 1$. For simplicity, we also assume that all agents act synchronously. We call this game the *reciprocation attitude and coefficient game (RACG)*. This game is analyzed in the central Theorems 24 and 25.

We first characterize the existence of pure NE in this game and subsequently look into their efficiency, by finding the price of anarchy and stability. Then, we consider the best response dynamics. We assume that $k_2 > k_1$ (strictly) in this section.

Theorem 24. *The only Nash equilibria of RACG are characterized as follows:*

$$\begin{aligned}
 (\text{fixed, fixed, } r_1 = \beta_2, r_2 = \beta_1) \text{ is an NE} & \iff 0 < \beta_1, \beta_2 < 1. \\
 (\text{float, fixed, } 0 < r_1, r_2 < 1, \beta_1 \leq r_2) \text{ is an NE} & \iff \beta_1 < 1 \text{ and } \beta_2 \leq 1. \\
 (\text{fixed, float, } 0 < r_1, r_2 < 1, r_1 \leq \beta_2) \text{ is an NE} & \iff \beta_1 \geq 1 \text{ and } \beta_2 > 0. \\
 (\text{float, float, } 0 < r_1, r_2 < 1) \text{ is an NE} & \iff \beta_1 = \beta_2 = 1.
 \end{aligned}$$

The proof is based on Theorem 22, which narrows down the set of possible Nash equilibria, on Proposition 9 and Proposition 10 about utility maximization, and on convergence results from the previous chapter (See Section 2.)

Proof. We go over all the NE for RAG from Theorem 22 and look at all the possible choices of r_1 and r_2 to have an equilibrium in the new game. No other equilibria exist, since if no condition of Theorem 22 is satisfied, then even deviating by changing only the attitude is possible.

We begin with *(fixed, fixed)*, an NE in RAG if and only if $\beta_1 \geq r_2$ and $\beta_2 \leq r_1$. Given these reciprocation attitudes, Proposition 9 implies that to prevent the only best choice of r_1 being 0 or 1, we must have $(r_2 - \beta_1) = 0$, and to avoid the situation where the only best choice of r_2 is 0 or 1, we must have $(\beta_2 - r_1) = 0$. This implies the necessity of the conditions for an NE with *fixed* attitudes. Theorem 22 and Proposition 9 imply that these conditions are also sufficient to prevent deviations of only the attitude or only the reciprocation coefficient. If agent j simultaneously deviates to another attitude and r_j , then Corollary 3 on page 94 implies that any $r_j > 0$ yields the same utility, and therefore, this deviation may be considered to consist of attitude only, which is known to be not profitable. This proves the sufficiency.

Consider the profile *(float, fixed)* now, an NE in RAG if and only if $\beta_1 \leq r_2$ and $\beta_2 \leq 1$. Since $r_2 < 1$, we conclude that $\beta_1 \leq r_2 < 1$, and we have the necessity of the conditions for an NE with *floating* and *fixed* attitudes. Theorem 22 implies that deviating in attitude only is not profitable. According to Corollary 3, any $r_1, r_2 \in (0, 1)$ suffice for a best

response, and so deviating in reciprocation coefficient only is not profitable as well. Consider a deviation of an agent to another attitude and reciprocation coefficient simultaneously. Unless this includes r_2 becoming less than β_1 , we still know from what we have just proven that for this new profile, a deviation by the attitude only would not benefit agent 2, and since changing r_2 has not been profitable, the whole deviation is not profitable. The only remaining option is agent 2 becoming *floating* and changing r_2 to be less than β_1 . This would yield agent 2 the utility of $(1 - \beta_2)(\frac{r_2}{r_1+r_2}k_1 + \frac{r_1}{r_1+r_2}k_2)$, by Theorem 16 on page 87, while he previously had, by Corollary 3, $(1 - \beta_2)k_2$. Since $1 - \beta_2 \geq 0$ and $k_2 > k_1$, the previous profit is not smaller than the new one.

Consider the profile (*fixed, float*), which is an NE in RAG if and only if $\beta_1 \geq 1$ and $\beta_2 \geq r_1$. Since $r_1 > 0$, we conclude that $\beta_2 \geq r_1 > 0$, and the necessity of the conditions for an NE with *fixed* and *float* attitudes is proven. As always in this proof, Theorem 22 implies that deviating in attitude only is not profitable. According to Corollary 3, any $r_1, r_2 \in (0, 1)$ suffice for a best response, and so deviating in reciprocation coefficient only is not profitable as well. Consider a deviation of an agent to another attitude and reciprocation coefficient simultaneously. Unless this includes r_1 becoming greater than β_2 , we still know from what we have just proven that for this new profile, a deviation by the attitude only would not benefit agent 1, and since changing r_1 has not been profitable, the whole deviation is not profitable. The only remaining option is agent 1 becoming *floating* and changing r_1 to be more than β_2 . This would yield agent 1 the utility of $(1 - \beta_1)(\frac{r_2}{r_1+r_2}k_1 + \frac{r_1}{r_1+r_2}k_2)$, by Theorem 16, while he previously had, by Corollary 3, $(1 - \beta_1)k_1$. Since $1 - \beta_1 \leq 0$ and $k_2 > k_1$, the previous profit is not smaller than the new one.

Finally, look at (*float, float*), an NE in RAG if and only if $\beta_1 \leq 1$ and $\beta_2 \geq 1$. Given these reciprocation attitudes, Proposition 10 implies that we must have $\beta_1 = 1$, in order to prevent the only best choices of r_1 from being 0 or 1, and $\beta_2 = 1$, to avoid the same problems with the best choices of r_2 . This proves the necessity. We prove the sufficiency now. Theorem 22 implies that deviating in attitude only is not profitable, and Proposition 10 shows that deviating in only the reciprocation coefficient is not profitable either. Finally, consider a deviation of agent i in both attitude and r_i . Then, Corollary 3 on page 94 implies that any $r_j > 0$ yields the same utility, and therefore, this deviation may be considered to consist of attitude only, which is known to be not profitable. This proves the sufficiency. \square

Remark 18 (Existence of NE). *When no characterizing condition holds, no NE exists. For instance, if $\beta_1 < 1 < \beta_2$, no characterizing condition holds, and therefore, no (pure) NE exists.*

Let us exemplify the theorem for the parameter values from Example 12.

Example 12 (Continued). *For $\beta_1 = 0.3, \beta_2 = 0.6$, there exist two equilibria: (fixed, fixed, $r_1 = 0.6, r_2 = 0.3$) and (float, fixed, $0 < r_1, r_2 < 1, 0.3 \leq r_2$).*

6.1. PoA AND PoS

We now look at the efficiency of these equilibria, proving the following key result.

Theorem 25. *The efficiency of the equilibria is as follows:*

Conditions:	PoA :	PoS :
$0 < \beta_1, \beta_2 < 1$	$\frac{\sum_{i=1, j \neq i}^2 (1-\beta_i) \frac{(1-\beta_j)k_i + \beta_j(1-\beta_i)k_j}{1-\beta_j\beta_i}}{(2-\beta_1-\beta_2)k_2}$	1
$\{\beta_1 < 1\} \wedge \{\beta_2 \leq 1\} \wedge \neg \{0 < \beta_1, \beta_2 < 1\}$	1	1
$\{\beta_1 \geq 1\} \wedge \{0 < \beta_2 \leq 1\} \wedge \neg \{\beta_1 = \beta_2 = 1\}$	$\frac{(1-\beta_1-\beta_2)k_1}{(1-\beta_1)k_1 + (1-\beta_2)k_2}$	$\frac{(1-\beta_1-\beta_2)k_1}{(1-\beta_1)k_1 + (1-\beta_2)k_2}$
$\{\beta_1 \geq 1\} \wedge \{\beta_2 > 1\}$	1	1
$\beta_1 = \beta_2 = 1$	1	1

We find the possible NE from Theorem 24, and compare their social welfare with the optimal social welfare, found based on the proof of Proposition 13. We only use the ideas of what one should minimize or maximize to maximize the social welfare from the proof of Proposition 13, since the proposition sets reciprocation coefficients to 0 and 1, so we cannot use it directly. To calculate the social welfare, we employ the definition of utility and the limit values from Theorem 15 on page 85, Theorem 16 on page 87, and Corollary 3 on page 94.

Proof. If $0 < \beta_1, \beta_2 < 1$, Theorem 24 implies that there exist exactly two Nash equilibria, namely (*fixed, fixed*, $r_1 = \beta_2, r_2 = \beta_1$) and (*float, fixed*, $0 < r_1, r_2 < 1, \beta_1 \leq r_2$). For the optimal social welfare, we need to maximize both $\lim_{t \rightarrow \infty} x(t)$ and $\lim_{t \rightarrow \infty} y(t)$, as does, for instance, the second NE above, yielding the social welfare of $(2 - \beta_1 - \beta_2)k_2$. Taking the ratios of the social welfare values gives row one in the table from the statement of the theorem.

If $\beta_1 < 1$ and $\beta_2 \leq 1$ but not $0 < \beta_1, \beta_2 < 1$, then there exists only the NE (*float, fixed*, $0 < r_1, r_2 < 1, \beta_1 \leq r_2$), according to Theorem 24. As to the optimal social welfare, we need to maximize both $\lim_{t \rightarrow \infty} x(t)$ and $\lim_{t \rightarrow \infty} y(t)$, and since the only NE does this, $\text{PoA} = \text{PoS} = 1$.

If $\beta_1 \geq 1$ and $1 \geq \beta_2 > 0$ but not $\beta_1 = \beta_2 = 1$, then the only NE, accordig to Theorem 24, is (*fixed, float*, $0 < r_1, r_2 < 1, r_1 \leq \beta_2$). To maximize the social welfare, we need to minimize $\lim_{t \rightarrow \infty} x(t)$ and maximize $\lim_{t \rightarrow \infty} y(t)$; obtaining the social welfare of $(1 - \beta_1)k_1 + (1 - \beta_2)k_2$. Taking the ratio of the social welfare values in the equilibria to $(1 - \beta_1)k_1 + (1 - \beta_2)k_2$, we obtain the third row.

If $\beta_1 \geq 1$ and $\beta_2 > 1$, Theorem 24 states that the only NE is (*fixed, float*, $0 < r_1, r_2 < 1, r_1 \leq \beta_2$). As to maximizing the social welfare, we need to minimize both $\lim_{t \rightarrow \infty} x(t)$ and $\lim_{t \rightarrow \infty} y(t)$, and since the only NE does this, $\text{PoA} = \text{PoS} = 1$, implying the penultimate row of our table.

If $\beta_1 = \beta_2 = 1$, Theorem 24 implies that there exist exactly the following NE: (*fixed, float*, $0 < r_1, r_2 < 1, r_1 \leq \beta_2$) and (*float, float*, $0 < r_1, r_2 < 1$). The social welfare is always zero, regardless of the strategy profile, and so every NE is optimal, i.e. $\text{PoA} = \text{PoS} = 1$. \square

And again, we exemplify the theorem for Example 12.

Example 12 (Continued). *For $\beta_1 = 0.3, \beta_2 = 0.6$, there exist two equilibria: (fixed, fixed, $r_1 = 0.6, r_2 = 0.3$) and (float, fixed, $0 < r_1, r_2 < 1, 0.3 \leq r_2$). The price of stability is 1, since*

the NE (float, fixed, $0 < r_1, r_2 < 1, 0.3 \leq r_2$) is socially optimal. The equilibrium (fixed, fixed, $r_1 = 0.6, r_2 = 0.3$) is not optimal, causing the price of anarchy to be

$$\frac{0.7 \frac{0.4k_1 + 0.6 \cdot 0.7k_2}{0.82} + 0.4 \frac{0.7k_2 + 0.3 \cdot 0.4k_1}{0.82}}{(2 - 0.3 - 0.6)k_2} = \frac{0.36}{k_2} \frac{k_1}{k_2} + \overline{0.63}.$$

This example shows that for small k_1/k_2 , there exists a socially inefficient equilibrium. In general, for a reciprocation attitude game, Theorem 23 implies that small enough β_1, β_2 guarantee that all the NE are optimal. In RACG, however, when $0 < \beta_1, \beta_2 < 1$, we see in the proof of Theorem 25 that along with a socially optimal NE, the social welfare of the NE (fixed, fixed, $r_1 = \beta_2, r_2 = \beta_1$) relative to the optimum is $\frac{\sum_{i=1,2;j \neq i} (1-\beta_i) \frac{(1-\beta_j)k_i + \beta_j(1-\beta_i)k_j}{1-\beta_j\beta_i}}{(2-\beta_1-\beta_2)k_2}$. The limit of this expression when the efforts of acting approach zero for both agents is

$$\lim_{\beta_1 \rightarrow 0, \beta_2 \rightarrow 0} \frac{\sum_{i=1,2;j \neq i} (1-\beta_i) \frac{(1-\beta_j)k_i + \beta_j(1-\beta_i)k_j}{1-\beta_j\beta_i}}{(2-\beta_1-\beta_2)k_2} = \frac{\sum_{i=1,2;j \neq i} k_i}{2k_2} = \frac{k_1 + k_2}{2k_2} = \frac{1}{2} \left(\frac{k_1}{k_2} + 1 \right).$$

That is, allowing more freedom (setting own reciprocation attitude and coefficient), we may lose up to half of the efficiency, if k_1/k_2 is small. However, Theorem 25 leaves a sparkle of hope: if at least one agent acts completely effortlessly or even enjoys it, meaning that $\beta_i \leq 0$, then all the NE are socially optimal. In this case, the only equilibrium is (float, fixed, $0 < r_1, r_2 < 1, \beta_1 \leq r_2$), so the less kind agent 1 mimics the kinder agent 2.

6.2. CONVERGING TO NE

In Section 5.2, we show that in RAG, the best response dynamics can move to any NE from any strategy profile. We want to prove an analogous proposition for RACG, but now, the non-compactness of the domain does not allow a best response to always exist. To show this, consider the following example. Pick β_1, β_2 such that $1 > \beta_2 > \beta_1 > 0$. Theorem 24 says that (fixed, fixed, β_2, β_1) is an NE. We show now that there exists a profile that does not possess a best response. Pick r_2 such that $\beta_1 > r_2 > 0$, and consider the profile where agent 2 plays fixed, r_2 . We now look what a best response of 1 would be. If agent 1 plays fixed, then Proposition 9 implies that he should also play $r_1 = 0$. This yields her the utility of $(1-r_2)k_2 + (r_2-\beta_1)k_1$. Since $r_1 = 0$ is not an option, we understand this as a limit. If agent 1 plays floating, she obtains the utility of $(1-\beta_1)k_2$, which is (strictly) smaller than $(1-r_2)k_2 + (r_2-\beta_1)k_1$. Consequently, there is no best response, since playing fixed and choosing r_1 that is approaching 0 is the best.

Having analyzed games of two agents, we next turn to the case of n agents.

7. *n* AGENTS

From now on, we handle the general case of agents interacting with many agents. We aim to model and analyze the game of setting own reciprocation coefficients. To find the equilibria of this game, we find best responses, and to quantify their efficiency we need to maximize the social welfare.

Let us warm up with the following example about the reactive agents following the stable ones. The reader is advised to skip the proof at the first reading.

Example 13. Assume the interaction graph G to be a clique. Assume that the agents act synchronously. If for all agents $i \in C \subseteq N, C \neq \emptyset$, we have $r_i = r'_i = 0$, and for all agents $j \notin C$ we have $r'_j = 1$ ($\Rightarrow r_j = 0$),⁶ then for each $j \notin C$ and $(j, k) \in E$, the limit $L_{j,k}$ exists and is equal to $\frac{\sum_{i \in C} k_i}{|C|}$.

This case does not fit the conditions of Theorem 20 on page 98, so that even the existence of the limits has to be proven here. The statement states that the agents in C are asymptotically determining for the other agent, in the sense that all the rest will asymptotically act as the arithmetic average of the agents in C . The idea of the proof is to first look at the limit of the sum of all the actions at time t , then at $\lim_{t \rightarrow \infty} \text{got}_j(t)$, and these imply the statement.

Proof. Since $r_i + r'_i = 0$, agent $i \in C$ will always act k_i on any agent. We assumed $0 < |C| \leq n$. If $|C| = n$, we have nothing to prove, so assume that $0 < |C| < n$.

Denote the sum of all the actions at time t as $S(t) \triangleq \sum_{i,j \in N} x_{i,j}(t)$. We first show that $\lim_{t \rightarrow \infty} S(t) = n(n-1) \frac{\sum_{i \in C} k_i}{|C|}$. Since for every agent besides those in C the sum of her actions is equal to the sum of the previous actions upon it, the total sum can only change because of the agents in C . At any time $t \in T$, the sum of all the actions besides those done by agents in C is $S(t) - \sum_{i \in C} (n-1)k_i$ and therefore any agent i in C gets $\frac{S(t) - \sum_{i \in C} (n-1)k_i}{n-1} + \sum_{l \in C \setminus \{i\}} k_l = \frac{S(t)}{n-1} - \sum_{i \in C} k_i + \sum_{l \in C \setminus \{i\}} k_l$, while her own actions always sum up to $(n-1)k_i$. Thus, after time t the difference from $|C|S(t)$ to $n(n-1) \sum_{i \in C} k_i$ has changed by $|C| \cdot (\sum_{i \in C} (n-1)k_i - (|C| \frac{S(t)}{n-1} - |C| \sum_{i \in C} k_i + (|C|-1) \sum_{l \in C} k_l)) = |C| \cdot ((n-1+|C|-|C|+1) \sum_{i \in C} k_i - |C| \frac{S(t)}{n-1}) = |C| \cdot ((n) \sum_{i \in C} k_i - |C| \frac{S(t)}{n-1}) = |C| \cdot (\frac{(n(n-1) \sum_{i \in C} k_i - |C|S(t))}{n-1})$. Thus, after each time slot, exactly $\frac{n-1-|C|}{n-1}$ of the difference $(n)(n-1) \sum_{i \in C} k_i - |C|S(t)$ remains, and therefore, $\lim_{t \rightarrow \infty} |C|S(t) = n(n-1) \sum_{i \in C} k_i \Rightarrow \lim_{t \rightarrow \infty} S(t) = n(n-1) \frac{\sum_{i \in C} k_i}{|C|}$.

Consider any agent j not in C now. We now prove that $\lim_{t \rightarrow \infty} \text{got}_j(t) = (n-1) \frac{\sum_{i \in C} k_i}{|C|}$. For any positive ϵ , consider a time t , such that from this time on $S(t)$ is in the ϵ -environment of its limit, namely of $n(n-1) \frac{\sum_{i \in C} k_i}{|C|}$. At time t , j 's total actions sum up to $\text{got}_j(t-1)$, and therefore it receives then $\frac{S(t) - \text{got}_j(t-1)}{n-1}$, which is in the ϵ -environment of $n \frac{\sum_{i \in C} k_i}{|C|} - \frac{\text{got}_j(t-1)}{n-1}$. Thus, $\text{got}_j(t) - \text{got}_j(t-1)$ is in the ϵ -environment of $n \frac{\sum_{i \in C} k_i}{|C|} - \frac{\text{got}_j(t-1)n}{n-1} = \frac{n}{n-1} ((n-1) \frac{\sum_{i \in C} k_i}{|C|} - \text{got}_j(t-1))$. Thus, the difference $(n-1) \frac{\sum_{i \in C} k_i}{|C|} - \text{got}_j(t-1)$ changes such that only $\frac{1}{n-1}$ of its absolute value remains and the sign changes, all up to ϵ . Since the ϵ is arbitrary, this implies that $\lim_{t \rightarrow \infty} \text{got}_j(t) = (n-1) \frac{\sum_{i \in C} k_i}{|C|}$.

Since every agent acts the same on all the neighbors, for any $j \notin C$ and any $l \neq j$ we have $x_{j,l} + \text{got}_j(t) = \frac{S(t)}{n-1}$ and therefore, $\lim_{t \rightarrow \infty} S(t) = n(n-1) \frac{\sum_{i \in C} k_i}{|C|}$ together with $\lim_{t \rightarrow \infty} \text{got}_j(t) = (n-1) \frac{\sum_{i \in C} k_i}{|C|}$ imply that $L_{j,l}$ exists and is equal to $\frac{\sum_{i \in C} k_i}{|C|}$. \square

8. UTILITY MAXIMIZATION

As the first step to studying strategic choice reciprocation habits, to be able to facilitate these decisions, we consider what reciprocation strategy an agent should adopt to max-

⁶With these reciprocation parameters, all the agents can be considered as either *fixed* or *floating*.

imize her utility. In the case of example 5 on page 79, this models a colleague changing her behavioral habits, to improve her own well-being, as a result of a psychologist's advice, for instance. In the case of example 9, this models a country choosing her defence policy, to improve her own well-being, as a result of a specialist's or a decision system's advice, for instance. We now analyze how i 's utility depends on r_i , assuming that the other parameters are set. Recall that r_i is the fraction of an agent's action, determined by responding to the previous action of the other agent she acts on. When all agents employ floating reciprocation, we prove:

Proposition 16. *Given a connected interaction graph with an odd cycle, where all agents act synchronously and with floating reciprocation, assume that for all agents i , $r'_i > 0$. Then, each i can maximize her utility by setting r_i to $1 - r'_i$, if*

$$(1 - \beta_i) \left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \cdot k_j \right) - k_i \left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \right) \right) \right) \quad (5.4)$$

is positive, to 0 if negative, and to an arbitrary value if zero. When this expression is not zero, these choices are the only optimal ones.

Equivalently, we may have looked at the sign of $(1 - \beta_i)(L - k_i)$, where L is the common limit from Theorem 20 on page 98.

As mentioned after Theorem 20, we almost always have an odd cycle, by having three people acting on each other. Assuming the positivity of r'_i is also very realistic, since completely ignoring the other side is rare.

The idea of the proof is maximizing the expression for i 's utility.

Proof. According to Theorem 20, all the actions converge to a common limit $L = \frac{\sum_{i \in N} \left(\frac{d(i)}{r_i + r'_i} \cdot k_i \right)}{\sum_{i \in N} \left(\frac{d(i)}{r_i + r'_i} \right)}$. The utility of i is thus $(d(i)L) - \beta_i(d(i)L) = (1 - \beta_i)d(i)L$. The derivative is

$$\begin{aligned} \frac{\partial u_i}{\partial r_i} &= (1 - \beta_i)d(i) \cdot \frac{\frac{-d(i)k_i}{(r_i + r'_i)^2} \left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \right) \right) + \sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \cdot k_j \right) \frac{d(i)}{(r_i + r'_i)^2}}{\left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \right) \right)^2} \\ &= (1 - \beta_i)d(i) \cdot \frac{\frac{d(i)}{(r_i + r'_i)^2} \left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \cdot k_j \right) - k_i \left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \right) \right) \right)}{\left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \right) \right)^2} \\ &= (1 - \beta_i)d(i) \cdot \frac{\frac{d(i)}{(r_i + r'_i)^2} \left(\sum_{j \in N \setminus \{i\}} \left(\frac{d(j)}{r_j + r'_j} \cdot k_j \right) - k_i \left(\sum_{j \in N \setminus \{i\}} \left(\frac{d(j)}{r_j + r'_j} \right) \right) \right)}{\left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \right) \right)^2}. \end{aligned}$$

Therefore, the derivative is zero either for all r_i or for none. In any case, the maximum is attained at an endpoint. To avoid a complicated substitution, we consider the derivative

sign:

$$\frac{\partial u_i}{\partial r_i} \geq 0 \iff (1 - \beta_i) \cdot \left(\sum_{j \in N \setminus \{i\}} \left(\frac{d(j)}{r_j + r'_j} \cdot k_j \right) - k_i \left(\sum_{j \in N \setminus \{i\}} \left(\frac{d(j)}{r_j + r'_j} \right) \right) \right) \geq 0,$$

and so when $(1 - \beta_i) \sum_{j \in N \setminus \{i\}} \left(\frac{d(j)}{r_j + r'_j} \cdot k_j \right) - k_i \left(\sum_{j \in N \setminus \{i\}} \left(\frac{d(j)}{r_j + r'_j} \right) \right)$ is nonnegative, i should choose the largest r_i , which is $1 - r'_i$, and she should choose 0 otherwise. When the derivative is not zero, these choices are the only optimal ones. \square

Remark 19 (Free actions). *For $\beta_i = 0$, which is the case when i just wants to receive more action weight, and does not mind acting herself, the result has the following intuitive appeal: The agents with original kindness smaller than the weighted average of the kindness values should choose to be very reciprocating, while the other agents should choose to be completely non-reciprocating, thereby remaining kind and pulling the other agents to act more.*

The next section looks into maximizing the social welfare. This is interesting both to advise the manager how to do it and to subsequently quantify the efficiency of the equilibria of the game of choosing interaction habits.

9. MAXIMIZING SOCIAL WELFARE

We now analyze how the manager can maximize the total utility by setting the r_i coefficients. This allows for advising managers and for analyzing the efficiency of equilibria in the next section. Recall that we assume w.l.o.g., that $k_n \geq \dots \geq k_1$.

When all agents employ *floating* reciprocation, we provide a way to maximize the social welfare.

Proposition 17. *Given a connected interaction graph with an odd cycle, where all agents act synchronously and employ floating reciprocation, assume that for all agents i , $r'_i > 0$. We can maximize the social welfare by the following procedure. First, set all r_i to $1 - r'_i$. Order the agents in a non-increasing (non-decreasing, for $\sum_{i \in N, j \in N(i) \setminus \{i\}} (1 - \beta_i) \leq 0$) order of kindness; w.l.o.g., let the obtained order be $n, n-1, \dots, 1$. Set r_n to zero. Go over the other agents in this order, and as long as*

$$\left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \cdot k_j \right) - k_i \left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \right) \right) \right)$$

is negative (positive, for $\sum_{i \in N, j \in N(i) \setminus \{i\}} (1 - \beta_i) \leq 0$), set r_i to zero.

Equivalently, we may have looked at the sign of $(L - k_i)$, for the L from Theorem 20 on page 98.

The conditions $r'_i > 0$ and the connected graph having an odd cycle (say, a triangle) are met in any reasonable setting where every agent reacts to the others, and there are at least 3 agents who directly influence each other.

The proof is by induction on the suggested procedure.

Proof. All the actions converge to a common limit $L = \frac{\sum_{i \in N} \left(\frac{d(i)}{r_i + r'_i} \cdot k_i \right)}{\sum_{i \in N} \left(\frac{d(i)}{r_i + r'_i} \right)}$, by Theorem 20.

The social welfare is $SW = L \sum_{i \in N, j \in N(i) \setminus \{i\}} (1 - \beta_i)$. Thus, we need to maximize L , if $\sum_{i \in N, j \in N(i) \setminus \{i\}} (1 - \beta_i) \geq 0$, and minimize it otherwise. W.l.o.g., assume that $\sum_{i \in N, j \in N(i) \setminus \{i\}} (1 - \beta_i) \geq 0$; in the other case, the proof is analogous.

We now prove that the described procedure maximizes the social welfare by induction on the handled agents. The induction basis is agent n , which, being the kindest, should obviously receive the highest weight. At an induction step, we compare the next agent, say i , to the weighted average of the other agents, and if k_i is larger, we assign r_i to zero. This maximizes the common limit L . The assignment $r_i = 0$ does not change the optimal assignment to the previous agents, since they are all at least as kind as i is. At an induction step where the first time an agent, say l , is not assigned $r_l = 0$, and, therefore, remains to be $1 - r'_l$, the weighted average of the other agents was not smaller than k_l , so, to maximize L , kindness k_l should have received the minimum possible weights. The rest of the agents are not kinder than l , so by leaving them with the current reciprocation coefficient, the social welfare keeps being maximized. We conclude that the final assignment of reciprocation coefficients maximizes the social welfare. \square

Remark 20 (Intuition for free actions). *For $\beta_1 = \beta_2 = \dots = \beta_n = 0$, the result is intuitive, since the less kind agents are set to reciprocate as much as possible, while the kinder agents are set to reciprocate as little as possible, thereby remaining kind and pulling the other agents to act more.*

In Section 8 we analyze utility maximization by an agent, which is required for analyzing NE. The current section allows for analyzing the efficiency of equilibria relatively to the maximum possible social welfare. Both topics are studied in the following section.

10. RECIPROCATION COEFFICIENT GAME

We have considered an agent choosing her reciprocation coefficient, each choice yielding certain (asymptotic) utility to the agent. Therefore, the situation is naturally modeled as a game where the strategies of each agent are the choices of her reciprocation coefficient and her utility is the asymptotic utility of the interaction. Recall that the utility of agent i is $\lim_{t \rightarrow \infty} \{ \sum_{j \in N(i)} x_{j,i}(t) - \beta_i \sum_{j \in N(i)} x_{i,j}(t) \}$. This is a one-shot game, the attitudes being chosen once, before the interaction commences. Analyzing this game allows predicting the situation, useful for supplying some advice to an external party (such as the manager who wants to influence the agents) or the agents themselves. Since usually at least three completely connected agents exist, and people do not ignore the others completely, we always assume we are given a connected interaction graph with an odd cycle, and for all agents i , $r'_i > 0$. To be able to analyze the process, we assume that all agents act synchronously and employ *floating* reciprocation. We call this game the *reciprocation coefficient game* (RCG) and summarize its analysis in Theorem 26.

We first characterize the existence of pure NE in this game and subsequently look into their efficiency, by finding the price of anarchy and stability. Proposition 16 allows us to characterize the existence of pure NE in this game, as follows.

Theorem 26. *The profile (r_1, r_2, \dots, r_n) is a Nash equilibria of reciprocation coefficient game if and only if for every $i \in N$ there holds*

$$r_i = 1 - r'_i \iff (1 - \beta_i) \cdot \left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \cdot k_j \right) - k_i \left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \right) \right) \right) > 0.$$

$$r_i = 0 \iff (1 - \beta_i) \cdot \left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \cdot k_j \right) - k_i \left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \right) \right) \right) < 0.$$

If $(1 - \beta_i) \left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \cdot k_j \right) - k_i \left(\sum_{j \in N} \left(\frac{d(j)}{r_j + r'_j} \right) \right) \right) = 0$, then any $r_i \in [0, 1 - r'_i]$ is satisfactory.

There always exists a pure NE, by Proposition 20.3 in [18].

10.1. PoA AND PoS

We now find how far the social welfare in an equilibrium is from the maximum possible social welfare. This indicates whether regulation is required. To this end, consider the famous measures of the efficiency of an equilibrium, namely price of anarchy [26] (PoA) and price of stability [27] (PoS). Using Theorem 26, we know for each given set of parameters whether a profile is an NE or not. Using Proposition 17, we know for each given set of parameters what the maximum social welfare for a given set of parameters is. Calculating the social welfare at each of the Nash equilibria and finding its ratio to the optimum social welfare enables us to find the price of anarchy and stability, to measure efficiency. We first prove that all the NE are optimal for similar β_i s, and then we prove that in any case, the best and the worst equilibria have the same social welfare, meaning that PoA = PoS.

First, we show the optimality of the NE for similar β_i s.

Proposition 18. *If all $1 - \beta_i$ s have the same sign, then a profile maximizes the social welfare if and only if it is an NE. In particular, PoA = 1.*

Proof. Since all $1 - \beta_i$ s have the same sign, then either all the agents need to maximize or they all need to minimize the common limit in order to maximize their utility. This implies that being socially optimal is equivalent to being individually optimal for everyone. This means that every social welfare maximizing profile is an NE. Additionally, the proof of Proposition 17 demonstrates that if a profile is not optimal for the social welfare, then there exists an agent, whose utility can be unilaterally increased. Therefore, an NE has to be socially optimal. \square

On the way to prove that PoA = PoS, we are going to prove that exactly one NE exists, up to agents with borderline kindness. First, we denote the agents with nonnegative $1 - \beta_i$ by P and the rest by $M \triangleq N \setminus P$. We begin by making the following observation.

Observation 9. *In an NE with the common limit L , all the agents $i \in P$ with k_i less than L play $r_i = 1 - r'_i$, while all $i \in P$ with k_i larger than L play $r_i = 0$. Anti-symmetrically, all the agents $i \in M$ with kindness less than L play $r_i = 0$, and all i in M with kindness larger than L play $r_i = 1 - r'_i$. The agents with kindness equal to L can play arbitrarily, preserving the equilibrium.*

Proof. Directly from Theorem 26. □

Now, we prove the uniqueness:

Proposition 19. *Given a RCG, it has a single equilibrium, up to arbitrary strategies of the agents with kindness values equal to the common limit L of the equilibrium.*

Proof. The observation we have just proven implies it is enough to prove the uniqueness of the common limit L . Assume to the contrary that two equilibria with $L' > L$ exist. If now all the agents of M adopt their strategy in the NE with L , then their utility will not increase, and it definitely will not increase if then all the agents in N best respond by adopting their strategy in the NE with L , which contradicts the decrease in the common limit. □

This essential uniqueness on NE immediately yields

Corollary 4. $\text{PoA} = \text{PoS}$.

Let us demonstrate our results on Example 9.

Example 9 (Continued). *Assume there are $n = 3$ neighboring countries where everyone acts on everyone else, and that $r'_1 = r'_2 = 0.3, r'_3 = 0.1, \beta_1 = 0, \beta_2 = 0.2, \beta_3 = 0.1$ (acting is cheap), $k_1 = 0, k_2 = 0.5, k_3 = 1$. In addition, all the countries act synchronously and employ floating reciprocation. By Proposition 19 and Proposition 18, the only NE is optimal. Let us find the NE directly, to exemplify how things work. By Proposition 17, to maximize the social welfare, we should make the countries choose $(0.7, 0.7, 0)$. Consider an NE (r_1, r_2, r_3) in this 3-player game. Since Formula (5.4) is positive for $i = 1$, we conclude that $r_1 = 1 - r'_1 = 0.7$. Since Formula (5.4) is negative for $i = 3$, and we have $r_3 = 0$. Now, we see that Formula (5.4) is positive, and we conclude that $r_2 = 0.7$. Thus, $(0.7, 0.7, 0)$ is the only NE, and we obtain again that $\text{PoA} = \text{PoS} = 1$.*

10.2. CONVERGING TO NE

In order to show stability of an equilibrium, we prove that the best response dynamic can converge to the NE from any starting profile, similarly to Proposition 15.

Proposition 20. *For any strategy profile there exists a succession of profiles starting from it and terminating at the (essentially) single NE, such that each profile in this succession is an optimum reaction of each agent to the others' strategies in the previous profile.*

The idea of the proof is that the limits L monotonically approach that of the NE. It has to stop, and once it does, we are at an equilibrium.

Proof. Analogously to Observation 9, in any best response to a profile with the common limit L , all the agents $i \in P$ with k_i less than L play $r_i = 1 - r'_i$, while all $i \in P$ with k_i larger than L play $r_i = 0$. Anti-symmetrically, all the agents $i \in M$ with kindness less than L play $r_i = 0$, and all i in M with kindness larger than L play $r_i = 1 - r'_i$. The agents with kindness equal to L can play arbitrarily, preserving the equilibrium.

We now look at a process of best responses. Once the common limit does not change, an NE has been achieved. We now show that the common limit changes monotonically.

Given any starting profile, let the agents best respond, obtaining a common limit L_1 . In the second best response, assume, w.l.o.g. that the new common limit L_2 is larger than L_1 . Assume to the contrary that the common limit after the third best response, L_3 , is smaller than L_2 . The agents in M will prefer the profile where the agents in P move to their strategies after the second best response, and even more if then the agents in M themselves move to play what they did after the second best response. As this contradicts $L_2 > L_3$, we conclude that the common limits may change only monotonically. The structure of best responses we described in the beginning of the proof implies that the changes are discrete. Since the common limit is bounded by the extremal kindness values, the process stops in finite time. As we note before, this is an NE. \square

Remark 21 (Other strategy spaces). *If, instead of choosing r_i 's, the agents could choose their r_i 's in some $[a, 1 - r_i]$, for a positive a , all the analysis of the game and the efficiency of its equilibria would be completely analogous, because of the symmetry of Equation (4.11) from page 98.*

To summarize, we have analyzed stable states (NE) of choosing habits and their impact on the social welfare (PoA and PoS). It is now reasonable to look back at the model we have been analyzing.

11. DISCUSSION OF THE MODEL

This section discusses reasonable changes to several definitions of our modeling. It is an explanation of what we do and what else can be done. We first look into the definition of utility in the theoretically infinite interaction, and then consider the habits that are being set.

11.1. UTILITY AGGREGATION

In all the games we consider, we model utility as the limit of the utility at time t . This is a far-sighted model, though the exponential convergence makes the required sight not so "far". Our game is not repeated, since the choice is made before the repeated interaction commences; still, utility is obtained at each interaction in a repeated manner. It is therefore interesting to compare the utility modeling in our model with the traditional modeling of preferences comparison in repeated games, like that in Osborne and Rubinstein [18, Chapter 8.3]. They consider the following three options:

Discounting with a factor $\delta \in (0, 1)$, where $u_i \triangleq (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_{i,t}$. The factor $1 - \delta$ promises that if $u_{i,t} \equiv 1$, then $u_i = 1$.

Limit of (arithmetic) means, where utility sequence $u_{i,t}$ is said to be strictly preferred to $v_{i,t}$, if $\liminf_{T \rightarrow \infty} \frac{\sum_{t=0}^{T-1} (u_{i,t} - v_{i,t})}{T}$ is positive.

Overtaking, where they do not define utilities, and $u_{i,t}$ is preferred to $v_{i,t}$, if $\liminf_{T \rightarrow \infty} \sum_{t=0}^T (u_{i,t} - v_{i,t})$ is positive.

As Osborne and Rubinstein mention, if one sequence of utilities $u_{i,t}$ is strictly preferred to another sequence $v_{i,t}$ in the limit of means, then there exists large enough

$\delta \in (0, 1)$, such that $u_{i,t}$ is also strictly larger in the sense of discounting with this δ . We prove now that our definition is not really different from these other ones.

Proposition 21. *Given a pair of converging sequences $\{u_{i,t}\}_{t=0}^{\infty}$ and $\{v_{i,t}\}_{t=0}^{\infty}$, regardless whether they can be obtained in some reciprocation process, the following relations hold between our model and the models above:*

1. *Preference in our model is equivalent to preference in the sense of limit of mean and in the sense of overtaking.*
2. *A strict preference in our model implies that there exists large enough $\delta \in (0, 1)$, such that the same preference in discounting holds. We can say nothing if equality holds in our model. In reverse, if a preference in discounting holds for every $\delta \in (a, 1)$, for some $a \in (0, 1)$, then the same preference or equality holds in our model.*
3. *For some values of δ , our model can neither imply discounting, nor be implied by it.*

Proof. Assume that sequence $u_{i,t}$ is strictly preferred to another sequence $v_{i,t}$ in our model, that is $\lim_{t \rightarrow \infty} u_{i,t} > \lim_{t \rightarrow \infty} v_{i,t}$. Therefore, $\lim_{t \rightarrow \infty} u_{i,t} - v_{i,t} > 0$. Convergence of a sequence implies the convergence of the sequence of its arithmetic means to the same limit (see, for instance, Example 1 on pp. 95 of [28]), so $\lim_{T \rightarrow \infty} \frac{\sum_{t=0}^{T-1} (u_{i,t} - v_{i,t})}{T}$ exists and is positive, which means the preference in the limit of means. By the remark above, this implies preference in discounting, for large enough $\delta \in (0, 1)$. To prove preference in overtaking as well, we write

$$\lim_{T \rightarrow \infty} \sum_{t=0}^{T-1} (u_{i,t} - v_{i,t}) = \lim_{T \rightarrow \infty} \frac{\sum_{t=0}^{T-1} (u_{i,t} - v_{i,t})}{T} T,$$

and since $\lim_{T \rightarrow \infty} \frac{\sum_{t=0}^{T-1} (u_{i,t} - v_{i,t})}{T} > 0$, the result is infinitely large, proving preference in overtaking.

By now, we have shown what our model implies in part 1 and part 2, if the original preference is strict. If the limits are equal then the analogous proof goes through, besides the case of discounting. If the limits are equal, the discounting utility can be either way. For example, the limits of $\{1, 1, 1, \dots, 1, \dots\}$ and $\{1.5, 0, 1, 1, \dots, 1, \dots\}$ are both 1. The discounted utilities differ only in their first part, which is (omitting $q - \delta$) $1 + \delta$ and 1.5, respectively. The comparison here depends on the comparison between δ and 1.5.

To prove the other direction, notice that since we assume convergence, every model provides some preference, strict or not. Therefore, the second direction ensues from the already shown direction of the implications together with the symmetry of the larger and smaller options, besides, perhaps, the case of being given a preference in the discounted model. In that case, the proved implication of strict preferences implies that if sequence $\{u_{i,t}\}_{t=0}^{\infty}$ is strictly preferred to sequence $\{v_{i,t}\}_{t=0}^{\infty}$ in the discounted sense for any large enough $\delta \in (0, 1)$, then $\lim_{t \rightarrow \infty} u_{i,t} \geq \lim_{t \rightarrow \infty} v_{i,t}$, since otherwise, $\{u_{i,t}\}_{t=0}^{\infty}$ would be strictly worse than $\{v_{i,t}\}_{t=0}^{\infty}$, by the already proven. However, the example above demonstrates that the equality can indeed take place. Finally, if the discounting model gives us equality, then so does our model, because otherwise, we have proven that the discounted model would not give equality. This completes the other direction of parts 1 and 2.

In order to show part 3, consider the following example. For a positive b , let one sequence be $\{u_{i,t}\}_{t=0}^{\infty} = (M, 1, 1, \dots, 1, \dots)$ and the other one $\{v_{i,t}\}_{t=0}^{\infty} = (1 + b, 1 + b, 1 + b, \dots, 1 + b, \dots)$. The discounted utility of $\{u_{i,t}\}_{t=0}^{\infty}$ is $M + \frac{\delta}{1-\delta}$, and the utility of $\{v_{i,t}\}_{t=0}^{\infty}$ is $\frac{1+b}{1-\delta}$. For any δ , if M is large enough, the first sum will be larger, despite that its limit is 1, which is smaller than the limit of the second sequence, $1 + b > 1$. \square

Parts 1 and 2 of Proposition 21 imply that when interaction converges, our utility aggregation indicates the other possible aggregations well, especially if the discounting is slow enough. We thus conclude that our definition is a good fit to the existing practices.

11.2. STRATEGY SPACE

When modeling games, we define, alongside with utilities, also the possible strategies, meaning the domain of parameters. We now discuss other possible models.

In a reciprocation attitude and coefficient game (RACG), we allow choosing r_i in $(0, 1)$. The whole segment $[0, 1]$ is, as we explain, often unrealistic; in addition, the reciprocation process may not converge at the edges. However, we can choose a closed segment $[a, b]$, for any $0 < a < b < 1$. This would limit the domain, but the compactness of the domain may facilitate existence of NE. This is an possible direction for future work. On the other hand, allowing the extreme points $r_i = 0$ or 1 with a proper handling of the cases of no convergence is also an alternative.

For two agents we are able to analyze the game of choosing the reciprocation coefficient for the not *floating* case too. However, analyzing all the possibilities would be too lengthy and would not convey new ideas.

Even more directions for continuing the research refer to the reciprocation coefficient games (RCG), where agent i chooses $r_i \in [0, 1 - r'_i]$. As we mention at the end of that section, we can instead allow choosing r'_i in some $[a, 1 - r_i]$, for a positive a , and obtain the symmetrical results. However, choosing r'_i in the half-open segment $(0, 1 - r_i]$ may be another story, where the lack of compactness jeopardizes the existence of NE. Another possible game would be allowing to choose r_i and r_i simultaneously.

We can never cover every possible model, but we believe our model sheds light on the general phenomena.

12. CONCLUSIONS AND FUTURE WORK

We first summarize the chapter and present its main conclusions, continuing to several interesting directions for future research.

We aim to predict and advise on strategic behavior in reciprocation, in both human-to-human and human-to-machine interactions. A reciprocal action is modeled as a balance between the inner self and a reaction to others' actions. We define an agent's utility asymptotically and prove the equivalence of this definition to slowly discounted and other classical utilities. We then consider choosing the reciprocation attitude or coefficient to maximize her own utility. We finally model the strategic behavior of the reciprocating agents in several games, characterize the NE and their efficiency. We also show that NE may always be achieved by a natural process, the best response dynamics [18, Section 2.2], besides in a RACG. This gives hope for achieving a situation that is stable to unilateral deviations without any regulation.

Our main advice is that both for maximizing own utility and for maximizing the social welfare, if contributing is cheaper than receiving, then, both in choosing the reciprocation attitude and coefficient, the kinder agent should be most stable (be *fixed* or have the reciprocation coefficient $r_i = 0$), and the opposite should be done if contributing is costlier than receiving. When contributing is much cheaper than receiving (β_i s are smaller than all the other parameters), then, for the reciprocation attitude game and for the reciprocation coefficient game, the price of anarchy is 1, so rationally reciprocating agents will play socially optimally. In such equilibria, the kinder agents are stable and the less kind agents follow the kinder ones. For the reciprocation attitude and coefficient game, the price of stability is 1, but the price of anarchy is positive, meaning that rationally reciprocating agents may play socially optimally, but may also play suboptimally, so that coordination would be useful.

Comparing Theorem 23 for choosing only the reciprocation attitudes to Theorem 25 for choosing the coefficients as well, we observe that more freedom of choice allows for a socially suboptimal equilibrium, achieving as little as about half of the optimal social welfare, if the kindness values are very different. This is an important pitfall, which emphasizes the importance of cooperation when more freedom and power lies at our disposal. Like Churchill said⁷: “Where there is great power there is great responsibility”.

The analysis also relates to some real-life phenomena. Our results regarding maximizing utility and social welfare show why in life, if acting is not too hard, then following the example of the kindest makes the individuals and the society thrive, which has already been observed [22]. Since being polite usually consists of words and simple gestures, and is therefore quite easy for many people, this explains why people choose this strategy with experience, becoming more polite, as is indeed observed [23]. In diplomacy, such as in Example 9, these results predict that diplomats will be polite to each other, since this does not take much effort. Being polite benefits the individual and the society by making people feel better easily.

We show that if the agents in some subset do not reciprocate at all, while the rest have perfect reciprocation to the neighborhood, then the actions will converge to the average of the kindness values of the non-reciprocal agents. This is a formal way to say that steadfast individuals can set the eventual group behavior at will, if the rest are not stable but follow others' behavior.

The current results inspire considering the games of choosing both r, r' or choosing the kindness. Changing kindness is less reasonable, since this seems a very basic quality of an agent, harder to change than how she reacts to others.

Many interesting directions for further research exist:

1. Modeling changes in the reciprocity coefficients, attitudes, or β s during the interaction and not only before it starts.
2. Modeling probabilistic reaction.
3. Looking how the manager can really influence the behavior of the agents.
4. Real agents often join and leave the interaction dynamically. For example, people

⁷This quote dates back to the French National Convention, 08/05/1793.

get born and immigrate to a country, some die and emigrate. Therefore, dynamic interaction is very interesting.

5. Another important question is how influential various agents are on the other agents.
6. We used others' research, based on real data, as a basis for the model. Therefore, verifying the model on relevant data, like the arms race actions, would be interesting.

Analyzing the strategic reciprocation analytically provides behavioral advice, predicts and explains reciprocation phenomena. It lays the foundation for more detailed modeling and analysis of reciprocation, required to even better anticipate and improve the individual utilities and the social welfare. Analyzing a single reciprocal interaction also serves as a building block for analyzing effort sharing between several such interactions, which we undertake in the following chapter.

REFERENCES

- [1] G. Polevoy, M. de Weerd, and C. Jonker, *The game of reciprocation habits*, in *Proceedings of the 2016 European Conference on Artificial Intelligence*, Frontiers in Artificial Intelligence and Applications, Vol. Volume 285: ECAI 2016 (2016) pp. 417–425.
- [2] A. Falk and U. Fischbacher, *A theory of reciprocity*, *Games and Economic Behavior* **54**, 293 (2006).
- [3] E. Fehr and S. Gächter, *Fairness and retaliation: The economics of reciprocity*, *Journal of Economic Perspectives* **14**, 159 (2000).
- [4] W. Güth, R. Schmittberger, and B. Schwarze, *An experimental analysis of ultimatum bargaining*, *Journal of Economic Behavior & Organization* **3**, 367 (1982).
- [5] W. J. Dixon, *Reciprocity in United States-Soviet relations: Multiple symmetry or issue linkage?* *American Journal of Political Science* **30**, pp. 421 (1986).
- [6] M. D. Ward, *Modeling the USA-USSR arms race*, *Transactions of The Society for Modeling and Simulation International* **43**, 196 (1984).
- [7] J. Gottman, C. Swanson, and J. Murray, *The mathematics of marital conflict: Dynamic mathematical nonlinear modeling of newlywed marital interaction*, *Journal of Family Psychology* **13**, 3 (1999).
- [8] H. Spencer, *Why reciprocation is key to building business relationships*, *Business 2 Community* (2012).
- [9] R. Axelrod, *The emergence of cooperation among egoists*, *American Political Science Review* **75**, 306 (1981).
- [10] R. Axelrod, *The evolution of cooperation*, Basic books (Basic Books, 1984).

- [11] R. L. Trivers, *The evolution of reciprocal altruism*, *The Quarterly Review of Biology* **46**, 35 (1971).
- [12] E. Fehr, U. Fischbacher, and S. Gächter, *Strong reciprocity, human cooperation, and the enforcement of social norms*, *Human Nature* **13**, 1 (2002).
- [13] J. C. Cox, D. Friedman, and S. Gjerstad, *A tractable model of reciprocity and fairness*, *Games and Economic Behavior* **59**, 17 (2007).
- [14] M. Rabin, *Incorporating fairness into game theory and economics*, *The American Economic Review* **83**, pp. 1281 (1993).
- [15] W. Mastenbroek, *Onderhandelen* (Het Spectrum, 1992).
- [16] D. Kahneman, *Thinking, Fast and Slow* (Farrar, Straus and Giroux, 2011).
- [17] A. Rubinstein, *Modeling Bounded Rationality*, 1st ed., Vol. 1 (The MIT Press, 1997).
- [18] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*, MIT Press Books, Vol. 1 (The MIT Press, 1994).
- [19] D. Abreu and A. Rubinstein, *The structure of Nash equilibrium in repeated games with finite automata*, *Econometrica* **56**, 1259 (1988).
- [20] A. Rubinstein, *Finite automata play the repeated prisoner's dilemma*, *Journal of Economic Theory* **39**, 83 (1986).
- [21] K. G. Binmore and L. Samuelson, *Evolutionary stability in repeated games played by finite automata*, *Journal of Economic Theory* **57**, 278 (1992).
- [22] R. Axelrod and W. Hamilton, *The evolution of cooperation*, *Science* **211**, 1390 (1981), <http://www.sciencemag.org/content/211/4489/1390.full.pdf>.
- [23] M. R. B. Giovanna Axia, *Linguistic politeness at different age levels*, *Child Development* **56**, 918 (1985).
- [24] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, *Algorithmic Game Theory* (Cambridge University Press, 2007).
- [25] J. Nash, *Non-Cooperative Games*, *The Annals of Mathematics Second Series*, **54**, 286 (1951).
- [26] E. Koutsoupias and C. Papadimitriou, *Worst-case equilibria*, in *16th Annual Symposium on Theoretical Aspects of Computer Science* (Trier, Germany, 1999) pp. 404–413.
- [27] E. Anshelevich, A. DasGupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden, *The price of stability for network design with fair cost allocation*, in *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on* (2004) pp. 295–304.
- [28] V. A. Ilyin and E. G. Poznyak, *Fundamentals of mathematical analysis (Osnovi matameticheskogo analiza), book 1*, 7th ed. (Physical and Mathematical Literature, 2005) in Russian.

6

RECIPROCATION EFFORT GAMES

Commerce, trade and exchange make other people more valuable alive than dead, and mean that people try to anticipate what the other guy needs and wants. It engages the mechanisms of reciprocal altruism, as the evolutionary biologists call it, as opposed to raw dominance.

Steven Pinker

Consider people dividing their time and effort between friends, interest clubs, and reading seminars. Consider computers dividing efforts between participating in several network calculations. These are all reciprocal interactions, and these processes eventually determine the utilities of the agents from these interactions. We prove that each of these processes converge, and using the converged values, we determine existence and efficiency of Nash equilibria of the game of allocating effort to such projects. We show that when no minimum effort is required to receive reciprocation, an equilibrium always exists, and if acting is either easy to everyone, or hard to everyone, then every equilibrium is socially optimal. If a minimal effort is needed to participate, we prove that not contributing at all is an equilibrium, and for two agents, also a socially optimal equilibrium can be found.

1. INTRODUCTION

In the last two chapters, we modeled and analyzed a *reciprocation process* and strategic approach to this process. Now, we finally combine this to analyzing *dividing effort between several such processes*.

In many real-world situations people invest effort in several interactions, where they have an effort budget they may use, such as in discretionary daily activities [1], daily communication between school pupils, sharing files over networks, or in business cooperation. In such an interaction, people tend to reciprocate, i.e., react on the past actions of others (sometimes, only if a certain minimum effort is invested) [2–4]. Interaction is an important sort of common projects, different from the classical value-creating projects. The classical projects obtain a value that is typically divided based on the individual investments, such as contributions to online communities [5], Wikipedia [6], political campaigns [7], paper co-authorship [8]. Interaction projects, where mutual actions are the main aspect, are also frequent in life, e.g., at daily communication between school pupils, when sharing files over networks, or in business cooperation.

In order to recommend how to divide one's limited efforts efficiently, we aim to predict stable strategies for these settings and estimate their efficiency. We study settings with and without a threshold for effort required to participate.

Dividing a budget of effort is studied in so-called shared effort games [10]. In these games players contribute to various projects, and given their contributions, each project attains a value, which is subsequently divided between the contributors. In order to support decisions regarding individually and publicly good stable strategy profiles in these games, the social welfare (total utility) of strategy profiles is important, and in particular of Nash equilibria (NE). For this, the price of anarchy (PoA) [11, 12] and stability (PoS) [13, 14] are the most famous efficiency measures of Nash equilibria. The social welfare of a strategy profile is the total utility in this profile. The price of anarchy is the ratio of the least social welfare in an equilibrium to the optimal social welfare, and the price of stability is the ratio of the social welfare in a best NE to the optimal social welfare.

Bachrach, Syrgkanis and Vojnović [10] bind the price of anarchy in shared effort games, but only when a player obtains at least a constant share of her marginal contribution to the project's value, which does not hold for a positive threshold. Chapter 2 analyzes the Nash equilibria, and price of anarchy and stability also in the case with a threshold. When the threshold is equal to the highest contribution, such shared effort games are equivalent to all-pay auctions. In all-pay auctions, only one contributor benefits from the project. Its equilibria are analyzed by Baye, Kovenock and de Vries [15] and many others. A famous example of an all-pay auction is the colonel Blotto game with two players [16], where the players spread their forces between several battlefields, winning a battle if having allocating it more forces than the opponent. The relative number of won battles determines the player's utility. Anshelevich and Hoefer [17] study an undirected graph, where nodes are the players that divide effort between the incident edges, which represent projects that reward the incident players equally. They especially concentrate on minimum-effort projects. The work proves the existence of an NE and shows that the price of anarchy is at most 2. To conclude, the models of [10] and Section 2 of Chapter 2 are the only general ones, the model of Section 2 being the only relevant analysis with a threshold. Therefore, we base the model of investing effort in reciprocal interac-

tions on the model of Section 2 of Chapter 2. In most existing work on contributing to projects, the value of a project is directly defined based on the contributions, such as in contributions to online communities [5], Wikipedia [6], political campaigns [7], and paper co-authorship [8]. In this paper we model the interaction within the projects and the result thereof explicitly.

Existing models of reciprocation (sometimes repeated) often consider why reciprocation has emerged. Other models assume reciprocation and study how reciprocation influences interactions. The following works consider the emergence of reciprocation. Some works, such as the famous works by Axelrod [18, 19], study and motivate direct evolution of reciprocal behavior. Others consider a more elaborate evolution, like Bicchieri's work on norm emergence [20, Chapter 6] or Van Segbroeck et al. [21], who study the evolution of fairness, which, in turn, engenders reciprocation. The classical work of Trivers [22] takes a detailed biological approach, describing how altruism-related emotions like guilt and suspicion have evolved. There are also other approaches to the nature of reciprocation, such as the *strong reciprocation*, where intricate social or emotional drives are considered [23]. Some other works take reciprocal behavior to be given, and analyze the development of certain interactions, modeling them as appropriate games [2, 24–26]. However, there was no model of reciprocal actions of various extent, before the model from Chapter 4.

The approach In Chapter 4, partially published in [27], we assume agents reciprocate and formally model and analyze lengthy repeated reciprocation. That chapter analyzes convergence in this model. This model is mainly inspired by similar models used to analyze arms race [28, 29] and spouses' interaction [30]. The model defines an action on an agent as a convex combination¹ between one's own last action, the considered other agent's and all the other agents' last actions. This is called the *floating* reciprocation attitude. The idea of humans behaving according to a convex combination appears also in other contexts, such as modeling altruism in several papers, like [31–34], and Chapter iii.2 of [35]. In Chapter 5, the utility of a reciprocating agent is defined as what an agent gets minus the effort her actions incur. That chapter analyzes strategical setting of reciprocation habits to maximize utility.

Summarizing, existing work on sharing effort concentrates on the case of equal sharing of linear project's value between all agents who contribute above the threshold, and this does not include the case when each project constitutes a network of reciprocal agents. Therefore, the area of dividing effort between reciprocal interactions, such as meeting friends or mutual advising, is yet to be modeled and analyzed. We now aim to model lengthy interactions with actions of various extent, as in [27].

We want to predict dividing efforts between reciprocal interactions to be able to recommend an efficient way to do it. We consider contributing to several interactions, and analyze the Nash equilibria and their efficiency. We also model first dividing effort between several interactions and then deciding how to reciprocate. In this case, we analyze the subgame perfect equilibria. We consider only pure equilibria throughout the chapter, even when we do not mention this explicitly. Since the strategies include all the ways to divide budget among the interactions, the set of pure strategies is already uncountably infinite.

¹A combination is convex if it has nonnegative weights that sum up to 1.

We define a model of several reciprocal interactions in Section 2, taking into account the fact that the maximum capacity of an agent to invest in the various projects may be not enough to satisfy the definition of every reciprocal interaction. This forces the agent to curb her investments in some interactions, complicating the process. We prove convergence of such an interaction with curbing in Section 3.

Next, to have an exact formula for the limits, we assume for the rest of the chapter that the maximum capacity for an agent to invest under the requirement of the reciprocation model is big enough, so no curbing is required. This assumption allows using the formula for the common limit (4.11) on page 98 from Chapter 4 and analyzing the situation. We model the minimum contribution threshold in three different ways, defining three games: one without a threshold, where an agent's utility from an interaction is simply her utility from the respective interaction, and two other ones with a threshold. In the second game, only those who contribute above the threshold obtain the utility from the interaction, but all the agents participate in all the interactions. In the third game, those who are below the threshold in an interaction, are not even allowed to participate in the respective interaction. The model to use depends on the concrete situation at hand. We characterize the equilibria of the first game, and find their efficiency in Section 4. Then, we analyze the second and third games in Section 5 and Section 6, respectively.

We finally define an extensive game that models dividing the capacity between the interaction and subsequently deciding how to behave in each ensuing interaction (choosing habits as in Chapter 5). We provide sufficient conditions for the existence of a subgame perfect equilibrium in this game. Section 8 concludes and outlines new interesting research directions.

2. MODEL

This section models dividing effort between reciprocal interactions. Inspired by the shared effort games models from [10] and Chapter 2, and adopting the reciprocation model from Chapter 4, we define a *reciprocation effort game*. First, we define a reciprocal process and the agents' utilities in this process. Next, we define a reciprocation effort game, comprising several such processes. The game includes dividing one's budget between these processes and curbing the reciprocal actions, if an agent's effort does not suffice the requirements of the reciprocation. Afterwards, we define two thresholded variations on this game, to model the minimal required investment. We conclude by defining a two-level extended game, which models dividing effort between interactions and then choosing the interaction habits.

Remark 22 (Notation). *As we mentioned in the beginning of the dissertation, there exist some notational differences between the comprising publications. Since letters x are used in this chapter to denote contribution, analogically to Chapter 2 and Chapter 3, we use another notation for actions, namely act . This is different from the notation of x for actions, employed Chapter 4 and Chapter 5.*

We begin by defining the reciprocation model from Chapter 4. Given agents $N = \{1, \dots, n\}$, at any time $t \in T \triangleq \{0, 1, 2, \dots\}$, every agent acts on any other agent. Denote the

weight of the action by agent $i \in N$ on another agent $j \in N$ at moment t by $\text{act}_{i,j}(t): T \rightarrow \mathbb{R}$. Since only the weight of an action is relevant, we usually write “action” while referring to its weight. For example, when interacting with colleagues, the weights of the actions of helping, nothing, or insulting are in the decreasing order.

In order to define how agents reciprocate, we need the following notation. We denote the total received action from all the other agents at time t by $\text{got}_i(t): T \rightarrow \mathbb{R}$; formally, $\text{got}_i(t) \triangleq \sum_{j \in N} \text{act}_{j,i}(t)$. The kindness of agent i is denoted by $k_i \in \mathbb{R}$. Kindness models constant inherent inclination to help others; in particular, it determines the first action of an agent, before others have acted. We model agent i 's inclination to mimic another agent's action and the actions of all the other participants in the project by reciprocation coefficients $r_i \in [0, 1]$ and $r'_i \in [0, 1]$ respectively, both staying constant for all interactions. r_i is the fraction of $\text{act}_{i,j}(t)$ that is determined by the previous action of j upon i , and r'_i is the fraction that is determined by $\frac{1}{n-1}$ th of the total contribution to i from all the other agents at the previous time. Consequently, $r_i + r'_i \leq 1$.

We now define the actions. First, there is nothing to react to, so the kindness determines the action: $\text{act}_{i,j}(0) \triangleq k_i$.

Definition 22. *Agent i 's action is a weighted average of her own last action (inertia), of that of the other agent j (direct reaction) and of the total action of the other agents divided over all the others (social reaction):*

$$\text{act}_{i,j}(t) \triangleq (1 - r_i - r'_i) \cdot \text{act}_{i,j}(t-1) + r_i \cdot \text{act}_{j,i}(t-1) + r'_i \cdot \frac{\text{got}_i(t-1)}{n-1}.$$

An agent's utility from a given reciprocation project at a given moment is the action one receives minus effort to act, following Chapter 5. This is classical (see, for example, the quasilinear preferences of auction theory [36, Chapter 9.3]). Formally, define the utility of agent i at time t , $u_{i,t}: \mathbb{R}^{\text{deg}(i)} \times \mathbb{R}^{\text{deg}(i)} \rightarrow \mathbb{R}$, as

$$u_{i,t} \left(\text{act}_{i,j}(t), \text{act}_{j,i}(t) \right) \triangleq \sum_{j \in N} \text{act}_{j,i}(t) - \beta_i \sum_{j \in N} \text{act}_{i,j}(t)$$

where the constant $\beta_i \in \mathbb{R}$ is the relative importance of the performed actions for i 's utility. The personal price of acting is higher, equal or lower than of receiving an action, if β_i is bigger, equal or smaller than 1, respectively. The minus in front of i 's actions subtracts the effort of acting from one's utility (unless β_i is negative, where that is added). Since the presence of negative actions would mess up this logic, by contributing to the utility, we assume that actions are always non-negative, which occurs if and only if all kindness values are non-negative. We can have negative influence, but we assume having added large enough a constant to all the actions, to avoid negative actions.

Every such interaction converges, as shown in Chapter 4. To model the utility in the long run, we define the asymptotic utility, or just the *utility*, of agent i , as the limit of her utilities as the time approaches infinity. In formulas, $u_i: (\mathbb{R}^{\text{deg}(i)})^\infty \times (\mathbb{R}^{\text{deg}(i)})^\infty \rightarrow \mathbb{R}$, as $u_i \left(\bigcup_{t'=0}^\infty \{ \text{act}_{i,j}(t), \text{act}_{j,i}(t) \} \right) \triangleq \lim_{t \rightarrow \infty} u_{i,t} \left(\text{act}_{i,j}(t), \text{act}_{j,i}(t) \right)$. This is the utility we consider in the chapter. When the parameters in the parentheses are clear from the context,

we may omit them. This definition of the utility of a process is equivalent to the discounted sum of utilities when the discounting is slow enough, as we prove in Section 11.1 of Chapter 5.

We are now ready to define a *reciprocation effort game*. Consider n players $N = \{1, \dots, n\}$ and m interactions $\Omega = \{1, 2, \dots, m\}$. Intuitively, the strategies are the contributions, and they determine the kindness values. Definition 22 describes the reciprocal processes. If the sum of an agent's actions exceeds the budgets at some time, then the actions are curbed as defined below. Interaction determines the utilities.

The reciprocation coefficients of the players, r_i, r'_i , and the β_i are given. The contribution of player $i \in N$ to interaction project $j \in \Omega$ at time $t \in T$ is defined as the sum of her actions in that interaction at that time, i.e. $x_j^i(t) \triangleq \sum_{j \in N \setminus \{i\}} \text{act}_{i,j}(t)$. Player i 's kindness k_i at the reciprocal interaction called project ω is determined by her contribution to that interaction at time zero, called just "the contribution", divided by the number of other agents who participate in the interaction at ω , to act on everyone else. Therefore, the sum of all the actions of agent i at time $t = 0$ is equal to her contributions to all the reciprocation projects, which are bounded by her budget b_i . This completes the definition of reciprocal interactions Ω .

An agent invests something in the beginning of a reciprocation, and from that time on it aims to proceed according to Definition 22. We naturally require that not only the sum of the contributions at $t = 0$, but also the sum of the contributions at any time $t > 0$ is within the acting agent's budget. Each player i has a normal budget $b_i > 0$ (or just a budget) to contribute from at $t = 0$ and an extended budget $B_i \geq b_i$ that can be used when the actions are required by the reciprocation process at $t > 0$, perhaps resulting in a higher summarized contribution than the voluntarily chosen at $t = 0$. We differentiate between these two budgets, since the need to reciprocate urges people to act [23]. Given this extended budget limitation on the actions at any time $t > 0$, if the total contribution of an agent required for the actions according to the reciprocation model exceeds her budget, then they will be curbed as follows. For every project j , agent i has curbing coefficient c_j^i , expressing how much she values this project. If, at time $t > 0$, the interaction model predicts contribution y_j^i to project j , then the actual contribution $x_j^i(t)$ will be the proportional normalization of y_j^i , meaning that

$$x_j^i(t) \triangleq B_i \cdot \frac{c_j^i y_j^i}{\sum_{l=1, \dots, n} c_l^i y_l^i}.$$

Note that $\sum_{j \in \Omega} x_j^i(t) = B_i$, as required in a normalization to B_i . The normalization coefficients are proportional to the curbing coefficients c_1^i, \dots, c_m^i . When these coefficients are all equal, we get the usual normalization, which is

$$x_j^i(t) \triangleq B_i \cdot \frac{y_j^i}{\sum_{l=1, \dots, n} y_l^i}.$$

The so curbed contribution to a project is divided between the actions on the other agents proportionally to the actions predicted by the reciprocation model. The idea

of a limited budget and of curbing actions to fit the budget comes from the time-displacement hypothesis [1, 37], which discusses taking time from one activity for the other. These works also mention that different activities get curbed in different rates, generally speaking: the discretionary activities like hobbies of meeting friends get curbed much more than the other ones, like work.

Formally, the strategy space of player i (i.e., the set of her possible actions) consists of her contributions (at time zero), determining her kindness values, and of her curbing coefficients: $\{x^i = (x_\omega^i)_{\omega \in \Omega} \in \mathbb{R}_+^{|\Omega|} \mid \sum_{\omega \in \Omega} x_\omega^i \leq b_i\} \cup \{c_1^i, \dots, c_m^i \mid c_j^i \geq 0\}$. The contributions at positive times will not be discussed, so saying “contribution” we always mean the contribution at $t = 0$. The utility agent i obtains from participating in project ω is defined as her asymptotic utility from the interaction with the other agents in that project. The utility $u_i(x)$ of a player $i \in N$ is defined to be the sum of the utilities it obtains from the various projects, completing the definition of a *reciprocation effort game*.

The convergence of the normal reciprocation is proven in Section 5 of Chapter 4, and we prove the convergence of the curbed reciprocation in Section 3. Afterwards, to be able to analyze the NE and efficiency of the problem, we simplify it by assuming that all the extended budgets B_i are big enough to prevent curbing at all. Therefore, the curbing coefficients become irrelevant, and we assume that the strategies are only $\{x^i = (x_\omega^i)_{\omega \in \Omega} \in \mathbb{R}_+^{|\Omega|} \mid \sum_{\omega \in \Omega} x_\omega^i \leq b_i\}$. Please note that an agent does not have to use up all her budget, so that the inequality $\sum_{\omega \in \Omega} x_\omega^i \leq b_i$ may be strict. The strategies of all the players except i are denoted x^{-i} . We denote the vector of all the contributions by $x = (x_\omega^i)_{\omega \in \Omega}^{i \in N}$.

We now define two variations on a shared effort game with reciprocation. First, following Section 2 of Chapter 2, we define a θ -*sharing mechanism*. This is relevant to many applications, like a minimum invested effort to be considered a coauthor, or a minimum effort to master a technology before working with it. Define, for every $\theta \in [0, 1]$, the players who get a share from project ω

$$N_\omega^\theta \triangleq \left\{ i \in N \mid x_\omega^i \geq \theta \cdot \max_{j \in N} x_\omega^j \right\},$$

which are those who bid at least θ fraction of the maximum contribution to ω .

We now define a *thresholded reciprocation effort game*, as a reciprocation effort game, where only the agents in N_ω^θ obtain the above utility from project ω ; other agents obtain nothing from that project. All agents interact, but only those who invest above the threshold obtain the revenue. The second variation we define is an *exclusive thresholded reciprocation effort game*, as a reciprocation effort game, where exclusively the agents in N_ω^θ interact. Others do not obtain utility and do not even interact. If an agent ends up participating alone at a project, he obtains zero utility from that project, since no interaction occurs.

The thresholded games model the fact that an agent needs to invest at least some minimum effort to have a fruitful interaction. Thresholded reciprocation effort games model situations where every agent interacts with everyone else, but she may keep the obtained utility only if she contributes above the threshold. This may happen in file sharing, or in any environment where agent accumulate utility, but may sometimes not

Reciprocal Interactions:

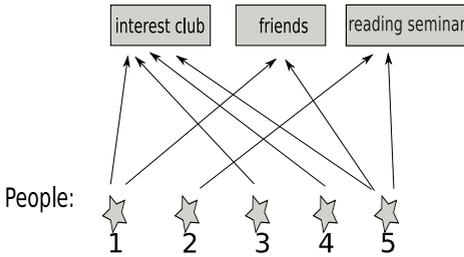


Figure 6.1: People divide their own effort between various interactions.

be able to use it up, such as in accumulating reputation in some area. Exclusive thresholded reciprocation effort games model situations when merely joining an interaction requires contributing enough, like the initial effort it takes to learn the required technology to contribute to Wikipedia.

We now give a concrete example of the model.

Example 14. Consider people choosing between the following activities for free time: Going to interest club, meeting friends, or going to a scientific reading seminar. This is illustrated in Figure 6.1. Each of these projects can be modeled as interaction, in the following way. Going to an interest club involves communication, which is interaction. Meeting friends means obviously interacting. At a reading seminar, people help each other understand science, though they may also influence each other in other ways, such as encouraging or insulting. A person divides her limited effort budget between these activities, in attempt to get as much as possible. Some people find helping others hard, while others do not, which is modeled by large or little β . If every contribution counts, without thresholds, we model the situation as a reciprocation effort game. If a minimal effort is required, and even mere participation requires overcoming the threshold, model this as an exclusive thresholded reciprocation effort game. Analyzing the games for existence of equilibria allows predicting whether a stable state exists, and looking into the efficiency of the equilibria provides some insight into how efficient these stable states are and whether the agents would benefit from regulation.

Finally, we model first dividing the budget and then deciding on how to interact in each given interaction as an *extended reciprocation effort game*. This is an extensive game with simultaneous moves, where at the root all the agents choose how to divide their budgets, and at the obtained subgames they choose the reciprocation coefficients r_i , to determine how to reciprocate. This models that once an agent sees what are the kindness values of the other agents, she chooses, how to reciprocate. In the example above, once a person has allotted some amount of effort to an activity and she sees how the others act, she decides on how to react to their actions from that time on. We do not allow a threshold in the extensive games, since this may leave us with only two agents interacting, and then, there is no convergence if $r_1 = r_2 = 1$.

We analyze convergence in the next section, which is basic for all the variations of the defined games.

3. GENERAL CONVERGENCE

In this section, we consider dividing effort between reciprocal projects, with or without threshold. In order to study asymptotic behavior, we ask whether the actions in all interactions converge, as time approaches infinity. The convergence of the normal reciprocation is proven in Chapter 4, and we now prove the convergence of the curbed reciprocation. Consider the undirected interaction graph $G = (N, E)$ of an interaction project, such that agent i can act on j and vice versa if and only if $(i, j) \in E$. In our model, we assume that this graph is a clique, but this is not necessary for the following theorem.

Theorem 27. *Consider dividing effort between reciprocal interactions, where every interaction has some connected interaction graph, and for all agents i , $r'_i > 0$. At every interaction, if there exists a cycle of an odd length in the interaction graph, or or at least one agent i has $r_i + r'_i < 1$, then, for all pairs of agents $i \neq j$, the limit $L_{i,j}$ exists.*

This result holds for any curbing, not only the one defined in Section 2.

In our model, we assume a completely connected graph, so if at least 3 agents interact, we have an odd cycle, namely a triangle. Therefore, then we only need to assume that for all agents i , $r'_i > 0$.

The proof expresses the dynamics as matrix multiplication. Without the curbing, the convergence of the powers of matrices is proven using the Perron-Frobenius theorem. Keeping convergence when curbing can occur requires the following definition and lemma.

Definition 23. *We remind that a square non-negative matrix A is called primitive, if there exists a positive l , such that $A^l > 0$ (see [38, Definition 1.1]).*

The following lemma, used to prove the theorem, has a value of its own as well. Given a convergent sequence of primitive matrices, the lemma shows that arbitrary squeezing the matrices keeps the convergence.

Lemma 10. *Given a vector $p(0) \in \mathbb{R}^d$, a primitive matrix $A \in \mathbb{R}^{d^2}$, such that $\lim_{t \rightarrow \infty} A^t$ exists, and a sequence of diagonal matrices $\{D(t)\}_{t=0}^{\infty}$, $D(t) = \text{diag}(\lambda_1(t), \dots, \lambda_d(t))$, where each $\lambda_i(t) \in (0, 1]$, define the sequence $\{p(t)\}_{t=0}^{\infty}$ by $p(t) \triangleq D(t)AD(t-1)A \dots D(1)Ap(0)$. Then, $\lim_{t \rightarrow \infty} p(t)$ exists.*

A particular case can be strengthened and be proven easily.

Remark 23. *If all the $D(t)$ matrices are scalar, meaning that for all t , $\lambda_1(t) = \lambda_2(t) = \dots = \lambda_n(t)$, then the lemma holds even without requiring that A is primitive. This simply follows by the fact that scalar matrices commute with any matrices, and therefore,*

$$p(t) = \left(\prod_{t'=1}^t D(t') \right) A^t p(0).$$

The scalar matrix $(\prod_{t'=1}^t D(t'))$ has on its diagonal the product $\prod_{t'=1}^t \lambda_i(t)$. This product forms a nondecreasing and bounded from below sequence, and therefore, it converges. Since we assume that $\{A^t\}_{t \in \mathbb{N}}$ converges, we conclude that their matrix product, being a sum of products of the elements, converges as well.

We now prove the original lemma.

Proof. Assume to the contrary, that $\{p(t)\}$ diverges. Define the sequence $\{p'(t)\}_{t=0}^\infty$ by $p'(t) \triangleq A^t p(0)$. Since $\{p(t)\}$ diverges and $\{p'(t)\}$ converges, they differ at some point, intuitively speaking. We now formalize this argument. Since $\{p(t)\}$ diverges and the space is complete, it is not a Cauchy sequence, and so there exists a positive ϵ , such that for each $N > 0$ there exist $n, m > N$, such that $|p(n) - p(m)| > \epsilon$. Since $\{p'(t)\}$ converges, it is a Cauchy sequence, so there exists $N > 0$, such that for all $n, m > N$ we have $|p'(n) - p'(m)| < \epsilon/2$. If $|p(n) - p(m)| > \epsilon$ and $|p'(n) - p'(m)| < \epsilon/2$, we cannot both have $|p(n) - p'(n)| < \epsilon/4$ and $|p(m) - p'(m)| < \epsilon/4$. Therefore, for some integer l , $|p(l) - p'(l)| > \delta$, for some $\delta > 0$, depending solely on ϵ . Since the product defining $p(l)$ is like that of $p'(l)$, but with more $D(t)$ matrices, and $D(t) = \text{diag}(\lambda_1(t), \dots, \lambda_d(t))$, where each $\lambda_i(t) \in (0, 1]$, we have $0 \leq p(l) \leq p'(l)$. Remembering this, and that matrix A is primitive, thereby propagating a change of an entry to every entry, we can choose l such that every coordinate of $p(l)$ will be at most α fraction of the corresponding coordinate of $p'(l)$, for some $\alpha < 1$. The α can be made to depend solely on ϵ , because of the boundedness of all the relevant vectors. So, we have $p(l) \leq \alpha A^l p(0)$.

By reiterating the same argument with $p'_1(t) \triangleq A^t p(l)$ and $p_1(t) \triangleq p(t+l)$, we find $l_1 > 0$, such that $p_1(l_1) \leq \alpha A^{l_1} p(l)$. Thus, $p(l_1+l) = p_1(l_1) \leq \alpha A^{l_1} p(l) \leq \alpha A^{l_1} \alpha A^l p(0) = \alpha^2 A^{l_1+l} p(0)$.

Continuing in this manner, and remembering the boundedness of $\{A^t p(0)\}$, which stems from its convergence, we prove that $\{p(t)\}$ converges to zero, contradictory to the assumption.² □

We are now ready to prove the theorem.

Proof. We consider one interaction and extend the proof of Theorem 20 on page 98. We recapitulate the used properties from there, to stay self-contained.

Basically, we express how each action depends on the actions in the previous time in a matrix, and prove the theorem by applying the famous Perron–Frobenius theorem [38, Theorem 1.1, 1.2] to this matrix, using the above lemma to cover the case of curbing actions. We define the dynamics matrix $A \in \mathbb{R}_+^{|E| \times |E|}$ as

$$A((i, j), (k, l)) \triangleq \begin{cases} (1 - r_i - r'_i) & k = i, l = j; \\ r_i + r'_i \frac{1}{|N^+(i)|} & k = j, l = i; \\ r'_i \frac{1}{|N^+(i)|} & k \neq j, l = i; \\ 0 & \text{otherwise.} \end{cases} \quad (6.1)$$

Assume that for each time $t \in T$, the column vector $\vec{p}(t) \in \mathbb{R}_+^{|E|}$ describes the actions at time t , in the sense that its (i, j) th coordinate contains $\text{act}_{i,j}(t)$ (for $(i, j) \in E$). Then, $\vec{p}(t+1) = D(t)A\vec{p}(t)$, where $D(t)$ is the diagonal matrix, describing the normalization, thus $D(t) = \text{diag}(\lambda_1, \dots, \lambda_{|E|})$. We call $\vec{p}(t)$ an action vector. Initially, $\vec{p}_{(i,j)}(0) = k_i$.

Further, we shall need to use the Perron-Frobenius theorem for primitive matrices. We now prepare to use it, and first we show that A is primitive. In the proof of Theorem 20

²The actual limit does not have to be zero; zero is just a result from the contradictory assumption.

from in Chapter 4, it is shown that A is irreducible and aperiodic, and therefore primitive by [38, Theorem 1.4]. Since the sum of every row is 1, the spectral radius is 1.

According to the Perron–Frobenius theorem for primitive matrices [38, Theorem 1.1], the absolute values of all the eigenvalues except one eigenvalue of 1 are strictly less than 1. The eigenvalue 1 has unique right and left eigenvectors, up to a constant factor. Both these eigenvectors are strictly positive. Therefore, [38, Theorem 1.2] implies that $\lim_{t \rightarrow \infty} A^t = \vec{1} \vec{v}'$, where \vec{v}' is the left eigenvector of the value 1, normalized such that $\vec{v}' \vec{1} = 1$.

Now, Lemma 10 implies that $L_{i,j}$ exists. □

This section analyzed convergence in the general case. From the next section on, we analyze a subcase of the situation, which lends itself to precise analysis and allows predicting the interaction and recommending how to divide one's effort. We consider three variations in the role of a contribution threshold, and then we analyze the extended version of the situation.

4. RECIPROCATION EFFORT GAME

This section studies existence and efficiency of Nash equilibria in the setting without a threshold. This analysis allows predicting behavior and recommending on the best practice.

From here and for the rest of the chapter, we assume that all the extended budgets B_i are big enough to prevent curbing.

We first completely analyze existence of NE, and then we find all the prices of anarchy and stability. This theorem characterizes the existence of equilibria.

Theorem 28. *Assume that for any agent i , $r'_i > 0$, and in addition, either $n > 2$ or $r_1 + r'_1 + r_2 + r'_2 < 2$. The set of all the NE is exactly all the strategy profiles where every agent with $\beta_i < 1$ somehow divides all her budget among the projects $\{1, \dots, m\}$, and every agent with $\beta_i > 1$ contributes nothing. These strategies are also dominant.*

Proof. Consider an arbitrary player l , and let her strategy (her contributions³) be $x^l = (x_1^l, \dots, x_m^l)$. According to Formula (4.11) on page 98, her utility from this strategy is

$$(n-1)(1-\beta_l) \left(\frac{\left(\frac{1}{r_l+r'_l} \cdot (x_1^l + \dots + x_m^l) \right)}{\sum_{i \in N} \left(\frac{1}{r_i+r'_i} \right)} + C \right),$$

for C that is independent of l 's strategy. Therefore, if $\beta_l < 1$, then l 's strategy is a best response to others' strategies if and only if l arbitrarily divides all her budget among the projects $\{1, \dots, m\}$. On the other hand, if $\beta_l > 1$, then a strategy is a best response if and only if all the contributions are zero. This is true for every agent l , proving that this is an NE. Since each agent is independent of the others, these strategies are also dominant. □

We immediately conclude the following about the existence of a NE.

³Contributions by default refer to the contributions at time zero.

Corollary 5. *There always exists an NE.*

The possible variations in an NE profile are what the agents with $\beta = 1$ do. This is important for analyzing the efficiency of the NE.⁴ To analyze efficiency, we define: $N^< \triangleq \{i \in N : \beta_i < 1\}$, $N^{\leq} \triangleq \{i \in N : \beta_i \leq 1\}$, $N^= \triangleq \{i \in N : \beta_i = 1\}$. We now analyze the efficiency of the most and the least efficient equilibria, comparing their social welfare to the maximum possible social welfare.

Proposition 22. *Under the assumptions of the theorem, if $(n > \sum_{i \in N} \beta_i)$, we have $\text{PoA} = \frac{\sum_{i \in N^<} \left(\frac{1}{r_i + r'_i} \cdot b_i \right)}{\sum_{i \in N} \left(\frac{1}{r_i + r'_i} \cdot b_i \right)}$, and the PoS is given by the same expression, where we use N^{\leq} instead of $N^<$. If $(n = \sum_{i \in N} \beta_i)$, we have $\text{PoA} = \text{PoS} = 1$. If $(n < \sum_{i \in N} \beta_i)$, then:*

If $N^< \neq \emptyset$, then we have $\text{PoA} = \text{PoS} = -\infty$.

If $N^< = \emptyset$, but $N^{\leq} \neq \emptyset$, then $\text{PoA} = -\infty$, but $\text{PoS} = 1$.

If $N^{\leq} = \emptyset$, then $\text{PoA} = \text{PoS} = 1$.

Proof. The possible social welfare values that an NE can achieve are exactly

$$(n-1)(n - \sum_{i \in N} \beta_i) \frac{\sum_{i \in N^<} \left(\frac{1}{r_i + r'_i} \cdot b_i \right) + \sum_{i \in N^=} \left(\frac{1}{r_i + r'_i} \cdot x^i \right)}{\sum_{i \in N} \left(\frac{1}{r_i + r'_i} \right)},$$

where $0 \leq x^i \leq b_i$. The optimum social welfare is

$$(n-1)(n - \sum_{i \in N} \beta_i) \frac{\sum_{i \in N} \left(\frac{1}{r_i + r'_i} \cdot b_i \right)}{\sum_{i \in N} \left(\frac{1}{r_i + r'_i} \right)},$$

if $(n > \sum_{i \in N} \beta_i)$, and 0 otherwise.

Thus, if $(n > \sum_{i \in N} \beta_i)$, we have

$$\text{PoA} = \frac{\sum_{i \in N^<} \left(\frac{1}{r_i + r'_i} \cdot b_i \right)}{\sum_{i \in N} \left(\frac{1}{r_i + r'_i} \cdot b_i \right)}, \text{ and } \text{PoS} = \frac{\sum_{i \in N^+} \left(\frac{1}{r_i + r'_i} \cdot b_i \right)}{\sum_{i \in N} \left(\frac{1}{r_i + r'_i} \cdot b_i \right)}.$$

If $(n = \sum_{i \in N} \beta_i)$, we have $\text{PoA} = \text{PoS} = 1$, since the social welfare is always zero.

If $(n < \sum_{i \in N} \beta_i)$, then we may get negative social welfare, since zero is optimal, while some NE yield a negative social welfare. Concretely, we have the following subcases:

If $N^< \neq \emptyset$, then we have $\text{PoA} = \text{PoS} = -\infty$, because any NE has the social welfare of at most

$$(n-1)(n - \sum_{i \in N} \beta_i) \frac{\sum_{i \in N^<} \left(\frac{1}{r_i + r'_i} \cdot b_i \right)}{\sum_{i \in N} \left(\frac{1}{r_i + r'_i} \right)}.$$

⁴One usually does not consider PoA and PoS for negative utilities, but we do.

If $N^< = \emptyset$, but $N^+ \neq \emptyset$, then $\text{PoA} = -\infty$ but $\text{PoS} = 1$. The reason is that an NE can have the social welfare of because any NE has the social welfare from zero and down to

$$(n-1)(n - \sum_{i \in N} \beta_i) \frac{\sum_{i \in N^+} \left(\frac{1}{r_i + r'_i} \cdot b_i \right)}{\sum_{i \in N} \left(\frac{1}{r_i + r'_i} \right)}.$$

If $N^+ = \emptyset$, then $\text{PoA} = \text{PoS} = 1$, since an NE has the social welfare of zero.

□

In particular, we have shown that if all the agents find acting easy (i.e., all $\beta_i < 1$), or if all agents really do not like acting (i.e., all $\beta_i > 1$), then $\text{PoA} = \text{PoS} = 1$, meaning that any NE is optimum for the society. Intuitively, this is because here, all the agents have similar preferences: either everyone wants to contribute and receive, or no one does. If the average agent finds non acting as important as being acted upon (i.e., $\sum_{i \in N} \beta_i = n$), every equilibrium is trivially optimal, since the agents do not care. We have also shown, that if the average agent finds not contributing more important than receiving (i.e., $\sum_{i \in N} \beta_i > n$), but still $\beta_i < 1$ for some agent i , then $\text{PoA} = \text{PoS} = -\infty$, so any NE is catastrophic to the society. Intuitively, this stems from the differences in the agents' preferences. Finally, we see that if $\sum_{i \in N} \beta_i > n$, some agents have $\beta_i = 1$, but none have $\beta_i < 1$, then $\text{PoA} = -\infty$ but $\text{PoS} = 1$. Here, regulation may help to play the optimal NE.

Theorem 28 implies that if all the projects have $\beta \leq 1$, then any dividing of all the budget in cooperating is always an NE. This is unintuitive, since usually, some groups are more efficient to interact with than some other groups. The reason for this is that the model assumes that all agents always interact at every project $\omega \in \{1, \dots, m\}$, and only their kindness depends on the strategy. The following section analyzes the more realistic model with a threshold.

5. THRESHOLDED RECIPROCATION EFFORT GAME

After analyzing the simpler non-thresholded case, we now turn to analyze the thresholded case, where everyone participates in all the interactions, but a agent receives her utility from reciprocation only if she contributes at least the threshold. In this section, we assume w.l.o.g. that $b_n \geq \dots \geq b_1$.

Before characterizing the existence of NE, we first provide some auxiliary definitions. Then, we use this fairly complex characterization to provide several sufficient conditions for the existence of an NE. We follow the existence results up with several efficiency results.

To characterize the set of all NE, we need several definitions. We say that agent l covers project $\omega \in \{1, \dots, m\}$, if l 's contribution to ω is at least the minimum contribution, required for l to obtain utility from this project (given the threshold).

Recall the famous KNAPSACK problem [39]. There, we are given a knapsack size $S > 0$ and a set of items $\{1, \dots, n\}$, each item attributed with a positive size s_i and a positive value v_i , and we have to choose a subset of all items $E \subseteq \{1, \dots, n\}$, such that $\sum_{i \in E} s_i \leq S$, while maximizing for $\sum_{i \in E} v_i$. KNAPSACK is weakly NP-hard, which intuitively means

that it is hard only is the involved numbers are large. We therefore do not use KNAPSACK as evidence of hardness, but simply as a known problem.

Consider any agent l in a profile, and attribute her with the following KNAPSACK problem, we call l -KNAPSACK. The items are the projects $\{1, \dots, m\}$, project ω 's size defined as the minimum contribution x_ω^l , required for l to cover this project, and ω 's value defined as the total contributions of other agents to ω : $\sum_{i \in N \setminus \{l\}} \left(\frac{1}{r_i + r_i'} \cdot x_\omega^i \right)$. The knapsack size S is defined to be b_l . Then, a strategy of an agent can be looked at as a feasible solution of l -KNAPSACK, if we consider projects that l covers to be chosen. Intuitively, this problem is to choose the best projects for l to cover.

We now characterize the set of all the NE.

Theorem 29. *Assume that for any agent i , $r_i' > 0$, and in addition, either $n > 2$ or $r_1 + r_1' + r_2 + r_2' < 2$. The set of all the NE is exactly all the strategy profiles where the strategy of every agent l with $\beta_l < 1$ is an optimum solution to the appropriate l -KNAPSACK, and all b_l is divided among the covered projects (no requirement, if no project can be covered), and every agent with $\beta_i > 1$ covers no project with a positive contribution.*

Proof. First, consider agent l with $\beta_l < 1$. Since at any covered project j , contribution x_j^l enters the utility from the project as

$$(n-1)(1-\beta_l) \left(\frac{\left(\frac{1}{r_l + r_l'} \cdot (x_j^l) \right)}{\sum_{i \in N} \left(\frac{1}{r_i + r_i'} \right)} \right),$$

there is no difference, to which project to contribute, besides getting above the threshold and using others' contributions as well. Thus, getting as much as possible from various interactions is exactly solving the appropriate l -KNAPSACK optimally and contributing everything to the covered projects only (unless l cannot cover any project, in which case he will obtain zero utility anyway).

Consider now agent p with $\beta_p > 1$. Since covering any project with positive contribution would result in negative utility, she may not cover any project with a positive contribution. This guarantees the maximum possible utility for her. \square

Notice, that this statement coincides with Theorem 28 when all thresholds are zero, as required. KNAPSACK is NP-hard [39], and therefore, finding a best response strategy is NP-hard. This might seem to justify the hardness of deciding where to invest one's effort. However, the existence of a fully polynomial time approximation scheme for KNAPSACK [39] invalidates this justification.

A natural question is whether there always exists an NE. We answer affirmatively for the case where the threshold is not too large, for the case where the budgets are close, and for the case of two agents. We begin with the case of the threshold being at most half.

Proposition 23. *Under the assumptions of the theorem, and assuming also that $\theta \leq 0.5$, there always exists an NE, such that the agents with the largest budget among those with $\beta_l < 1$ contribute equally to all the projects $\{1, \dots, m\}$, each agent with $\beta_l > 1$ does whatever we choose, as long as she covers no project with a positive contribution, and the agents*

with $\beta = 1$ do whatever we choose. Moreover, such an NE can be computed in time $O(nt)$, where t is the time to calculate a best response of an agent.

Proof. We provide the following algorithm, which arrives at such an NE. We begin by letting every agent with $\beta_l < 1$ spread her budget equally between the projects $\{1, \dots, m\}$, every agent l with $\beta_l > 1$ covers no project with a positive contribution, and every agent l with $\beta_l = 1$ contributes whatever we choose. Now, order the agents with $\beta_l < 1$ in any order such that the ones with the largest budget come last.

For each agent with $\beta_l < 1$ on order, besides the ones with the largest budget, we deal with each one as follows. Let the agent change her strategy, so as to best respond to the others' strategies. Since $\theta \leq 0.5$, it is possible to do this without contributing more than $\frac{b_n}{m}$ to a single project, because a contribution of at least $0.5 \frac{b_n}{m}$ to a project can be used to cover another project. This condition ensures that the strategy of the agents we have dealt with remains a best response to the rest, since a project that receives more contribution does not cease being covered, if it was covered before, and a project that loses contribution is not easier to cover than before, since there is more contribution there from the agents with largest budget, i.e. b_n .

By induction on the algorithm, all the agents we have dealt with best respond. The agents with the largest budget among those with $\beta_l < 1$ are always best responding, since they always cover all the projects. The agents with $\beta_l \geq 1$ best respond as well. In conclusion, this algorithm arrives at a situation where all the agents best respond, which means that it arrives at an NE. Since we have not altered the strategies of the agents with $\beta_l < 1$ with the largest budget, they contribute equally to all the projects $\{1, \dots, m\}$. \square

We now prove the case of close budgets.

Proposition 24. *Under the assumptions of the theorem, assume also that for any two agents $k < l$ (implying $b_k \leq b_l$) with $\beta_k < 1, \beta_l < 1$ we have $b_k \geq \theta b_l$. Then, there always exists an NE, where all the agents with $\beta_l < 1$ spread their budgets equally among the projects $\{1, \dots, m\}$, each agent with $\beta_l > 1$ does whatever we choose, as long as she covers no project with a positive contribution, and the agents with $\beta = 1$ do whatever we choose.*

Proof. Since the budgets of any two agents with $\beta < 1$ are in the factor of θ within each other, by spreading her budget equally between $\{1, \dots, m\}$, each agent with $\beta < 1$ covers all the projects $\{1, \dots, m\}$. Since every agent with $\beta_i > 1$ covers no project with a positive contribution, Theorem 29 implies that the profile is an NE. \square

This is the case of one budget being much larger than all the rest.

Proposition 25. *If $\beta_n \leq 1$, then, under the assumptions of the theorem, if $\theta b_n > m b_{n-1}$, then there exists an NE.*

Proof. Consider the profile when all the agents besides n contribute arbitrarily, while agent n spreads b_n equally between all the projects.

This is an NE, for the following reasons. Agent n is obviously getting her best possible utility. All the other agents do not want to deviate, because they would not be able to contribute enough to reach the threshold anyway. \square

The following proposition proves the existence of an NE for two agents.

Proposition 26. *Under the assumptions of the theorem, and assuming also that $n = 2$, there always exists an NE. If both agents have $\beta_i \leq 1$, we can ensure an NE such that the agent with the largest budget covers all the projects, while the other agent covers all projects but one.*

Proof. If at least one of the agents has $\beta_i > 1$, let her contribute nothing, and the other agent may divide her budget optimally for herself, resulting in an NE.

Assume now that both agents have $\beta_i \leq 1$. If $b_1 \geq \theta b_2$, then our proposition immediately stems from Proposition 24. Otherwise, let agent 1 contribute whole $\frac{b_1}{m-1}$ to projects $1, \dots, m-1$, and let agent 2 contribute $\frac{1}{\theta} \frac{b_1}{m-1}$ to each of these projects and $b_2 - \frac{1}{\theta} b_1$ to project m . Now, agent 2 covers everything, and thus obtains the maximum possible utility. Agent 1 obtains the maximum possible utility as well, since she cannot cover all projects, and the all besides one projects that she covers have the highest possible limit of the actions. Therefore, this is an NE. \square

The next question is the efficiency of the equilibria. Since the characterization of an NE is not simple, we can only prove the following partial results. First, we analyze the case of close budgets.

Proposition 27. *Under the assumptions of the theorem, and assuming also that $b_1 \geq \theta b_n$ and $\forall l \in N : \beta_l \leq 1$, we have PoS = 1.*

Proof. Consider the NE from Proposition 24, where all the agents equally spread their budget among the projects $\{1, \dots, m\}$. Since $b_1 \geq \theta b_n$, all the agents cover all these projects, and thus, the optimum social welfare of

$$(n-1) \left(n - \sum_{i \in N} \beta_i \right) \frac{\sum_{i \in N} \left(\frac{1}{r_i + r'_i} \cdot b_i \right)}{\sum_{i \in N} \left(\frac{1}{r_i + r'_i} \right)}$$

is achieved. Therefore, we have PoS = 1. \square

Now, we consider the case of one budgets being much larger than all the rest.

Proposition 28. *Under the assumptions of the theorem, and assuming also that $\beta_n \leq 1$ and $\theta b_n > m b_{n-1}$, we have PoS = 1 and PoA = $\frac{\frac{1}{r_i + r'_i} b_n}{\sum_{i \in N} \frac{1}{r_i + r'_i} b_i}$.*

Proof. Consider the proof of Proposition 25. The maxim social welfare is achieved when every agent spreads her budget equally, yielding the price of stability of 1. The minimum social welfare is achieved when agent n spread her budget equally and the rest contribute

nothing, yielding the price of anarchy of PoA = $\frac{\frac{1}{r_i + r'_i} b_n}{\sum_{i \in N} \frac{1}{r_i + r'_i} b_i}$. \square

At finally, we analyze case of two agents.

Proposition 29. Assume the theorem, and that $n = 2$.

If for every $i = 1, 2$ we have $\beta_i > 1$, then $\text{PoA} = \text{PoS} = 1$.

If $\beta_l \leq 1$, but $\beta_p > 1$, then

$$\text{PoA} = \text{PoS} = \begin{cases} \frac{\frac{1}{r_l+r'_l} \cdot (b_l)}{\sum_{i=1}^2 \frac{1}{r_i+r'_i} \cdot b_i} & \text{if } \beta_1 + \beta_2 \leq 2, \\ -\infty & \text{otherwise.} \end{cases}$$

If both $\beta_i \leq 1$, then $\text{PoS} = 1$, and

$$\text{PoA} = \begin{cases} \frac{(1-\beta_2) \left(\frac{1}{r_2+r'_2} b_2 \right)}{(1-\beta_1) \left(\frac{1}{r_1+r'_1} b_1 + \frac{1}{r_2+r'_2} \frac{b_1}{\theta} \right) + (1-\beta_2) \left(\sum_{i=1}^2 \frac{1}{r_i+r'_i} b_i \right)} & \text{if } b_1 < \frac{\theta b_2}{m}, \\ \frac{(1-\beta_1) \left(\frac{1}{r_1+r'_1} b_1 + \frac{1}{r_2+r'_2} \theta b_1 \right) + (1-\beta_2) \left(\sum_{i=1}^2 \frac{1}{r_i+r'_i} b_i \right)}{(1-\beta_1) \left(\frac{1}{r_1+r'_1} b_1 + \frac{1}{r_2+r'_2} \frac{b_1}{\theta} \right) + (1-\beta_2) \left(\sum_{i=1}^2 \frac{1}{r_i+r'_i} b_i \right)} & \text{if } \frac{\theta b_2}{m} \leq b_1 < \theta b_2, \\ 1 & \text{otherwise.} \end{cases}$$

Proof. First, assume that for every i , $\beta_i > 1$. Then, the only NE has both agents contributing nothing, and this is optimal, implying $\text{PoA} = \text{PoS} = 1$.

Next, assume that $\beta_l \leq 1$, but $\beta_p > 1$. In any NE, agent p contributes nothing, while agent l divides her budget in a way that is optimal to her. This achieves the social welfare of

$$(n-1)(2-\beta_1-\beta_2) \left(\frac{\frac{1}{r_l+r'_l} \cdot (b_l)}{\sum_{i \in N} \left(\frac{1}{r_i+r'_i} \right)} \right).$$

The optimal social welfare depends on whether $\beta_1 + \beta_2 \leq 2$. If yes, it is

$$(n-1)(2-\beta_1-\beta_2) \left(\frac{\sum_{i=1}^2 \frac{1}{r_i+r'_i} \cdot b_i}{\sum_{i \in N} \left(\frac{1}{r_i+r'_i} \right)} \right),$$

and if not, then it is zero. Dividing the social welfare in an NE by the optimal one yields the result.

We finally prove the case that both $\beta_i \leq 1$. Let us handle the price of stability first. Consider the equilibria in the proof of Proposition 26. If $b_1 \geq \theta b_2$, then we achieve the maximum possible social welfare. If $b_1 < \theta b_2$, then we also achieve the maximum possible social welfare, since 2 covers everything, and 1 covers the maximum it ever can. Thus, these NE there are optimal, and thus, $\text{PoS} = 1$.

Consider now the price of anarchy. If $b_1 < \frac{\theta b_2}{m}$, then in any NE, agent 2 covers all the projects. This minimum bound on social welfare is achieved in the NE where agent 2 contributes $\frac{b_2}{m}$ to every project, and agent 1 contributes nothing, since he cannot achieve the threshold anywhere. Thus, the minimum social welfare at an NE

here is $(n-1)(1-\beta_2) \frac{\left(\frac{1}{r_2+r'_2} b_2 \right)}{\sum_{i=1,2} \frac{1}{r_i+r'_i}}$. If $\frac{\theta b_2}{m} \leq b_1 < \theta b_2$, then in any NE, agent 2 covers

all the projects, and agent 1 covers at least one project, and the contribution from 2 at the projects that 1 covers is at least θb_1 . This minimum bound on social welfare is achieved in the NE where agent 1 contributes all b_1 to project 2 and agent 2 contributes there θb_1 , while contributing the rest to project 1. Thus, the minimum social

welfare at an NE here is $(n-1)(1-\beta_1) \frac{\left(\frac{1}{r_1+r'_1} b_1 + \frac{1}{r_2+r'_2} \theta b_1\right)}{\frac{1}{\sum_{i=1,2} \frac{1}{r_i+r'_i}}} + (n-1)(1-\beta_2) \frac{\left(\sum_{i=1,2} \frac{1}{r_i+r'_i} b_i\right)}{\frac{1}{\sum_{i=1,2} \frac{1}{r_i+r'_i}}}$.

To compute the price of anarchy, note that the optimum social welfare for $b_1 < \theta b_2$ is $(n-1)(1-\beta_1) \frac{\left(\frac{1}{r_1+r'_1} b_1 + \frac{1}{r_2+r'_2} \frac{1}{\theta} b_1\right)}{\frac{1}{\sum_{i=1,2} \frac{1}{r_i+r'_i}}} + (n-1)(1-\beta_2) \frac{\left(\sum_{i=1,2} \frac{1}{r_i+r'_i} b_i\right)}{\frac{1}{\sum_{i=1,2} \frac{1}{r_i+r'_i}}}$. Dividing the minimum values of social welfare in an NE by this optimum, we obtain the price of anarchy.

If $b_1 \geq \theta b_2$, then both agents cover every project in any NE, so PoA = 1. \square

The different prices of anarchy and stability shown in Propositions 28 and 29 indicate that regulation may improve the social welfare.

In the next section, we analyze the case where not only obtaining utility, but even participating requires contributing at least the threshold.

6. EXCLUSIVE THRESHOLDED RECIPROCATION EFFORT GAME

We now analyze the case where only the agents who contribute at least the threshold may interact. This models the situations when a minimum effort is necessary, such as becoming a member of a file sharing community or starting a firm. In this section, we assume w.l.o.g. that $b_n \geq \dots \geq b_1$ and all $\beta_i \leq 1$. We first observe that existence of an equilibrium is easy, since no-one contributing constitutes an NE. Then, we show that less trivial equilibria exist as well. Finally, the harder question of equilibrium efficiency is answered for two agents.

We first notice a trivial equilibrium.

Observation 11. *The profile where all agents contribute nothing is an NE.*

Proof. In this profile, any agent who deviates by contributing a positive amount to a project will be the only one to interact there, so her utility will still be zero. \square

We call an NE where at any project, at most one agent interacts (reaches the threshold) and positively contributes there, a Zero NE. There may be multiple Zero NE. We have just shown that a Zero NE always exists. A natural question is whether there exist non-Zero NE as well. We answer affirmatively.

Theorem 30. *Assume that all agents have $\beta_i \leq 1$. Assume that for any agent i , $r'_i > 0$ and in addition, for any pair of agents i, j we have $r_i + r'_i + r_j + r'_j < 2$. There exists a non-Zero NE.*

Proof. Consider the profile where all agents $1, \dots, n-1$ contribute their whole respective budgets to project 1, and agent n contributes $\min\left\{b_n, \frac{b_{n-1}}{\theta}\right\}$ to project 1, and nothing to other projects.

This is an NE, for the following reasons. Any agent would be alone at any project other than 1. At project 1, the only agent who perhaps can increase her contribution is n , but she will stay alone, if she does. \square

The next question is the efficiency of the equilibria. Since we always have the Zero NE, and by contributing to the same project the same positive amounts we achieve a positive social welfare, we always have $\text{PoA} = 0$. Regarding the price of stability, we immediately know that it is positive, since there always exists a non-Zero NE. We now show that the price of stability for two agents is 1, meaning that there exists a socially optimal NE.

Proposition 30. *For $n = 2$ and under the assumptions of the theorem, $\text{PoS} = 1$.*

Proof. When we have only two players, we can assume w.l.o.g. that in a profile with maximum social welfare, a project that receives a positive contribution, receives it from both agents. Therefore, social welfare is maximized by maximizing the total contribution to the projects where interaction occurs.

Then, the following profile maximizes the social welfare. Agent 1 spreads her budget equally between all the projects. If $b_1 \geq \theta b_2$, then agent 2 divides her budget equally between all the projects, and otherwise, she contributes $\frac{1}{\theta} \frac{b_1}{m}$ to every project. Since this profile constitutes an NE, we conclude that $\text{PoS} = 1$. \square

Till now, we analyzed dividing the budget between reciprocal interactions. We now turn to analyze dividing the budget and then deciding on the interaction habits.

7. EXTENSIVE RECIPROCATION EFFORT GAME

We want to analyze the situation in more detail, explicitly modeling the process of effort dividing and deciding on a reciprocation habit. To this end, we consider *extensive reciprocation effort games*, where no threshold exists. At the root of the game, all the agents divide their budgets, defining the kindness at the various interactions, and subsequently choose the reciprocation coefficient r_i , separately for each interaction. This models the situation, when people divide effort between several interactions, and once they find themselves in an interaction, each player chooses her reciprocation habits. Since we allow choosing reciprocation coefficients r_i , we cannot guarantee that $r_i + r'_i < 1$ for at least one player, and to have convergence, we therefore need to have at least 3 participants in each interaction, which create an odd cycle, namely a triangle, in the interaction graph. Otherwise, we may have no convergence, and thus, a meaningless game. We, therefore, analyze the non-thresholded case, and assume that $n \geq 3$.

This is an extensive-form game, since the game has two stages, and we therefore consider its subgame-perfect equilibria (SPE). An SPE is a strategy profile, such that at each possible stage of playing, no player has an incentive to deviate.

Analyzing the existence of SPE requires the following definition, adapted from Chapter 2.4 of [40].

Definition 24. *Utility function u_i in a game $(N, (A_i), (u_i))$ is quasi-concave, if for every $a^* \in A \triangleq \prod_{i=1}^n A_i$, the set $\{a_i \in A_i : u_i(a_i, a_{-i}^*) \geq u_i(a^*)\}$ is convex.*

We prove the existence of an SPE under a condition on what the agents perceive as their utility at the beginning (root) of the game. The existence of an SPE without this condition is left open.

Theorem 31. *Assume that for any agent i , $r'_i > 0$, and in addition, $n \geq 3$. If all $1 - \beta_i$ s have the same sign, and the agents perceive their utility at the root of the game tree as quasi-concave (instead of u_i , they perceive it as $f_i(u_i)$, for a continuous f_i , and $f_i(u_i)$ is quasi-concave), then there always exists an SPE in the extensive reciprocation effort game.*

The proof first constructively finds a NE in each subgame, and then employs Proposition 20.3 from [40] to prove that also the root of the game tree possesses a Nash equilibrium.

Proof. We begin by defining the Nash equilibria in each subgame. In any subgame obtained after dividing efforts between the interactions, a player faces m independent choices, each choice being a choice of her reciprocation coefficient r_i for an interaction project. At any project, consider the profile that maximizes the social welfare, described in Proposition 17 on page 138 of Chapter 5. Since all $1 - \beta_i$ s have the same sign, Proposition 18 on page 140 states that any social welfare maximizing profile is also an NE. This provides an NE in each subgame.

To complete an SPE profile, it now remains to find an NE in the root of the game tree. We prove its existence using Proposition 20.3 from [40]. First, the set of actions of a player consists of all the divisions of her budget between the project, which is a nonempty, compact and convex. Since the equilibria we consider are the largest possible convex combinations of the kindness values, attainable with r_i s in the segments $[0, 1 - r'_i]$, the original utility u_i of an agent from a project depends continuously on her kindness in that project, and since f_i is continuous function, the perceived utility $f_i(u_i)$ is continuous as well. Since we assume that at the root of the game, agents perceive their utilities as quasi-concave, Proposition 20.3 from [40] implies the existence of an NE. This completes the statement that an SPE exists. \square

This section finalizes our analysis of dividing effort between reciprocal interactions.

8. CONCLUSIONS AND FURTHER RESEARCH

In order to predict investing effort in several reciprocal interactions, we define a game that models dividing efforts between several reciprocal projects. We allowed for either having no contribution threshold, or for a threshold required to obtain utility or even to participate in an interaction.

We first recapitulate the results of the chapter and emphasize the important conclusions. Then, we provide some ideas on the future extensions of this research.

We show that any effort dividing between reciprocal interactions results in convergent interactions, regardless how actions are curbed to fit the budgets. Then, we assume that no curbing is required, and analyze the existence and efficiency of Nash equilibria of the game. We prove that when no contribution threshold exists, there always exists an equilibrium, and if acting is easy to everyone (for all i , $\beta_i < 1$) or hard to everyone (for all i , $\beta_i > 1$), then every NE is socially optimal.

We also show that any dividing of all the budget when acting is easy to everyone is a Nash equilibrium. The result may seem surprising. Intuitively, this happens because everyone participates in each interaction, and the concrete dividing of the budget does not matter to the social welfare. However, life does not often provide such situations.

If a minimum contribution is required to enjoy reciprocation, but everyone may participate, then we prove that there exists an equilibrium, if the threshold is at most half, or if the budgets are in the threshold factor from each other, or if only two agents exist. For close budgets, a socially optimal NE exists. In the case of two agents, the gap between the optimal and the suboptimal NE can be large, so regulation may be instrumental to play the efficient equilibrium, and not the less efficient ones.

If a minimum contribution is necessary even to participate in interaction, we show that the situation where no-one contributes is an equilibrium. This models the case where people are very passive, and no-one can start interaction project on his own. In addition to this trivial equilibrium, we find an equilibrium where all the agents contribute to the same project. This describes the case when people interact with each other on the same topic. Such a situation is clearly not the only option, since people often have many friendships [41]. Indeed, for two agents, there exists an equilibrium which is socially optimal.

In several cases, we see that the choices of strategies by the agents who are indifferent significantly influence the social welfare. For instance, this happens in the case without threshold to agents for whom acting and receiving action are equally important. Making such agents do what benefits the society can increase the social welfare.

We model the process of first dividing the effort and once everyone has invested in the projects, setting the reciprocation habit, as an extensive game. We provide a sufficient condition for the existence of a subgame perfect equilibrium.

There are many more interesting directions to extend the research. First, we know that the curbed interactions converge, while their equilibria remain veiled. In the extensive game setting, the existence proof lays the basis to continue and analyze the efficiency of the subgame perfect equilibria. Second, looking at interactions in large groups where not everyone can act on everyone else would be a natural generalization of our work. We assumed that two agents who interact in multiple projects, interact in these projects independently. Modeling the dependency between these interactions is interesting. Since reality contains not only reciprocation projects, analyzing a mixed set of projects, only some of which are interaction projects, would model reality better.

As for the complexity of finding an NE, we have shown that for thresholded reciprocation effort games, finding a best response is NP-hard, and we conjecture that the seemingly similar problems of finding a (pure) NE or even determining whether it exists is NP-hard as well.

This work models and analyzes a ubiquitous class of interactions and lays the basis for further research, aimed to provide more advice to the agents and to the manager who wants to maximize the social welfare.

REFERENCES

- [1] N. H. Nie and D. S. Hillygus, *Where does internet time come from?: A reconnaissance*, IT & Society **1**, 1 (2002).

- [2] A. Falk and U. Fischbacher, *A theory of reciprocity*, [Games and Economic Behavior](#) **54**, 293 (2006).
- [3] E. Fehr and S. Gächter, *Fairness and retaliation: The economics of reciprocity*, [Journal of Economic Perspectives](#) **14**, 159 (2000).
- [4] W. Güth, R. Schmittberger, and B. Schwarze, *An experimental analysis of ultimatum bargaining*, [Journal of Economic Behavior & Organization](#) **3**, 367 (1982).
- [5] F. M. Harper, S. X. Li, Y. Chen, and J. A. Konstan, *Social comparisons to motivate contributions to an online community*, in *Proceedings of the 2nd international conference on Persuasive technology* (Springer-Verlag, Berlin, Heidelberg, 2007) pp. 148–159.
- [6] A. Forte and A. Bruckman, *Why do People Write for Wikipedia? Incentives to Contribute to Open-Content Publishing*, in *GROUP Workshop: Sustaining Community: The Role and Design of Incentive Mechanisms in Online Systems*. Sanibel Island, FL (2005) pp. 6–9.
- [7] R. Siegel, *All-pay contests*, [Econometrica](#) **77**, 71 (2009).
- [8] J. Kleinberg and S. Oren, *Mechanisms for (mis)allocating scientific credit*, in [Proceedings of the 43rd annual ACM symposium on Theory of computing](#) (ACM, New York, NY, USA, 2011) pp. 529–538.
- [9] G. Polevoy, S. Trajanovski, and M. M. de Weerd, *Nash Equilibria in Shared Effort Games*, in *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems, AAMAS '14* (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2014) pp. 861–868.
- [10] Y. Bachrach, V. Syrgkanis, and M. Vojnović, *Efficiency and the Redistribution of Welfare*, Tech. Rep. (Microsoft Research, 2011).
- [11] E. Koutsoupias and C. Papadimitriou, *Worst-case equilibria*, in *16th Annual Symposium on Theoretical Aspects of Computer Science* (Trier, Germany, 1999) pp. 404–413.
- [12] C. Papadimitriou, *Algorithms, Games, and the Internet*, in [Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing](#), STOC '01 (ACM, New York, NY, USA, 2001) pp. 749–753.
- [13] A. S. Schulz and N. S. Moses, *On the performance of user equilibria in traffic networks*, in [Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms](#), SODA '03 (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2003) pp. 86–87.
- [14] E. Anshelevich, A. DasGupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden, *The price of stability for network design with fair cost allocation*, in [Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on](#) (2004) pp. 295–304.

- [15] M. R. Baye, D. Kovenock, and C. G. de Vries, *The all-pay auction with complete information*, *Economic Theory* **8**, 291 (1996).
- [16] B. Roberson, *The Colonel Blotto game*, *Economic Theory* **29**, 1 (2006).
- [17] E. Anshelevich and M. Hoefer, *Contribution games in networks*, *Algorithmica* **63**, 51 (2011).
- [18] R. Axelrod, *The emergence of cooperation among egoists*, *American Political Science Review* **75**, 306 (1981).
- [19] R. Axelrod, *The evolution of cooperation*, Basic books (Basic Books, 1984).
- [20] C. Bicchieri, *The grammar of society: The nature and dynamics of social norms* (Cambridge University Press, 2006).
- [21] S. Van Segbroeck, J. M. Pacheco, T. Lenaerts, and F. C. Santos, *Emergence of fairness in repeated group interactions*, *Phys. Rev. Lett.* **108**, 158104 (2012).
- [22] R. L. Trivers, *The evolution of reciprocal altruism*, *The Quarterly Review of Biology* **46**, 35 (1971).
- [23] E. Fehr, U. Fischbacher, and S. Gächter, *Strong reciprocity, human cooperation, and the enforcement of social norms*, *Human Nature* **13**, 1 (2002).
- [24] J. C. Cox, D. Friedman, and S. Gjerstad, *A tractable model of reciprocity and fairness*, *Games and Economic Behavior* **59**, 17 (2007).
- [25] M. Dufwenberg and G. Kirchsteiger, *A theory of sequential reciprocity*, *Games and Economic Behavior* **47**, 268 (2004).
- [26] M. Rabin, *Incorporating fairness into game theory and economics*, *The American Economic Review* **83**, pp. 1281 (1993).
- [27] G. Polevoy, M. de Weerd, and C. Jonker, *The convergence of reciprocation*, in *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '16 (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2016) pp. 1431–1432.
- [28] W. J. Dixon, *Reciprocity in United States-Soviet relations: Multiple symmetry or issue linkage?* *American Journal of Political Science* **30**, pp. 421 (1986).
- [29] M. D. Ward, *Modeling the USA-USSR arms race*, *Transactions of The Society for Modeling and Simulation International* **43**, 196 (1984).
- [30] J. Gottman, C. Swanson, and J. Murray, *The mathematics of marital conflict: Dynamic mathematical nonlinear modeling of newlywed marital interaction*, *Journal of Family Psychology* **13**, 3 (1999).
- [31] P.-A. Chen, M. David, and D. Kempe, *Better vaccination strategies for better people*, in *Proceedings of the 11th ACM Conference on Electronic Commerce*, EC '10 (ACM, New York, NY, USA, 2010) pp. 179–188.

- [32] P.-A. Chen, B. de Keijzer, D. Kempe, and G. Schäfer, *The robust price of anarchy of altruistic games*, in *Internet and Network Economics*, Lecture Notes in Computer Science, Vol. 7090, edited by N. Chen, E. Elkind, and E. Koutsoupias (Springer Berlin Heidelberg, 2011) pp. 383–390.
- [33] M. Hoefer and A. Skopalik, *Altruism in atomic congestion games*, *ACM Trans. Econ. Comput.* **1**, 21:1 (2013).
- [34] M. Rahn and G. Schäfer, *Bounding the inefficiency of altruism through social contribution games*, in *Web and Internet Economics*, Lecture Notes in Computer Science, Vol. 8289, edited by Y. Chen and N. Immorlica (Springer Berlin Heidelberg, 2013) pp. 391–404.
- [35] J. O. Ledyard, *Public Goods: A Survey of Experimental Research*, Public Economics 9405003 (EconWPA, 1994).
- [36] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, *Algorithmic Game Theory* (Cambridge University Press, 2007).
- [37] N. H. Nie and D. S. Hillygus, *The impact of internet use on sociability: Time-diary findings*, *It & Society* **1**, 1 (2002).
- [38] E. Seneta, *Non-negative Matrices and Markov Chains*, Springer Series in Statistics (Springer, 2006).
- [39] V. Vazirani, *Approximation Algorithms* (Springer, 2001).
- [40] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*, MIT Press Books, Vol. 1 (The MIT Press, 1994).
- [41] S. T. Ennett and K. E. Bauman, *Adolescent social networks: Friendship cliques, social isolates, and drug use risk*, *Improving prevention effectiveness*, 47 (2000).

7

CONCLUSION

*It is not the strongest or the most intelligent who will survive
but those who can best manage change.*

Leon C. Megginson, 1963

We first describe the general topic of investing in projects, our answers to the research questions presented in Section 1.3 of Chapter 1, and our main conclusions from the conducted research. The conclusions that are especially relevant to the SHINE project are explained as well. We proceed to suggest extensions to our work and further directions and approaches to research in the area of participation in projects.

1. MAIN RESULTS

This thesis studies the ubiquitous domain of investing in projects. In particular, we consider *value-creating* projects, where a value is created and shared among the contributors, and *interaction* projects, where the participants gain or lose by reciprocating. The value-creating projects model, for example, writing for Wikipedia [1] and co-authoring articles [2]. The reciprocation projects model interactions such as arguing people [3] and superpowers involved in an arms race [4–6].

Our main goal is to analyze investment in projects by rational self interested agents. This allows to both understand existing phenomena, such as the fact that following the kindest person often brings prosperity to the group, and to advise on how to invest in the projects in order to maximize the individual utility and the social welfare. The social welfare is the sum of all the individual utilities. We approach this task by looking at the stable states of the investment, where no investor has an incentive to change her current investments. These states are to be reasonably expected to occur, motivating our interest in them. These states, called Nash equilibria [7], are the most famous stable states in game theory. In order to advise how to maximize the social welfare, we look for the stable states that have high social welfare, i.e. that are socially efficient.

Therefore, the main research question, presented in Section 1.3, is what the Nash equilibria (NE) and their efficiencies are. This research question decomposes to further questions about the existence and efficiency of NE in games with value-creating projects and in games with reciprocation projects. In addition, investment of effort in reciprocation projects initiates several questions about reciprocation; namely, we ask how it uncurls and what are the equilibria and their efficiencies for the game of agents who set their habits or reciprocation before the interaction begins. The general approach to answering these questions is modeling the situation as a game and looking for NE. Studying the efficiency translates to analyzing the prices of anarchy (PoA) [8, 9] and stability (PoS) [10, 11]. Recall that the price of anarchy is the ratio of the socially least profitable NE to the maximum possible social welfare, while the price of stability is the ratio of the socially best equilibrium to the optimal social welfare. We now describe how we answer each of the research questions from Section 1.3 in detail. Together, these answers describe when agents act optimally for the society whenever they are in an equilibrium, when their actions may need a regulation to be socially efficient, and which regulation they need then.

The first research question (question 1 on page 8 from Section 1.3) is: “What are the NE of shared effort games with equal sharing of a linear project’s value to everyone who contributes above a threshold? How efficient are these equilibria?” Here, the agents contribute by dividing their budgets between the projects. Each project obtains a value defined by the appropriate project function coefficient times the total contribution it receives. Aiming to answer this question, Chapter 2 characterizes the existence of NE for two agents. We show that any NE attains more than a half of the optimal social welfare, so regulation is not desperately needed here. For the case of more than two agents, we employ a fictitious play similar to the one originated by Brown [12], searching for Nash equilibria. Only the ratio of the largest (linear) project function coefficients and the ratio of the two highest budgets determine the existence and efficiency of NE. Thus, only these ratios need to be modified if someone wants to influence the existence and efficiency

of equilibria. The efficiency depends on the ratio of the project coefficients piece-wise linearly, the slope increasing with the threshold. For more than two agents, the efficiency is sometimes less than half, so regulation would be useful. We hypothesize that an NE exists if and only if at least one of several sets of certain conditions on the ratio of the budgets holds, and for every such condition, some conditions of being smaller or equal to a function of the threshold (and not of the budgets) on the ratio of project function coefficients hold together.

The second question (question 2 from Section 1.3) refers to requirements that projects have to face, in order to obtain their value. It is: “What changes in the answer to the first question, if a project obtains its value only if it survives a competition between the projects?” Aiming to answer this, Chapter 3 extends Chapter 2 by modeling quota and success threshold on the projects. A *quota* is the number of the projects with the highest values that actually obtain their value, and a *success threshold* is the threshold such that only the projects with at least this value obtain their values. We provide sufficient conditions for having an NE and analyze their efficiency. We show that setting an appropriate success threshold is a more powerful tool to guarantee that an optimal profile can be an equilibrium, than setting a quota. This is an important guide to conference organizers, though since the optimality promise on the NE refers only to the price of stability, while the price of anarchy is small, additional regulation is required to avoid inefficient equilibria. The found equilibria partially answer question 2 from Section 1.3.

At this point, we start treating reciprocation projects, since this is the second important class of projects we consider. The first two research questions refer to a single reciprocal interaction. This is relevant both as a preparation to analyzing dividing time between several reciprocal interactions and as analyzing reciprocation by itself. First, we analyze a single process and its limits, and then we consider strategic choice of reciprocation habits. Finally, we consider dividing effort between several reciprocation projects, similarly to what we do with value receiving projects.

We model agents, each of which is endowed with internal inclination, called *kindness* and *reciprocation coefficients*, defining the extent of reaction to the others' actions. An agent can be either *fixed*, where she reacts to others' actions and takes into account own kindness, or *floating*, where she also reacts to the others' actions but takes into account last own action instead of own kindness. The *fixed* attitude models being loyal to one's own inclinations, while the *floating* agents get carried away with the process. Each agent influences a subset of other agents, defined by her neighborhood in the interaction graph.

Question 3 is what the actions will become in a reciprocal interaction in the long run. We answer this in Chapter 4, where we prove exponential convergence of reciprocation and find the limits for pairwise interaction and for the case where at most one agent has the *fixed* reciprocation attitude. In the latter case, the limit is common to all the actions of all the agents. The results show that persistence (either in the sense of being *fixed* or having a low reciprocation coefficient) causes the interaction to resemble one's kindness. The same effect of setting the tone of the interaction occurs from being able to influence many agents. Interestingly, converging to a limit resembles converging to own style of behavior, known in life [13], and a common limit represents an organizational culture [14]. These results answer question 3.

Next, we study question 4, which inquires which reciprocation habits are the most efficient in the long run. A habit is modeled by the reciprocation attitude and/or coefficient of the agent. Thus, aiming to analyze reciprocation from a strategic perspective, Chapter 5 defines the utilities of the interacting agents and suggests the exact habits that maximize one's utility. When the agents choose either the reciprocation attitude or coefficient, then both for maximizing own utility and the social welfare, the less kind agents should be more flexible (which means employ the *floating* attitude or have a large reciprocation coefficient), if acting is cheaper than receiving actions, and that otherwise, the reverse holds. This fits the intuition that the kinder/less kind agents should pull the less kind/kinder ones to act more/less, thereby increasing the utility and the social welfare, when acting is easy/hard. This explains why following the kindest brings prosperity [15]. If acting is much cheaper than being acted upon, then the price of anarchy is 1, so agents act optimally in any NE. If the agents may choose both attitude and coefficient, then an optimal equilibrium indeed exists, but also a suboptimal one does, requiring regulation. The described results answer research question 4.

Finally, incorporating reciprocation in shared effort games, question 5 asks which Nash equilibria exist in the games of dividing effort between reciprocal interactions, and how efficient they are. Chapter 6 studies this. Among other results, in the case where no threshold of investment in a reciprocation project exists, we show that an NE always exists, and if acting is cheaper than receiving an action, then any NE is optimal. This means that provided no requirements exist on minimum effort in interaction, regulation is not required for optimality. Furthermore, modeling the process of first dividing own efforts between the interactions, and subsequently choosing the reciprocation coefficients at each interaction, we prove the existence of a subgame perfect equilibrium [16] in this extensive game, motivating future research of this two-step process. The results of Chapter 6 answer question 5 about existence and efficiency of NE for the analyzed important cases.

As we have discussed, we answer all the research questions by discovering the Nash equilibria and their efficiencies; sometimes, the answers only pertain to subcases of the model. This partially answers the overarching research question regarding the Nash equilibria and their efficiencies in project participation, advising on the need for regulation as follows. When the best efficiency of an equilibrium, which is the price of stability (PoS), is low, then the only way to improve the social welfare is to force the agents to go against their interest or to fundamentally change the situation. When the price of stability is high but the price of anarchy (PoA) is low, then we only need to regulate so that the agents play the efficient NE. On the other hand, when even the PoA is high, no regulation is needed, besides, perhaps, suggesting (any) NE to the participants.

We find many stable situations and their efficiencies in the studied scenarios and provide a solid piece of advice on what an agent or the manager needs to do. We still leave some space for extensions, some of which appear in the next section. But first we show some conclusions about the SHINE project, which was an important inspiration to our research questions.

1.1. THE CONNECTION TO THE SHINE PROJECT

The SHINE¹ project for information requiring and gathering by self-organizing agents is described in Section 1.2. It can be roughly modeled as a value-creating project, as well as an interaction, the actions being information providing. Let us see which general conclusions are especially interesting with relation to SHINE.

Assume first every SHINE-like crowdsourcing project is modeled as a value-creating project. Regarding putting effort in and sharing rewards from crowdsourcing projects, we provide conditions for the existence of socially optimal equilibria. For instance, under certain conditions, Section 3 of Chapter 2 shows that people would benefit from equally participating in the most promising projects, which gives yet another motivation to make our project be the best. A second place can be no good. If such projects need to be in the quota of a certain number of the best projects in order to remain active, then we saw in Section 3 of Chapter 3 that there are many stable situations, some of which may be very efficient, like everyone participating in a most profitable project, but some may be very inefficient, such as everyone participating in a least profitable project. Therefore, good media coverage that will explain which projects are most profitable is crucial for making the society play the efficient equilibria.

A crowdsourcing project like SHINE constitutes an interaction between its participants. Providing correct information is a positive act by the provider on the requester, while providing wrong information is a negative act. To model interaction in a crowdsourcing project, we model and predict the uncurling of a reciprocation process. In Section 5 of Chapter 4, we prove that the eventual influence of one's kindness is directly proportional to the number of people she influences and inversely proportional to her reactivity. Therefore, to design a system that maximizes mutual help, the kinder persons should be able to interact with as many people as possible. Once the designers of SHINE have somehow found the kinder persons, these persons can be given more peers to ask and to provide information to.

Regarding strategic choice of reciprocation habits, when everyone wants to have more help around, then Section 5, Section 6, and Section 10 of Chapter 5 show that the kinder people should keep helping, while “pulling” the less kind ones to help others more. Public campaign for this goal would facilitate interaction. If helping is easy and every person chooses only one habit of hers, then Section 5 and Section 10 of Chapter 5 demonstrate that any Nash equilibrium is socially optimal, so regulation is not needed here. However, when one can choose several of her habits at once, then Section 6 of Chapter 5 shows that there can be suboptimal Nash equilibria as well, so the system suggesting to act according to the optimal equilibrium may be expedient.

Consider people who split efforts between several interactions where we do model these projects as interactions, and helping is easy for most people, but some people find helping as hard as not being helped. Then, convincing these people to still help others will improve the social welfare, as Chapter 6 shows.

¹See <http://www.participatorysystems.nl/2013/02/shine/>, <http://mmi.tudelft.nl/~birna/projects/projects.html>, and <http://shine.tudelft.nl/home/>.

2. FUTURE RESEARCH

We lay the foundation of studying investment in several projects. We have concentrated on projects with a minimum threshold required to receive utility from them. Further steps of modeling and predicting participation in projects on top of our work are described next.

2.1. EXTENSIONS

Aiming to the goal of analyzing investing in projects, our work models two crucial kinds of projects, motivating further analysis of these kinds of projects and modeling and analysis of other kinds as well. First, it would be nice to complete the characterizations of the existence and the efficiency of equilibria that we have done. By characterizing the efficiency of equilibria we mean, for instance, finding the price of anarchy and stability, as they define the least and the most efficient NE. An interesting direction for future work is to model personal interactions between the agents, based on what they contribute to the various projects. This would allow the investments in distinct projects by the same agents to influence each other. We now go over the individual chapters and describe what can be extended in each chapter.

In Chapter 2, we form the basis of investing in projects by characterizing the existence of NE for 2 agents, for a thresholded equal sharing of projects with values linear in the total contribution. For more agents, we provide sufficient existence conditions. Extending the model to other, non-thresholded-equal-value-sharing mechanisms, and to non-linear project functions would allow modeling a wider range of real projects. For example, the sharing of the value of a project can be proportional to the individual investments [17]. In addition, the value of a project can depend on the total investment in an exponential manner, like the hardness of a code depends on the number of digits in the key.

Chapter 3 expands on the model from Chapter 2, and this complication allows for fewer theoretical results: we find sufficient conditions for NE, instead of characterizations. Analyzing these models further is an interesting direction to go to. Moreover, still more refinements can facilitate modeling the following scenarios even further. For example, projects like papers and books often have an upper bound on the maximal number of participants. For instance, writing a paper with ten co-authors is quite complicated, though not impossible. Furthermore, agents like people can only contribute to only a reasonably bounded number of projects, as we all have natural limitations. Modeling these upper bounds can improve the predictive power of our results; for instance, a bound on the number of the contributors to a project would rule out the equilibria where everyone invests in the same project.

Chapters 4 and 5 thoroughly analyze our model of reciprocation. This model is very basic and general, therefore allowing for many extensions, as we now present.

1. Introducing new parameters to the model allows modeling certain situations more precisely:
 - (a) We model reciprocation. One could also model more complicated motives to act, such as liking one's photo on a social medium because of both recip-

- rocating and actually finding the photo likeable. An action here is a combination of reacting and liking.
- (b) It would be also interesting to model agents with different kindness values towards different agents, and even different reciprocation coefficients, to represent the personal preferences.
 - (c) We have considered actions to be absolute. Allowing actions to be interpreted differently by different agents may fit to situations where various agents interpret or experience actions differently.
2. Adding dynamics and probabilistic behavior can improve the modeling of situations where parameters change and are not certain:
- (a) Another interesting venue would be changing the reciprocation parameters (aka habits) multiple times in the process, to model the people or countries changing while reciprocating. This requires time-dependent parameters.
 - (b) Considering a dynamically changing set of reciprocating agents is also important, since in life, people join groups (like the colleagues at a workplace who are hired and dismissed) and leave them.
 - (c) Noisy interaction may benefit from a probabilistic modeling of the reactions. This models real reactions, which usually are not fully predictable.
3. We may also ask how much influence various agents have on the total interaction. We have indeed seen that when everyone is *floating* (taking into account last own action instead of own kindness), then the share of one's kindness in the common limit is proportional to the number of agents she acts on and reversely proportional to the sum of one's reciprocation coefficients, i.e. how reciprocal she is. We have also seen that when *fixed* (taking into account own kindness instead of last own action) agents exist, then the kindness of the *floating* agents does not matter. Nonetheless, the exact influence of the *fixed* agents on the limits of the actions is yet to be discovered.

Chapter 6 considers an interesting kind of public projects, namely reciprocation, representing meeting friends of playing group sports. In life, there are many kinds of public projects, with various rules for each project, so modeling them would be the next step.

A basic approach of our research is that we sometimes take the sociological and behavioral knowledge as an inspiration for modeling a part of reality, and afterwards, we compare our results with the existing behavioral knowledge. For instance, we used the models of arguing people as an inspiration of the model of reciprocation. Conducting an experimental study to check the models and their predictions in details would allow to delineate where our work is applicable. An experimental analysis of how people, countries and other agents behave can also suggest how the model should be refined to allow for analyzing certain applications.

In addition to directly extending our models, we next discuss further approaches of investing in projects for practical usage. These directions are reminiscent of what we have done, but model other situations.

2.2. OTHER DIRECTIONS

This work models agents who divide their budgets of efforts between projects and gain from these projects. We now consider two other avenues of modeling investment in projects. The first avenue is a complementary — almost dual — modeling where the roles change, and next, we introduce dynamics to the model. The dual models look at the process from another angle, and the dynamic extensions allow for a more realistic modeling of the changes in time.

Consider the following dual points of view on investing in projects:

1. Projects select the agents who may participate. This can be done by an approval voting, conducted among the current participants in the project.
2. This thesis assumes a given definition of utilities and analyzes the resulting strategies. Another option would be to consider how to divide a project's profits so that the agents do not want to contribute differently, similarly to the model of cooperative games. Cooperative games with distributing resources among several coalitions have been considered in [18–20], and one can augment that model with project-dependent value functions.

Additionally, one can consider the following models that progress over time:

1. Introducing the time dimension to the model, where agents plan their investments ahead, in a repeated game. The information about a project becomes known to an agent only once she has invested there. This models the experience agents gain with learning.
2. Considering the evolutionary dynamics of several projects types and several agent types. A project type defines the project function that defines the value of the project based on the investments it receives, and an agent type defines the agent's investment policy. Similarly to the real life, the best projects and agents survive and continue to the next round.

2.3. OTHER APPROACHES TO PARTICIPATION IN PROJECTS

We have looked at project participation from the game-theoretic point of view. Based on this strategic analysis, let us now glance at some other aspects of investing effort that are practically important.

A natural question is whether our research can facilitate investing in certain projects in practice. The answer should be affirmative for the projects that are properly captured by our models. However, we have to consider the human factor. The optimistic point is that people do incorporate time management advice they get [21]. However, an important obstacle still remains to be overcome, namely, the personal predisposition to time management influences the actual management [22]. Therefore, psychologists may help with assisting people to properly realize the scientific recommendations.

We modeled people who are free to divide their efforts, bounding sometimes only the set of the addressed projects. In reality, investment in projects may have to face legal requirements, and modeling and analyzing these is a necessary element of practical deployment of any ideas. Another practical consideration is the public opinion, which

may have effects on the emotional part of the project values: people enjoy contributing to a project about which they and most others feel good.

To conclude, this thesis predicts the development of scenarios of dividing resources between common projects and rigorously advises how to act there efficiently. Such advice can be used by people, by automatic decision support systems, or interacting computers. This modeling and analysis provides many insights and lays the basis for further modeling and analysis, which will allow providing even better advice.

REFERENCES

- [1] A. Forte and A. Bruckman, *Why do People Write for Wikipedia? Incentives to Contribute to Open-Content Publishing*, in *GROUP Workshop: Sustaining Community: The Role and Design of Incentive Mechanisms in Online Systems*. Sanibel Island, FL (2005) pp. 6–9.
- [2] J. Kleinberg and S. Oren, *Mechanisms for (mis)allocating scientific credit*, in *Proceedings of the 43rd annual ACM symposium on Theory of computing* (ACM, New York, NY, USA, 2011) pp. 529–538.
- [3] J. Gottman, C. Swanson, and J. Murray, *The mathematics of marital conflict: Dynamic mathematical nonlinear modeling of newlywed marital interaction*, *Journal of Family Psychology* **13**, 3 (1999).
- [4] R. Axelrod, *The emergence of cooperation among egoists*, *American Political Science Review* **75**, 306 (1981).
- [5] W. J. Dixon, *Reciprocity in United States-Soviet relations: Multiple symmetry or issue linkage?* *American Journal of Political Science* **30**, pp. 421 (1986).
- [6] M. D. Ward, *Modeling the USA-USSR arms race*, *Transactions of The Society for Modeling and Simulation International* **43**, 196 (1984).
- [7] J. Nash, *Non-Cooperative Games*, *The Annals of Mathematics Second Series*, **54**, 286 (1951).
- [8] E. Koutsoupias and C. Papadimitriou, *Worst-case equilibria*, in *16th Annual Symposium on Theoretical Aspects of Computer Science* (Trier, Germany, 1999) pp. 404–413.
- [9] C. Papadimitriou, *Algorithms, Games, and the Internet*, in *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing*, STOC '01 (ACM, New York, NY, USA, 2001) pp. 749–753.
- [10] A. S. Schulz and N. S. Moses, *On the performance of user equilibria in traffic networks*, in *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '03 (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2003) pp. 86–87.
- [11] E. Anshelevich, A. DasGupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden, *The price of stability for network design with fair cost allocation*, in *Foundations of*

- Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on* (2004) pp. 295–304.
- [12] G. W. Brown, *Iterative solution of games by fictitious play*, Activity analysis of production and allocation **13**, 374 (1951).
- [13] B. W. Roberts, K. E. Walton, and W. Viechtbauer, *Patterns of mean-level change in personality traits across the life course: a meta-analysis of longitudinal studies*. Psychological bulletin **132**, 1 (2006).
- [14] G. Hofstede, *Culture and organizations*, *International Studies of Management & Organization* **10**, 15 (1980).
- [15] R. Axelrod and W. Hamilton, *The evolution of cooperation*, *Science* **211**, 1390 (1981), <http://www.sciencemag.org/content/211/4489/1390.full.pdf>.
- [16] R. Selten, *Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragertragheit*, *Zeitschrift für die gesamte Staatswissenschaft* **121**, 301 (1965).
- [17] Y. Bachrach, V. Syrgkanis, and M. Vojnović, *Efficiency and the Redistribution of Welfare*, Tech. Rep. (Microsoft Research, 2011).
- [18] Y. Zick and E. Elkind, *Arbitrators in overlapping coalition formation games*, in *AA-MAS*, edited by L. Sonenberg, P. Stone, K. Tumer, and P. Yolum (IFAAMAS, 2011) pp. 55–62.
- [19] G. Chalkiadakis, E. Elkind, E. Markakis, M. Polukarov, and N. R. Jennings, *Cooperative games with overlapping coalitions*, *J. Artif. Int. Res.* **39**, 179 (2010).
- [20] Y. Zick, G. Chalkiadakis, and E. Elkind, *Overlapping coalition formation games: charting the tractability frontier*, in *AAMAS*, edited by W. van der Hoek, L. Padgham, V. Conitzer, and M. Winikoff (IFAAMAS, 2012) pp. 787–794.
- [21] J. H. Richards, *Time management—a review*, *Work & Stress* **1**, 73 (1987), <http://dx.doi.org/10.1080/02678378708258485>.
- [22] B. J. Claessens, W. van Eerde, C. G. Rutte, and R. A. Roe, *A review of the time management literature*, *Personnel Review* **36**, 255 (2007), <http://dx.doi.org/10.1108/00483480710726136>.

ACKNOWLEDGEMENTS

Now it is the dangerous part. You forget someone, and it is over.

First and foremost, I am deeply grateful to my supervisor Mathijs de Weerd for his great guidance in research, writing and presenting, and in the proper academic life. He was always patient and highly helpful, listening to my thoughts and answering my questions. Mathijs, you helped me both conduct the research and enjoy it and the rest of my life here! I am also indebted to my promotors: Catholijn Jonker and Cees Witteveen. Catholijn, your advice and ideas made a huge impact on most parts of the work! Cees, the discussions on the thesis, from the large parts to the little details, immensely improved the quality of this dissertation!

I was extremely lucky with the people surrounding the research and gradually becoming my friends. Michel and Peter (Bosman), for their advice and interesting discussions, Peter (Novak), for our fervent arguments and deep discussions and for your intransigence, Kiriakos (Simon), for his support and humor. Shruti, your friendship and our discussions were invaluablely supportive! The new PhD candidate, Longjian Piao, exceptionally combines a highly motivated professional, an interesting interlocutor and a great person. Thanks Leon, Adriaan and Boris, for your help when I was fresh here. Thanks João, Tiago, Bob, Jaume, for your great company! Thanks Matthijs, Erwin, Frits and Rens for your help and advice! This includes watering my plant when I went to conferences. Especial thanks go to Rens for proofreading the summary for Pure. Stojan, Thomas, and Qingzhi, you were great to cooperate with! I especially remark your advice, Stojan, thanks a lot! Diederik, Reyhan, and Tim, you are charming as people as well as scientists. Joris, it was fun to talk to you! Shemara, Janneke (Hermans), Paulo, Munire, Stephen and Erik, it is a pleasure to acknowledge your generous assistance.

I had much more superb advice that I can describe here. To mention a few, Johan Pouwelse, Alexandru Iosup, Georgi Gaydadjiev, Birna van Riemsdijk, Virginia Dignum, Martijn Warnier, and John-Jules Meyer gave me advice I am very happy about! Orr, thanks for your efforts helping me with a mathematical question.

I especially thank my parents and grandmother, whose limitless support made all my achievements possible. Sasha (Landau), your friendship and wise advise are always invaluable. Alexey and Nikita (Lenchenkov), being able to discuss life with you and have the ability to openly dissect the intricacies of being were the best joys a scientist can have. Andrey and Dima (Verkhovsky), I highly value our ancient friendship! Dima and Yulia (Kononchuk), your warm and crisp company was a great retreat from the problems life had. Especial thanks go to Natalia who helped us all meet and made the Slavic life in Delft really gezellig! Ksenia, your imperturbable criticism and frank and quiet voice made me aspire to achieve more! Nikita (de Keijzer), thanks for your sensitivity, sharp thought, and for proofreading the Dutch version of the summary (de samenvatting) and of the propositions. Roma, your friendship was an invaluable asset! Ivan, your sober and original voice was highly welcome. Emil and Irina, thanks for the serious and funny time

we shared! Misha (Belonosov), thanks for your calm support and exhilarating discussions! Sveta (Lavskaya), you are a great, open and positive friend! Maiia (Chvaleva), your crisp thinking and personal warmth are highly appreciated! Yuliya (Miloslavina), your warm and highly unique personality and your company were exceptionally important! Daniya, your support made me feel much better about life! Anja (Envaeva), thanks for proofreading the Dutch version and for interviewing me about my work.

I am glad to have met many good people, whose presence in my life is highly esteemed: Vova (Korzinov), Denis, Mark (Khait), Anastasiia Pravedna, Nelli, Anastasia Holovchenko, Jenya (Platonova), Arina, Liza (Mirgorodskaya), Reinaldo, Yulia Yadgarova, Elena (Rogova), Nick Gaiko, Inna and Sergey, Maria (Zamiralova), Katerina, Ilya and Maia, Ievgen and Kateryna, Misha and Maria, Janneke (van der Zwaan), Jonas, and many others.

CURRICULUM VITÆ

Gleb POLEVOY

Gleb Polevoy was born in Donetsk, Ukraine (a part of the Soviet Union then) on June 09, 1982. He moved, together with his family, to Israel in 1996. He was awarded excellency every semester of the high school. At high school, he won the 1-st place at the Israeli astrophysics championship. His interests included mathematics, computer science, and physics.

Having begun studying at the Technion, I.I.T., he obtained five Technion President's excellency awards and after three years he was awarded a B.A. in mathematics and computer science (Summa cum Laude). Afterwards, he was obliged to be a software engineer in the army for 6 years, during which he studied for an MSc in computer science at the Technion, I.I.T. He defended a thesis called "Bandwidth Allocation in Cellular Networks with Multiple Interferences" under the supervision of Prof. dr. R. Bar-Yehuda.

Afterwards, he was looking for a suitable area for a PhD, visiting various courses and studying several topics on his own. This was a very engaging and enriching search, during which he worked and published with Prof. Reuven Cohen on network algorithms and worked with Rann Smorodinsky and Moshe Tennenholtz on competitive signaling of sellers to buyers. The last paper was his first game-theoretic work, while the interest in game theory had begun earlier on.

He chose a PhD in game theory with applications in the SHINE project about information gathering on demand. This was conducted under the supervision of Mathijs de Weerd, and with the promotors Catholijn Jonker and Cees Witteveen from Delft University of Technology. This work involved modeling of dividing efforts and setting interaction habits as non-cooperative games. He is about to begin a post-doctorate position at the University of Amsterdam, which is about algorithms in network security and network analysis.

LIST OF PUBLICATIONS

11. G. Polevoy, M.M. de Weerd, and C.M. Jonker. The game of reciprocation habits. In *Proceedings of the 2016 European Conference on Artificial Intelligence*, volume Volume 285: ECAI 2016 of *Frontiers in Artificial Intelligence and Applications*, pages 417–425, 2016.
10. G. Polevoy, M. de Weerd, and C. Jonker. The convergence of reciprocation. In *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '16, pages 1431–1432, Richland, SC, 2016. International Foundation for Autonomous Agents and Multiagent Systems.
9. R. Cohen and G. Polevoy. Inter-datacenter scheduling of large data flows. *Cloud Computing, IEEE Transactions on*, PP(99):1–1, 2015.
8. R. Bar-Yehuda, G. Polevoy, and D. Rawitz. Bandwidth allocation in cellular networks with multiple interferences. *Discrete Applied Mathematics*, 194:23 – 36, 2015.
7. G. Polevoy, S. Trajanovski, and M. M. de Weerd. Nash equilibria in shared effort games. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14, pages 861–868, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.
6. T. C. King, Q. Liu, G. Polevoy, M. de Weerd, V. Dignum, M. B. van Riemsdijk, and M. Warnier. Request driven social sensing. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14, pages 1651–1652, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.
5. G. Polevoy and M. M. de Weerd. Improving human interaction in crowdsensing. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14, pages 1741–1742, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.
4. G. Polevoy, R. Smorodinsky, and M. Tennenholtz. Signaling competition and social welfare. *ACM Trans. Econ. Comput.*, 2(1):1:1–1:16, Mar. 2014.
3. R. Cohen, I. Nudelman, and G. Polevoy. On the admission of dependent flows in powerful sensor networks. *Networking, IEEE/ACM Transactions on*, 21(5):1461–1471, Oct 2013.
2. R. Cohen, I. Nudelman, and G. Polevoy. On the admission of dependent flows in powerful sensor networks. In *Infocom'2012, Orlando, Florida*, Mar. 2012.
1. R. Bar-Yehuda, G. Polevoy, and D. Rawitz. Bandwidth allocation in cellular networks with multiple interferences. In *DIALM-PODC*, pages 33–42, 2010.

REFERENCES

- [1] *Linux website*, <https://www.linux.com/> (2016).
- [2] A. Forte and A. Bruckman, *Why do People Write for Wikipedia? Incentives to Contribute to Open-Content Publishing*, in *GROUP Workshop: Sustaining Community: The Role and Design of Incentive Mechanisms in Online Systems*. Sanibel Island, FL (2005) pp. 6–9.
- [3] R. Ganti, F. Ye, and H. Lei, *Mobile crowdsensing: current state and future challenges*, *IEEE Communications Magazine* **49**, 32 (2011).
- [4] J. Kleinberg and S. Oren, *Mechanisms for (mis)allocating scientific credit*, in *Proceedings of the 43rd annual ACM symposium on Theory of computing* (ACM, New York, NY, USA, 2011) pp. 529–538.
- [5] R. Axelrod, *The emergence of cooperation among egoists*, *American Political Science Review* **75**, 306 (1981).
- [6] W. J. Dixon, *Reciprocity in United States-Soviet relations: Multiple symmetry or issue linkage?* *American Journal of Political Science* **30**, pp. 421 (1986).
- [7] M. D. Ward, *Modeling the USA-USSR arms race*, *Transactions of The Society for Modeling and Simulation International* **43**, 196 (1984).
- [8] J. Gottman, C. Swanson, and J. Murray, *The mathematics of marital conflict: Dynamic mathematical nonlinear modeling of newlywed marital interaction*, *Journal of Family Psychology* **13**, 3 (1999).
- [9] A. Falk and U. Fischbacher, *A theory of reciprocity*, *Games and Economic Behavior* **54**, 293 (2006).
- [10] E. Fehr and S. Gächter, *Fairness and retaliation: The economics of reciprocity*, *Journal of Economic Perspectives* **14**, 159 (2000).
- [11] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*, MIT Press Books, Vol. 1 (The MIT Press, 1994).
- [12] A. Rubinstein, *Modeling Bounded Rationality*, 1st ed., Vol. 1 (The MIT Press, 1997).
- [13] J. Nash, *Non-Cooperative Games*, *The Annals of Mathematics Second Series*, **54**, 286 (1951).
- [14] Y. Chen, F. M. Harper, J. Konstan, and S. X. Li, *Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens*, *American Economic Review* **100**, 1358 (2010).
- [15] M. A. Raghunandan and C. A. Subramanian, *Sustaining cooperation on networks: An analytical study based on evolutionary game theory*, in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS'12 (2012) pp. 913–920.

- [16] W. Güth, R. Schmittberger, and B. Schwarze, *An experimental analysis of ultimatum bargaining*, *Journal of Economic Behavior & Organization* **3**, 367 (1982).
- [17] J. Sobel, *Interdependent preferences and reciprocity*, *Journal of Economic Literature* **43**, 392 (2005).
- [18] A. Begel, R. DeLine, and T. Zimmermann, *Social media for software engineering*, in *Proceedings of the FSE/SDP Workshop on Future of Software Engineering Research*, FoSER '10 (ACM, New York, NY, USA, 2010) pp. 33–38.
- [19] M. Allahbakhsh, B. Benatallah, A. Ignjatovic, H. R. Motahari-Nezhad, E. Bertino, and S. Dustdar, *Quality control in crowdsourcing systems: Issues and directions*, *IEEE Internet Computing* **17**, 76 (2013).
- [20] Y. Wang and D. R. Fesenmaier, *Assessing motivation of contribution in online communities: An empirical investigation of an online travel community*, *Electronic Markets* **13**, 33 (2003), <http://www.tandfonline.com/doi/pdf/10.1080/1019678032000052934>.
- [21] M. Bagnoli and M. McKee, *Voluntary contribution games: Efficient private provision of public goods*, *Economic Inquiry* **29**, 351 (1991).
- [22] M. Bagnoli and B. Lipman, *Provision of public goods: Fully implementing the core through private contributions*, *Review of Economic Studies* **56**, 583 (1989).
- [23] N. Kaufmann, T. Schulze, and D. Veit, *More than fun and money. worker motivation in crowdsourcing - a study on mechanical turk*, in *AMCIS* (2011).
- [24] F. M. Harper, S. X. Li, Y. Chen, and J. A. Konstan, *Social comparisons to motivate contributions to an online community*, in *Proceedings of the 2nd international conference on Persuasive technology* (Springer-Verlag, Berlin, Heidelberg, 2007) pp. 148–159.
- [25] P. J. Ludford, D. Cosley, D. Frankowski, and L. Terveen, *Think different: increasing online community participation using uniqueness and group dissimilarity*, in *Proceedings of the SIGCHI conference on Human factors in computing systems* (ACM, New York, NY, USA, 2004) pp. 631–638.
- [26] R. Siegel, *All-pay contests*, *Econometrica* **77**, 71 (2009).
- [27] L. S. Shapley, *A value for n-person games. Contribution to the Theory of Games*, *Annals of Mathematics Studies* **2**, 28 (1953).
- [28] H. Moulin, *On additive methods to share joint costs**, *Japanese Economic Review* **46**, 303 (1995).
- [29] Y. Bachrach, V. Syrgkanis, and M. Vojnović, *Efficiency and the Redistribution of Welfare*, Tech. Rep. (Microsoft Research, 2011).
- [30] M. R. Baye, D. Kovenock, and C. G. de Vries, *The all-pay auction with complete information*, *Economic Theory* **8**, 291 (1996).

- [31] B. Roberson, *The Colonel Blotto game*, *Economic Theory* **29**, 1 (2006).
- [32] E. Anshelevich and M. Hoefer, *Contribution games in networks*, *Algorithmica* **63**, 51 (2011).
- [33] R. Smorodinsky and M. Tennenholtz, *Overcoming free riding in multi-party computations – the anonymous case*, *Games and Economic Behavior* **55**, 385 (2006).
- [34] P. A. Bernstein, D. DeWitt, A. Heuer, Z. Ives, C. S. Jensen, H. Meyer, M. T. Özsu, R. T. Snodgrass, K.-Y. Whang, and J. Widom, *Database publication practices*, in *Proceedings of the 31st International Conference on Very Large Data Bases*, VLDB '05 (VLDB Endowment, 2005) pp. 1241–1245.
- [35] J. R. Douceur, *Paper rating vs. paper ranking*, *SIGOPS Oper. Syst. Rev.* **43**, 117 (2009).
- [36] R. Axelrod, *The evolution of cooperation*, Basic books (Basic Books, 1984).
- [37] W. Güth, M. Yaari, and U. Witt, *Explaining Reciprocal Behavior in Simple Strategic Games: An Evolutionary Approach*, (University of Michigan Press, Ann Arbor, 1992).
- [38] R. Sethi and E. Somanathan, *Preference evolution and reciprocity*, *Journal of Economic Theory* **97**, 273 (2001).
- [39] R. Axelrod and W. Hamilton, *The evolution of cooperation*, *Science* **211**, 1390 (1981), <http://www.sciencemag.org/content/211/4489/1390.full.pdf>.
- [40] J. A. Fletcher and M. Zwick, *Unifying the theories of inclusive fitness and reciprocal altruism*, *The American Naturalist* **168**, 252 (2006).
- [41] J. Berg, J. Dickhaut, and K. McCabe, *Trust, reciprocity, and social history*, *Games and Economic Behavior* **10**, 122 (1995).
- [42] C. Bicchieri, *The grammar of society: The nature and dynamics of social norms* (Cambridge University Press, 2006).
- [43] S. Van Segbroeck, J. M. Pacheco, T. Lenaerts, and F. C. Santos, *Emergence of fairness in repeated group interactions*, *Phys. Rev. Lett.* **108**, 158104 (2012).
- [44] R. L. Trivers, *The evolution of reciprocal altruism*, *The Quarterly Review of Biology* **46**, 35 (1971).
- [45] H. Gintis, *Game Theory Evolving: A Problem-centered Introduction to Modeling Strategic Behavior*, Economics / Princeton University Press (Princeton University Press, 2000).
- [46] E. Fehr, U. Fischbacher, and S. Gächter, *Strong reciprocity, human cooperation, and the enforcement of social norms*, *Human Nature* **13**, 1 (2002).

- [47] U. Segal and J. Sobel, *Tit for tat: Foundations of preferences for reciprocity in strategic settings*, *Journal of Economic Theory* **136**, 197 (2007).
- [48] J. C. Cox, D. Friedman, and S. Gjerstad, *A tractable model of reciprocity and fairness*, *Games and Economic Behavior* **59**, 17 (2007).
- [49] M. Dufwenberg and G. Kirchsteiger, *A theory of sequential reciprocity*, *Games and Economic Behavior* **47**, 268 (2004).
- [50] M. Rabin, *Incorporating fairness into game theory and economics*, *The American Economic Review* **83**, pp. 1281 (1993).
- [51] D. Kahneman, *Thinking, Fast and Slow* (Farrar, Straus and Giroux, 2011).
- [52] *Wikipedia:Ten Simple Rules for Editing Wikipedia*, https://en.wikipedia.org/wiki/Wikipedia:Ten_Simple_Rules_for_Editing_Wikipedia (2016).
- [53] G. W. Brown, *Iterative solution of games by fictitious play*, *Activity analysis of production and allocation* **13**, 374 (1951).
- [54] G. Polevoy, S. Trajanovski, and M. M. de Weerd, *Nash Equilibria in Shared Effort Games*, in *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems, AAMAS '14* (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2014) pp. 861–868.
- [55] G. Polevoy, M. de Weerd, and C. Jonker, *The convergence of reciprocation*, in *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems, AAMAS '16* (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2016) pp. 1431–1432.
- [56] G. Polevoy, M. de Weerd, and C. Jonker, *The game of reciprocation habits*, in *Proceedings of the 2016 European Conference on Artificial Intelligence*, *Frontiers in Artificial Intelligence and Applications*, Vol. Volume 285: ECAI 2016 (2016) pp. 417–425.
- [57] E. Koutsoupias and C. Papadimitriou, *Worst-case equilibria*, in *16th Annual Symposium on Theoretical Aspects of Computer Science* (Trier, Germany, 1999) pp. 404–413.
- [58] C. Papadimitriou, *Algorithms, Games, and the Internet*, in *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing*, STOC '01 (ACM, New York, NY, USA, 2001) pp. 749–753.
- [59] A. S. Schulz and N. S. Moses, *On the performance of user equilibria in traffic networks*, in *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '03 (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2003) pp. 86–87.

- [60] E. Anshelevich, A. DasGupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden, *The price of stability for network design with fair cost allocation*, in *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on* (2004) pp. 295–304.
- [61] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, *Algorithmic Game Theory* (Cambridge University Press, 2007).
- [62] J. Kleinberg and E. Tardos, *Balanced outcomes in social exchange networks*, in *Proceedings of the 40th annual ACM symposium on Theory of computing* (ACM, New York, NY, USA, 2008) pp. 295–304.
- [63] *Course grading policy*, <http://www.physics.umd.edu/courses/Phys121/Griffin/grading.html> (2016).
- [64] B. D. Nguyen and K. M. Nielsen, *What death can tell: Are executives paid for their contributions to firm value?* AFA Denver Meetings Paper (2011).
- [65] E. Koutsoupias and C. H. Papadimitriou, *Worst-case equilibria*, *Computer Science Review* **3**, 65 (2009).
- [66] S.-f. Cheng, D. M. Reeves, Y. Vorobeychik, and M. P. Wellman, *Notes on Equilibria in Symmetric Games*, in *Proceedings of the 6th International Workshop On Game Theoretic And Decision Theoretic Agents* (2004) pp. 1–5.
- [67] D. Monderer and L. S. Shapley, *Potential games*, *Games and Economic Behavior* **14**, 124 (1996).
- [68] I. L. Glicksberg, *A Further Generalization of the Kakutani Fixed Point Theorem, with Application to Nash Equilibrium Points*, *Proceedings of the American Mathematical Society* **3**, 170 (1952).
- [69] E. Maskin and P. Dasgupta, *The existence of equilibrium in discontinuous economic games, part i (theory)*, *Review of Economic Studies* **53**, 1 (1986), reprinted in K. Binmore and P. Dasgupta (eds.), *Economic Organizations as Games*, Oxford: Basil Blackwell, 1986, pp. 48-82.
- [70] R. J. Aumann, *Subjectivity and correlation in randomized strategies*, *Journal of Mathematical Economics* **1**, 67 (1974).
- [71] H. Moulin and J. P. Vial, *Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon*, *International Journal of Game Theory* **7**, 201 (1978).
- [72] J. Robinson, *An iterative method of solving a game*, *Annals of Mathematics Second Series*, **54**, pp. 296 (1951).
- [73] J. M. Danskin, *Fictitious play for continuous games revisited*, *International Journal of Game Theory* **10**, 147 (1981).

- [74] Y. Zick and E. Elkind, *Arbitrators in overlapping coalition formation games*, in *AA-MAS*, edited by L. Sonenberg, P. Stone, K. Tumer, and P. Yolum (IFAAMAS, 2011) pp. 55–62.
- [75] Y. Zick, G. Chalkiadakis, and E. Elkind, *Overlapping coalition formation games: charting the tractability frontier*, in *AAMAS*, edited by W. van der Hoek, L. Padgham, V. Conitzer, and M. Winikoff (IFAAMAS, 2012) pp. 787–794.
- [76] S. Mannor, *Advanced topics in systems, learning and control*, Lecture notes (2008).
- [77] D. Monderer and L. S. Shapley, *Fictitious play property for games with identical interests*, *Journal of Economic Theory* **68**, 258 (1996).
- [78] L. Shapley and R. C. S. M. CALIF, *Some Topics in Two-person Games*, Memorandum (Rand Corporation) (Rand Corporation, 1963).
- [79] G. K. Batchelor, *Preoccupations of a journal editor*, *Journal of Fluid Mechanics* **106**, 1 (1981).
- [80] M. A. Nowak and K. Sigmund, *Evolution of indirect reciprocity*, *Nature* **437**, 1291 (2005).
- [81] B. W. Roberts, K. E. Walton, and W. Viechtbauer, *Patterns of mean-level change in personality traits across the life course: a meta-analysis of longitudinal studies*. *Psychological bulletin* **132**, 1 (2006).
- [82] G. Hofstede, *Culture and organizations*, *International Studies of Management & Organization* **10**, 15 (1980).
- [83] P. Breitman and C. Hatch, *How to Say No Without Feeling Guilty: And Say Yes to More Time, More Joy, and what Matters Most to You* (Broadway Books, 2000).
- [84] W. Ury, *The Power of a Positive No: How to Say No and Still Get to Yes* (Random House Publishing Group, 2007).
- [85] H. McAllister, *Who You Are Is What You Do: Making Choices About Life After School*, illustrated ed. (Wilkins Farago Pty Ltd, 2013).
- [86] V. Blondel, J. Hendrickx, A. Olshevsky, and J. Tsitsiklis, *Convergence in multiagent coordination, consensus, and flocking*, in *Decision and Control, 2005 and 2005 European Control Conference. CDC-ECC '05. 44th IEEE Conference on* (2005) pp. 2996–3000.
- [87] L. Moreau, *Stability of multiagent systems with time-dependent communication links*, *Automatic Control, IEEE Transactions on* **50**, 169 (2005).
- [88] J. Tsitsiklis, D. Bertsekas, and M. Athans, *Distributed asynchronous deterministic and stochastic gradient optimization algorithms*, *Automatic Control, IEEE Transactions on* **31**, 803 (1986).

- [89] M. H. DeGroot, *Reaching a consensus*, [Journal of the American Statistical Association](#) **69**, 118 (1974).
- [90] D. Bindel, J. Kleinberg, and S. Oren, *How bad is forming your own opinion?* in [Foundations of Computer Science \(FOCS\), 2011 IEEE 52nd Annual Symposium on](#) (2011) pp. 57–66.
- [91] J. S. Rosenschein and G. Zlotkin, *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers* (MIT Press, Cambridge, MA, USA, 1994).
- [92] A. Rubinstein, *Perfect equilibrium in a bargaining model*, [Econometrica](#) **50**, 97 (1982).
- [93] B. K. Butler and P. H. Siegel, *Sharp bounds on the spectral radius of nonnegative matrices and digraphs*, [Linear Algebra and its Applications](#) **439**, 1468 (2013).
- [94] E. Seneta, [Non-negative Matrices and Markov Chains](#), Springer Series in Statistics (Springer, 2006).
- [95] E. Hewitt and K. Stromberg, [Real and Abstract Analysis: A Modern Treatment of the Theory of Functions of a Real Variable](#), Graduate Texts in Mathematics (Springer New York, 1975).
- [96] D. Pruitt, [Negotiation Behavior](#), Library and Information Science (Academic Press, 1981).
- [97] P.-A. Chen, B. de Keijzer, D. Kempe, and G. Schäfer, *The robust price of anarchy of altruistic games*, in [Internet and Network Economics](#), Lecture Notes in Computer Science, Vol. 7090, edited by N. Chen, E. Elkind, and E. Koutsoupias (Springer Berlin Heidelberg, 2011) pp. 383–390.
- [98] M. Hoefer and A. Skopalik, *Altruism in atomic congestion games*, [ACM Trans. Econ. Comput.](#) **1**, 21:1 (2013).
- [99] M. Rahn and G. Schäfer, *Bounding the inefficiency of altruism through social contribution games*, in [Web and Internet Economics](#), Lecture Notes in Computer Science, Vol. 8289, edited by Y. Chen and N. Immorlica (Springer Berlin Heidelberg, 2013) pp. 391–404.
- [100] J. O. Ledyard, [Public Goods: A Survey of Experimental Research](#), Public Economics 9405003 (EconWPA, 1994).
- [101] H. Spencer, *Why reciprocity is key to building business relationships*, Business 2 Community (2012).
- [102] W. Mastenbroek, [Onderhandelen](#) (Het Spectrum, 1992).
- [103] D. Abreu and A. Rubinstein, *The structure of Nash equilibrium in repeated games with finite automata*, [Econometrica](#) **56**, 1259 (1988).

- [104] A. Rubinstein, *Finite automata play the repeated prisoner's dilemma*, *Journal of Economic Theory* **39**, 83 (1986).
- [105] K. G. Binmore and L. Samuelson, *Evolutionary stability in repeated games played by finite automata*, *Journal of Economic Theory* **57**, 278 (1992).
- [106] M. R. B. Giovanna Axia, *Linguistic politeness at different age levels*, *Child Development* **56**, 918 (1985).
- [107] V. A. Ilyin and E. G. Poznyak, *Fundamentals of mathematical analysis (Osnovi matematicheskogo analiza)*, book 1, 7th ed. (Physical and Mathematical Literature, 2005) in Russian.
- [108] N. H. Nie and D. S. Hillygus, *Where does internet time come from?: A reconnaissance*, *IT & Society* **1**, 1 (2002).
- [109] P.-A. Chen, M. David, and D. Kempe, *Better vaccination strategies for better people*, in *Proceedings of the 11th ACM Conference on Electronic Commerce*, EC '10 (ACM, New York, NY, USA, 2010) pp. 179–188.
- [110] N. H. Nie and D. S. Hillygus, *The impact of internet use on sociability: Time-diary findings*, *It & Society* **1**, 1 (2002).
- [111] V. Vazirani, *Approximation Algorithms* (Springer, 2001).
- [112] S. T. Ennett and K. E. Bauman, *Adolescent social networks: Friendship cliques, social isolates, and drug use risk*, *Improving prevention effectiveness*, 47 (2000).
- [113] R. Selten, *Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragefragheit*, *Zeitschrift fur die gesamte Staatswissenschaft* **121**, 301 (1965).
- [114] G. Chalkiadakis, E. Elkind, E. Markakis, M. Polukarov, and N. R. Jennings, *Cooperative games with overlapping coalitions*, *J. Artif. Int. Res.* **39**, 179 (2010).
- [115] J. H. Richards, *Time management—a review*, *Work & Stress* **1**, 73 (1987), <http://dx.doi.org/10.1080/02678378708258485> .
- [116] B. J. Claessens, W. van Eerde, C. G. Rutte, and R. A. Roe, *A review of the time management literature*, *Personnel Review* **36**, 255 (2007), <http://dx.doi.org/10.1108/00483480710726136> .

SIKS DISSERTATION SERIES

2009

- 1 Rasa Jurgelenaite (RUN) *Symmetric Causal Independence Models*
- 2 Willem Robert van Hage (VU) *Evaluating Ontology-Alignment Techniques*
- 3 Hans Stol (UvT) *A Framework for Evidence-based Policy Making Using IT*
- 4 Josephine Nabukenya (RUN) *Improving the Quality of Organisational Policy Making using Collaboration Engineering*
- 5 Sietse Overbeek (RUN) *Bridging Supply and Demand for Knowledge Intensive Tasks—Based on Knowledge, Cognition, and Quality*
- 6 Muhammad Subianto (UU) *Understanding Classification*
- 7 Ronald Poppe (UT) *Discriminative Vision-Based Recovery and Recognition of Human Motion*
- 8 Volker Nannen (VU) *Evolutionary Agent-Based Policy Analysis in Dynamic Environments*
- 9 Benjamin Kanagwa (RUN) *Design, Discovery and Construction of Service-oriented Systems*
- 10 Jan Wielemaker (UVA) *Logic programming for knowledge-intensive interactive applications*
- 11 Alexander Boer (UVA) *Legal Theory, Sources of Law & the Semantic Web*
- 12 Peter Massuthe (TUE, Humboldt-Universitaet zu Berlin) *Operating Guidelines for Services*
- 13 Steven de Jong (UM) *Fairness in Multi-Agent Systems*
- 14 Maksym Korotkiy (VU) *From ontology-enabled services to service-enabled ontologies (making ontologies work in e-science with ONTO-SOA)*
- 15 Rinke Hoekstra (UVA) *Ontology Representation—Design Patterns and Ontologies that Make Sense*
- 16 Fritz Reul (UvT) *New Architectures in Computer Chess*
- 17 Laurens van der Maaten (UvT) *Feature Extraction from Visual Data*
- 18 Fabian Groffen (CWI) *Armada, An Evolving Database System*
- 19 Valentin Robu (CWI) *Modeling Preferences, Strategic Reasoning and Collaboration in Agent-Mediated Electronic Markets*
- 20 Bob van der Vecht (UU) *Adjustable Autonomy: Controlling Influences on Decision Making*
- 21 Stijn Vanderlooy (UM) *Ranking and Reliable Classification*
- 22 Pavel Serdyukov (UT) *Search For Expertise: Going beyond direct evidence*
- 23 Peter Hofgesang (VU) *Modelling Web Usage in a Changing Environment*
- 24 Annerieke Heuvelink (VUA) *Cognitive Models for Training Simulations*
- 25 Alex van Ballegooij (CWI) *"RAM: Array Database Management through Relational Mapping"*
- 26 Fernando Koch (UU) *An Agent-Based Model for the Development of Intelligent Mobile Services*
- 27 Christian Glahn (OU) *Contextual Support of social Engagement and Reflection on the Web*
- 28 Sander Evers (UT) *Sensor Data Management with Probabilistic Models*
- 29 Stanislav Pokraev (UT) *Model-Driven Semantic Integration of Service-Oriented Applications*
- 30 Marcin Zukowski (CWI) *Balancing vectorized query execution with bandwidth-optimized storage*
- 31 Sofiya Katrenko (UVA) *A Closer Look at Learning Relations from Text*
- 32 Rik Farenhorst (VU) and Remco de Boer (VU) *Architectural Knowledge Management: Supporting Architects and Auditors*
- 33 Khiet Truong (UT) *How Does Real Affect Affect Affect Recognition In Speech?*
- 34 Inge van de Weerd (UU) *Advancing in Software Product Management: An Incremental Method Engineering Approach*
- 35 Wouter Koelewijn (UL) *Privacy en Politiegegevens; Over geautomatiseerde normatieve informatie-uitwisseling*
- 36 Marco Kalz (OUN) *Placement Support for Learners in Learning Networks*
- 37 Hendrik Drachsler (OUN) *Navigation Support for Learners in Informal Learning Networks*
- 38 Riina Vuorikari (OU) *Tags and self-organisation: a metadata ecology for learning resources in a multilingual context*

- 39 Christian Stahl (TUE, Humboldt-Universitaet zu Berlin) *Service Substitution—A Behavioral Approach Based on Petri Nets*
- 40 Stephan Raaijmakers (UvT) *Multinomial Language Learning: Investigations into the Geometry of Language*
- 41 Igor Berezhnyy (UvT) *Digital Analysis of Paintings*
- 42 Toine Bogers (UvT) *Recommender Systems for Social Bookmarking*
- 43 Virginia Nunes Leal Franqueira (UT) *Finding Multi-step Attacks in Computer Networks using Heuristic Search and Mobile Ambients*
- 44 Roberto Santana Tapia (UT) *Assessing Business-IT Alignment in Networked Organizations*
- 45 Jilles Vreeken (UU) *Making Pattern Mining Useful*
- 46 Loredana Afanasiev (UvA) *Querying XML: Benchmarks and Recursion*

2010

- 1 Matthijs van Leeuwen (UU) *Patterns that Matter*
- 2 Ingo Wassink (UT) *Work flows in Life Science*
- 3 Joost Geurts (CWI) *A Document Engineering Model and Processing Framework for Multimedia documents*
- 4 Olga Kulyk (UT) *Do You Know What I Know? Situational Awareness of Co-located Teams in Multidisplay Environments*
- 5 Claudia Hauff (UT) *Predicting the Effectiveness of Queries and Retrieval Systems*
- 6 Sander Bakkes (UvT) *Rapid Adaptation of Video Game AI*
- 7 Wim Fikkert (UT) *Gesture interaction at a Distance*
- 8 Krzysztof Siewicz (UL) *Towards an Improved Regulatory Framework of Free Software. Protecting user freedoms in a world of software communities and eGovernments*
- 9 Hugo Kielman (UL) *A Politiele gegevensverwerking en Privacy, Naar een effectieve waarborging*
- 10 Rebecca Ong (UL) *Mobile Communication and Protection of Children*
- 11 Adriaan Ter Mors (TUD) *The world according to MARP: Multi-Agent Route Planning*
- 12 Susan van den Braak (UU) *Sensemaking software for crime analysis*
- 13 Gianluigi Folino (RUN) *High Performance Data Mining using Bio-inspired techniques*
- 14 Sander van Splunter (VU) *Automated Web Service Reconfiguration*
- 15 Lianne Bodenstaff (UT) *Managing Dependency Relations in Inter-Organizational Models*
- 16 Sicco Verwer (TUD) *Efficient Identification of Timed Automata, theory and practice*
- 17 Spyros Kotoulas (VU) *Scalable Discovery of Networked Resources: Algorithms, Infrastructure, Applications*
- 18 Charlotte Gerritsen (VU) *Caught in the Act: Investigating Crime by Agent-Based Simulation*
- 19 Henriette Cramer (UvA) *People's Responses to Autonomous and Adaptive Systems*
- 20 Ivo Swartjes (UT) *Whose Story Is It Anyway? How Improv Informs Agency and Authorship of Emergent Narrative*
- 21 Harold van Heerde (UT) *Privacy-aware data management by means of data degradation*
- 22 Michiel Hildebrand (CWI) *End-user Support for Access to Heterogeneous Linked Data*
- 23 Bas Steunebrink (UU) *The Logical Structure of Emotions*
- 24 Dmytro Tykhonov *Designing Generic and Efficient Negotiation Strategies*
- 25 Zulfiqar Ali Memon (VU) *Modelling Human-Awareness for Ambient Agents: A Human Mindreading Perspective*
- 26 Ying Zhang (CWI) *XRPC: Efficient Distributed Query Processing on Heterogeneous XQuery Engines*
- 27 Marten Voulon (UL) *Automatisch contracteren*
- 28 Arne Koopman (UU) *Characteristic Relational Patterns*
- 29 Stratos Idreos (CWI) *Database Cracking: Towards Auto-tuning Database Kernels*
- 30 Marieke van Erp (UvT) *Accessing Natural History—Discoveries in data cleaning, structuring, and retrieval*
- 31 Victor de Boer (UVA) *Ontology Enrichment from Heterogeneous Sources on the Web*
- 32 Marcel Hiel (UvT) *An Adaptive Service Oriented Architecture: Automatically solving Interoperability Problems*
- 33 Robin Aly (UT) *Modeling Representation Uncertainty in Concept-Based Multimedia Retrieval*
- 34 Teduh Dirgahayu (UT) *Interaction Design in Service Compositions*
- 35 Dolf Trieschnigg (UT) *Proof of Concept: Concept-based Biomedical Information Retrieval*
- 36 Jose Janssen (OU) *Paving the Way for Lifelong Learning: Facilitating competence development through a learning path specification*
- 37 Niels Lohmann (TUE) *Correctness of services and their composition*
- 38 Dirk Fahland (TUE) *From Scenarios to components*

- 39 Ghazanfar Farooq Siddiqui (VU) *Integrative modeling of emotions in virtual agents*
- 40 Mark van Assem (VU) *Converting and Integrating Vocabularies for the Semantic Web*
- 41 Guillaume Chaslot (UM) *Monte-Carlo Tree Search*
- 42 Sybren de Kinderen (VU) *Needs-driven service bundling in a multi-supplier setting—the computational e3-service approach*
- 43 Peter van Kranenburg (UU) *A Computational Approach to Content-Based Retrieval of Folk Song Melodies*
- 44 Pieter Bellekens (TUE) *An Approach towards Context-sensitive and User-adapted Access to Heterogeneous Data Sources, Illustrated in the Television Domain*
- 45 Vasilios Andrikopoulos (UvT) *A theory and model for the evolution of software services*
- 46 Vincent Pijpers (VU) *e3alignment: Exploring Inter-Organizational Business-ICT Alignment*
- 47 Chen Li (UT) *Mining Process Model Variants: Challenges, Techniques, Examples*
- 48 *Withdrawn*
- 49 Jahn-Takeshi Saito (UM) *Solving difficult game positions*
- 50 Bouke Huurnink (UVA) *Search in Audiovisual Broadcast Archives*
- 51 Alia Khairia Amin (CWI) *Understanding and supporting information seeking tasks in multiple sources*
- 52 Peter-Paul van Maanen (VU) *Adaptive Support for Human-Computer Teams: Exploring the Use of Cognitive Models of Trust and Attention*
- 53 Edgar Meij (UVA) *Combining Concepts and Language Models for Information Access*

2011

- 1 Botond Cseke (RUN) *Variational Algorithms for Bayesian Inference in Latent Gaussian Models*
- 2 Nick Tinnemeier (UU) *Organizing Agent Organizations. Syntax and Operational Semantics of an Organization-Oriented Programming Language*
- 3 Jan Martijn van der Werf (TUE) *Compositional Design and Verification of Component-Based Information Systems*
- 4 Hado van Hasselt (UU) *Insights in Reinforcement Learning; Formal analysis and empirical evaluation of temporal-difference learning algorithms*
- 5 Base van der Raadt (VU) *Enterprise Architecture Coming of Age—Increasing the Performance of an Emerging Discipline.*
- 6 Yiwen Wang (TUE) *Semantically-Enhanced Recommendations in Cultural Heritage*
- 7 Yujia Cao (UT) *Multimodal Information Presentation for High Load Human Computer Interaction*
- 8 Nieske Vergunst (UU) *BDI-based Generation of Robust Task-Oriented Dialogues*
- 9 Tim de Jong (OU) *Contextualised Mobile Media for Learning*
- 10 Bart Bogaert (UvT) *Cloud Content Contention*
- 11 Dhaval Vyas (UT) *Designing for Awareness: An Experience-focused HCI Perspective*
- 12 Carmen Bratosin (TUE) *Grid Architecture for Distributed Process Mining*
- 13 Xiaoyu Mao (UvT) *Airport under Control. Multiagent Scheduling for Airport Ground Handling*
- 14 Milan Lovric (EUR) *Behavioral Finance and Agent-Based Artificial Markets*
- 15 Marijn Koolen (UvA) *The Meaning of Structure: the Value of Link Evidence for Information Retrieval*
- 16 Maarten Schadd (UM) *Selective Search in Games of Different Complexity*
- 17 Jiyin He (UVA) *Exploring Topic Structure: Coherence, Diversity and Relatedness*
- 18 Mark Ponsen (UM) *Strategic Decision-Making in complex games*
- 19 Ellen Rusman (OU) *The Mind 's Eye on Personal Profiles*
- 20 Qing Gu (VU) *Guiding service-oriented software engineering—A view-based approach*
- 21 Linda Terlouw (TUD) *Modularization and Specification of Service-Oriented Systems*
- 22 Junte Zhang (UVA) *System Evaluation of Archival Description and Access*
- 23 Wouter Weerkamp (UVA) *Finding People and their Utterances in Social Media*
- 24 Herwin van Welbergen (UT) *Behavior Generation for Interpersonal Coordination with Virtual Humans On Specifying, Scheduling and Realizing Multimodal Virtual Human Behavior*
- 25 Syed Waqar ul Qounain Jaffry (VU) *Analysis and Validation of Models for Trust Dynamics*
- 26 Matthijs Aart Pontier (VU) *Virtual Agents for Human Communication—Emotion Regulation and Involvement-Distance Trade-Offs in Embodied Conversational Agents and Robots*
- 27 Aniel Bhulai (VU) *Dynamic website optimization through autonomous management of design patterns*
- 28 Rianne Kaptein (UVA) *Effective Focused Retrieval by Exploiting Query Context and Document Structure*
- 29 Faisal Kamiran (TUE) *Discrimination-aware Classification*
- 30 Egon van den Broek (UT) *Affective Signal Processing (ASP): Unraveling the mystery of emotions*
- 31 Ludo Waltman (EUR) *Computational and Game-Theoretic Approaches for Modeling Bounded Rationality*

- 32 Nees-Jan van Eck (EUR) *Methodological Advances in Bibliometric Mapping of Science*
- 33 Tom van der Weide (UU) *Arguing to Motivate Decisions*
- 34 Paolo Turrini (UU) *Strategic Reasoning in Interdependence: Logical and Game-theoretical Investigations*
- 35 Maaïke Harbers (UU) *Explaining Agent Behavior in Virtual Training*
- 36 Erik van der Spek (JU) *Experiments in serious game design: a cognitive approach*
- 37 Adriana Burlutiu (RUN) *Machine Learning for Pairwise Data, Applications for Preference Learning and Supervised Network Inference*
- 38 Nyree Lemmens (UM) *Bee-inspired Distributed Optimization*
- 39 Joost Westra (UU) *Organizing Adaptation using Agents in Serious Games*
- 40 Viktor Clerc (VU) *Architectural Knowledge Management in Global Software Development*
- 41 Luan Ibraimi (UT) *Cryptographically Enforced Distributed Data Access Control*
- 42 Michal Sindlar (UU) *Explaining Behavior through Mental State Attribution*
- 43 Henk van der Schuur (UU) *Process Improvement through Software Operation Knowledge*
- 44 Boris Reuderink (UT) *Robust Brain-Computer Interfaces*
- 45 Herman Stehouwer (UvT) *Statistical Language Models for Alternative Sequence Selection*
- 46 Beibei Hu (TUD) *Towards Contextualized Information Delivery: A Rule-based Architecture for the Domain of Mobile Police Work*
- 47 Azizi Bin Ab Aziz (VU) *Exploring Computational Models for Intelligent Support of Persons with Depression*
- 48 Mark Ter Maat (UT) *Response Selection and Turn-taking for a Sensitive Artificial Listening Agent*
- 49 Andreea Niculescu (UT) *Conversational interfaces for task-oriented spoken dialogues: design aspects influencing interaction quality*

2012

- 1 Terry Kakeeto (UvT) *Relationship Marketing for SMEs in Uganda*
- 2 Muhammad Umair (VU) *Adaptivity, emotion, and Rationality in Human and Ambient Agent Models*
- 3 Adam Vanya (VU) *Supporting Architecture Evolution by Mining Software Repositories*
- 4 Jurriaan Souer (UU) *Development of Content Management System-based Web Applications*
- 5 Marijn Plomp (UU) *Maturing Interorganisational Information Systems*
- 6 Wolfgang Reinhardt (OU) *Awareness Support for Knowledge Workers in Research Networks*
- 7 Rianne van Lambalgen (VU) *When the Going Gets Tough: Exploring Agent-based Models of Human Performance under Demanding Conditions*
- 8 Gerben de Vries (UVA) *Kernel Methods for Vessel Trajectories*
- 9 Ricardo Neisse (UT) *Trust and Privacy Management Support for Context-Aware Service Platforms*
- 10 David Smits (TUE) *Towards a Generic Distributed Adaptive Hypermedia Environment*
- 11 J.C.B. Rantham Prabhakara (TUE) *Process Mining in the Large: Preprocessing, Discovery, and Diagnostics*
- 12 Kees van der Sluijs (TUE) *Model Driven Design and Data Integration in Semantic Web Information Systems*
- 13 Suleman Shahid (UvT) *Fun and Face: Exploring non-verbal expressions of emotion during playful interactions*
- 14 Evgeny Knutov (TUE) *Generic Adaptation Framework for Unifying Adaptive Web-based Systems*
- 15 Natalie van der Wal (VU) *Social Agents. Agent-Based Modelling of Integrated Internal and Social Dynamics of Cognitive and Affective Processes.*
- 16 Fiemke Both (VU) *Helping people by understanding them—Ambient Agents supporting task execution and depression treatment*
- 17 Amal Elgammal (UvT) *Towards a Comprehensive Framework for Business Process Compliance*
- 18 Eltjo Poort (VU) *Improving Solution Architecting Practices*
- 19 Helen Schonenberg (TUE) *What's Next? Operational Support for Business Process Execution*
- 20 Ali Bahramisharif (RUN) *Covert Visual Spatial Attention, a Robust Paradigm for Brain-Computer Interfacing*
- 21 Roberto Cornacchia (TUD) *Querying Sparse Matrices for Information Retrieval*
- 22 Thijs Vis (UvT) *Intelligence, politie en veiligheidsdienst: verenigbare grootheden?*
- 23 Christian Muehl (UT) *Toward Affective Brain-Computer Interfaces: Exploring the Neurophysiology of Affect during Human Media Interaction*
- 24 Laurens van der Werff (UT) *Evaluation of Noisy Transcripts for Spoken Document Retrieval*
- 25 Silja Eckartz (UT) *Managing the Business Case Development in Inter-Organizational IT Projects: A Methodology and its Application*
- 26 Emile de Maat (UVA) *Making Sense of Legal Text*
- 27 Hayrettin Gurkok (UT) *Mind the Sheep! User Experience Evaluation & Brain-Computer Interface Games*

- 28 Nancy Pascall (UvT) *Engendering Technology Empowering Women*
- 29 Almer Tigelaar (UT) *Peer-to-Peer Information Retrieval*
- 30 Alina Pommeranz (TUD) *Designing Human-Centered Systems for Reflective Decision Making*
- 31 Emily Bagarukayo (RUN) *A Learning by Construction Approach for Higher Order Cognitive Skills Improvement, Building Capacity and Infrastructure*
- 32 Wietske Visser (TUD) *Qualitative multi-criteria preference representation and reasoning*
- 33 Rory Sie (OUN) *Coalitions in Cooperation Networks (COCOON)*
- 34 Pavol Jancura (RUN) *Evolutionary analysis in PPI networks and applications*
- 35 Evert Haasdijk (VU) *Never Too Old To Learn—On-line Evolution of Controllers in Swarm- and Modular Robotics*
- 36 Denis Ssebugwawo (RUN) *Analysis and Evaluation of Collaborative Modeling Processes*
- 37 Agnes Nakakawa (RUN) *A Collaboration Process for Enterprise Architecture Creation*
- 38 Selmar Smit (VU) *Parameter Tuning and Scientific Testing in Evolutionary Algorithms*
- 39 Hassan Fatemi (UT) *Risk-aware design of value and coordination networks*
- 40 Agus Gunawan (UvT) *Information Access for SMEs in Indonesia*
- 41 Sebastian Kelle (OU) *Game Design Patterns for Learning*
- 42 Dominique Verpoorten (OU) *Reflection Amplifiers in self-regulated Learning*
- 43 *Withdrawn*
- 44 Anna Tordai (VU) *On Combining Alignment Techniques*
- 45 Benedikt Kratz (UvT) *A Model and Language for Business-aware Transactions*
- 46 Simon Carter (UVA) *Exploration and Exploitation of Multilingual Data for Statistical Machine Translation*
- 47 Manos Tsagkias (UVA) *Mining Social Media: Tracking Content and Predicting Behavior*
- 48 Jorn Bakker (TUE) *Handling Abrupt Changes in Evolving Time-series Data*
- 49 Michael Kaisers (UM) *Learning against Learning—Evolutionary dynamics of reinforcement learning algorithms in strategic interactions*
- 50 Steven van Kervel (TUD) *Ontology driven Enterprise Information Systems Engineering*
- 51 Jeroen de Jong (TUD) *Heuristics in Dynamic Sceduling: a practical framework with a case study in elevator dispatching*

2013

- 1 Viorel Milea (EUR) *News Analytics for Financial Decision Support*
- 2 Erietta Liarou (CWI) *MonetDB/DataCell: Leveraging the Column-store Database Technology for Efficient and Scalable Stream Processing*
- 3 Szymon Klarman (VU) *Reasoning with Contexts in Description Logics*
- 4 Chetan Yadati (TUD) *Coordinating autonomous planning and scheduling*
- 5 Dulce Pumareja (UT) *Groupware Requirements Evolutions Patterns*
- 6 Romulo Goncalves (CWI) *The Data Cyclotron: Juggling Data and Queries for a Data Warehouse Audience*
- 7 Giel van Lankveld (UvT) *Quantifying Individual Player Differences*
- 8 Robbert-Jan Merk (VU) *Making enemies: cognitive modeling for opponent agents in fighter pilot simulators*
- 9 Fabio Gori (RUN) *Metagenomic Data Analysis: Computational Methods and Applications*
- 10 Jeewanie Jayasinghe Arachchige (UvT) *A Unified Modeling Framework for Service Design.*
- 11 Evangelos Pournaras (TUD) *Multi-level Reconfigurable Self-organization in Overlay Services*
- 12 Marian Razavian (VU) *Knowledge-driven Migration to Services*
- 13 Mohammad Safiri (UT) *Service Tailoring: User-centric creation of integrated IT-based homecare services to support independent living of elderly*
- 14 Jafar Tanha (UVA) *Ensemble Approaches to Semi-Supervised Learning Learning*
- 15 Daniel Hennes (UM) *Multiagent Learning—Dynamic Games and Applications*
- 16 Eric Kok (UU) *Exploring the practical benefits of argumentation in multi-agent deliberation*
- 17 Koen Kok (VU) *The PowerMatcher: Smart Coordination for the Smart Electricity Grid*
- 18 Jeroen Janssens (UvT) *Outlier Selection and One-Class Classification*
- 19 Renze Steenhuizen (TUD) *Coordinated Multi-Agent Planning and Scheduling*
- 20 Katja Hofmann (UvA) *Fast and Reliable Online Learning to Rank for Information Retrieval*
- 21 Sander Wubben (UvT) *Text-to-text generation by monolingual machine translation*
- 22 Tom Claassen (RUN) *Causal Discovery and Logic*
- 23 Patricio de Alencar Silva (UvT) *Value Activity Monitoring*
- 24 Haitham Bou Ammar (UM) *Automated Transfer in Reinforcement Learning*
- 25 Agnieszka Anna Latoszek-Berendsen (UM) *Intention-based Decision Support. A new way of representing*

- and implementing clinical guidelines in a Decision Support System*
- 26 Alireza Zarghami (UT) *Architectural Support for Dynamic Homecare Service Provisioning*
- 27 Mohammad Huq (UT) *Inference-based Framework Managing Data Provenance*
- 28 Frans van der Sluis (UT) *When Complexity becomes Interesting: An Inquiry into the Information eXperience*
- 29 Iwan de Kok (UT) *Listening Heads*
- 30 Joyce Nakatumba (TUE) *Resource-Aware Business Process Management: Analysis and Support*
- 31 Dinh Khoa Nguyen (UvT) *Blueprint Model and Language for Engineering Cloud Applications*
- 32 Kamakshi Rajagopal (OUN) *Networking For Learning: The role of Networking in a Lifelong Learner's Professional Development*
- 33 Qi Gao (TUD) *User Modeling and Personalization in the Microblogging Sphere*
- 34 Kien Tjin-Kam-Jet (UT) *Distributed Deep Web Search*
- 35 Abdallah El Ali (UvA) *Minimal Mobile Human Computer Interaction*
- 36 Than Lam Hoang (TUE) *Pattern Mining in Data Streams*
- 37 Dirk Börner (OUN) *Ambient Learning Displays*
- 38 Eelco den Heijer (VU) *Autonomous Evolutionary Art*
- 39 Joop de Jong (TUD) *A Method for Enterprise Ontology based Design of Enterprise Information Systems*
- 40 Pim Nijssen (UM) *Monte-Carlo Tree Search for Multi-Player Games*
- 41 Jochem Liem (UVA) *Supporting the Conceptual Modelling of Dynamic Systems: A Knowledge Engineering Perspective on Qualitative Reasoning*
- 42 Léon Planken (TUD) *Algorithms for Simple Temporal Reasoning*
- 43 Marc Bron (UVA) *Exploration and Contextualization through Interaction and Concepts*

2014

- 1 Nicola Barile (UU) *Studies in Learning Monotone Models from Data*
- 2 Fiona Tullyano (RUN) *Combining System Dynamics with a Domain Modeling Method*
- 3 Sergio Raul Duarte Torres (UT) *Information Retrieval for Children: Search Behavior and Solutions*
- 4 Hanna Jochmann-Mannak (UT) *Websites for children: search strategies and interface design—Three studies on children's search performance and evaluation*
- 5 Jurriaan van Reijssen (UU) *Knowledge Perspectives on Advancing Dynamic Capability*
- 6 Damian Tamburri (VU) *Supporting Networked Software Development*
- 7 Arya Adriansyah (TUE) *Aligning Observed and Modeled Behavior*
- 8 Samur Araujo (TUD) *Data Integration over Distributed and Heterogeneous Data Endpoints*
- 9 Philip Jackson (UvT) *Toward Human-Level Artificial Intelligence: Representation and Computation of Meaning in Natural Language*
- 10 Ivan Salvador Razo Zapata (VU) *Service Value Networks*
- 11 Janneke van der Zwaan (TUD) *An Empathic Virtual Buddy for Social Support*
- 12 Willem van Willigen (VU) *Look Ma, No Hands: Aspects of Autonomous Vehicle Control*
- 13 Arlette van Wissen (VU) *Agent-Based Support for Behavior Change: Models and Applications in Health and Safety Domains*
- 14 Yangyang Shi (TUD) *Language Models With Meta-information*
- 15 Natalya Mogles (VU) *Agent-Based Analysis and Support of Human Functioning in Complex Socio-Technical Systems: Applications in Safety and Healthcare*
- 16 Krystyna Milian (VU) *Supporting trial recruitment and design by automatically interpreting eligibility criteria*
- 17 Kathrin Dentler (VU) *Computing healthcare quality indicators automatically: Secondary Use of Patient Data and Semantic Interoperability*
- 18 Mattijs Ghijsen (VU) *Methods and Models for the Design and Study of Dynamic Agent Organizations*
- 19 Vinicius Ramos (TUE) *Adaptive Hypermedia Courses: Qualitative and Quantitative Evaluation and Tool Support*
- 20 Mena Habib (UT) *Named Entity Extraction and Disambiguation for Informal Text: The Missing Link*
- 21 Cassidy Clark (TUD) *Negotiation and Monitoring in Open Environments*
- 22 Marieke Peeters (UU) *Personalized Educational Games—Developing agent-supported scenario-based training*
- 23 Eleftherios Sidirourgos (UvA/CWI) *Space Efficient Indexes for the Big Data Era*
- 24 Davide Ceolin (VU) *Trusting Semi-structured Web Data*
- 25 Martijn Lappenschaar (RUN) *New network models for the analysis of disease interaction*
- 26 Tim Baarslag (TUD) *What to Bid and When to Stop*

- 27 Rui Jorge Almeida (EUR) *Conditional Density Models Integrating Fuzzy and Probabilistic Representations of Uncertainty*
- 28 Anna Chmielowiec (VU) *Decentralized k-Clique Matching*
- 29 Jaap Kabbedijk (UU) *Variability in Multi-Tenant Enterprise Software*
- 30 Peter de Cock (UvT) *Anticipating Criminal Behaviour*
- 31 Leo van Moergestel (UU) *Agent Technology in Agile Multiparallel Manufacturing and Product Support*
- 32 Naser Ayat (UvA) *On Entity Resolution in Probabilistic Data*
- 33 Tesfa Tegegne (RUN) *Service Discovery in eHealth*
- 34 Christina Manteli (VU) *The Effect of Governance in Global Software Development: Analyzing Transactive Memory Systems.*
- 35 Joost van Ooijen (UU) *Cognitive Agents in Virtual Worlds: A Middleware Design Approach*
- 36 Joos Buijs (TUE) *Flexible Evolutionary Algorithms for Mining Structured Process Models*
- 37 Maral Dadvar (UT) *Experts and Machines United Against Cyberbullying*
- 38 Danny Plass-Oude Bos (UT) *Making brain-computer interfaces better: improving usability through post-processing.*
- 39 Jasmina Maric (UvT) *Web Communities, Immigration, and Social Capital*
- 40 Walter Omona (RUN) *A Framework for Knowledge Management Using ICT in Higher Education*
- 41 Frederic Hogenboom (EUR) *Automated Detection of Financial Events in News Text*
- 42 Carsten Eijckhof (CWI/TUD) *Contextual Multidimensional Relevance Models*
- 43 Kevin Vlaanderen (UU) *Supporting Process Improvement using Method Increments*
- 44 Paulien Meesters (UvT) *Intelligent Blauw. Met als ondertitel: Intelligence-gestuurde politiezorg in gebiedsgebonden eenheden.*
- 45 Birgit Schmitz (OUN) *Mobile Games for Learning: A Pattern-Based Approach*
- 46 Ke Tao (TUD) *Social Web Data Analytics: Relevance, Redundancy, Diversity*
- 47 Shangsong Liang (UVA) *Fusion and Diversification in Information Retrieval*

2015

- 1 Niels Netten (UvA) *Machine Learning for Relevance of Information in Crisis Response*
- 2 Faiza Bukhsh (UvT) *Smart auditing: Innovative Compliance Checking in Customs Controls*
- 3 Twan van Laarhoven (RUN) *Machine learning for network data*
- 4 Howard Spoelstra (OUN) *Collaborations in Open Learning Environments*
- 5 Christoph Bösch (UT) *Cryptographically Enforced Search Pattern Hiding*
- 6 Farideh Heidari (TUD) *Business Process Quality Computation—Computing Non-Functional Requirements to Improve Business Processes*
- 7 Maria-Hendrike Peetz (UvA) *Time-Aware Online Reputation Analysis*
- 8 Jie Jiang (TUD) *Organizational Compliance: An agent-based model for designing and evaluating organizational interactions*
- 9 Randy Klaassen (UT) *HCI Perspectives on Behavior Change Support Systems*
- 10 Henry Hermans (OUN) *OpenU: design of an integrated system to support lifelong learning*
- 11 Yongming Luo (TUE) *Designing algorithms for big graph datasets: A study of computing bisimulation and joins*
- 12 Julie M. Birkholz (VU) *Modi Operandi of Social Network Dynamics: The Effect of Context on Scientific Collaboration Networks*
- 13 Giuseppe Procaccianti (VU) *Energy-Efficient Software*
- 14 Bart van Straalen (UT) *A cognitive approach to modeling bad news conversations*
- 15 Klaas Andries de Graaf (VU) *Ontology-based Software Architecture Documentation*
- 16 Changyun Wei (UT) *Cognitive Coordination for Cooperative Multi-Robot Teamwork*
- 17 André van Cleeff (UT) *Physical and Digital Security Mechanisms: Properties, Combinations and Trade-offs*
- 18 Holger Pirk (CWI) *Waste Not, Want Not!—Managing Relational Data in Asymmetric Memories*
- 19 Bernardo Tabuenca (OUN) *Ubiquitous Technology for Lifelong Learners*
- 20 Lois Vanhée (UU) *Using Culture and Values to Support Flexible Coordination*
- 21 Sibren Fetter (OUN) *Using Peer-Support to Expand and Stabilize Online Learning*
- 22 Zhemín Zhu (UT) *Co-occurrence Rate Networks*
- 23 Luit Gazendam (VU) *Cataloguer Support in Cultural Heritage*
- 24 Richard Berendsen (UVA) *Finding People, Papers, and Posts: Vertical Search Algorithms and Evaluation*
- 25 Steven Woudenbergh (UU) *Bayesian Tools for Early Disease Detection*
- 26 Alexander Hogenboom (EUR) *Sentiment Analysis of Text Guided by Semantics and Structure*

- 27 Sándor Héman (CWI) *Updating compressed column-stores*
- 28 Janet Bagorogoza (TiU) *Knowledge Management and High Performance; The Uganda Financial Institutions Model for HPO*
- 29 Hendrik Baier (UM) *Monte-Carlo Tree Search Enhancements for One-Player and Two-Player Domains*
- 30 Kiavash Bahreini (OUN) *Real-time Multimodal Emotion Recognition in E-Learning*
- 31 Yakup Koç (TUD) *On Robustness of Power Grids*
- 32 Jerome Gard (UL) *Corporate Venture Management in SMEs*
- 33 Frederik Schadd (UM) *Ontology Mapping with Auxiliary Resources*
- 34 Victor de Graaff (UT) *Geosocial Recommender Systems*
- 35 Junchao Xu (TUD) *Affective Body Language of Humanoid Robots: Perception and Effects in Human Robot Interaction*

2016

- 1 Syed Saiden Abbas (RUN) *Recognition of Shapes by Humans and Machines*
- 2 Michiel Christiaan Meulendijk (UU) *Optimizing medication reviews through decision support: prescribing a better pill to swallow*
- 3 Maya Sappelli (RUN) *Knowledge Work in Context: User Centered Knowledge Worker Support*
- 4 Laurens Rietveld (VU) *Publishing and Consuming Linked Data*
- 5 Evgeny Sherkhonov (UVA) *Expanded Acyclic Queries: Containment and an Application in Explaining Missing Answers*
- 6 Michel Wilson (TUD) *Robust scheduling in an uncertain environment*
- 7 Jeroen de Man (VU) *Measuring and modeling negative emotions for virtual training*
- 8 Matje van de Camp (TiU) *A Link to the Past: Constructing Historical Social Networks from Unstructured Data*
- 9 Archana Nottamkandath (VU) *Trusting Crowdsourced Information on Cultural Artefacts*
- 10 George Karafotias (VUA) *Parameter Control for Evolutionary Algorithms*
- 11 Anne Schuth (UVA) *Search Engines that Learn from Their Users*
- 12 Max Knobbout (UU) *Logics for Modelling and Verifying Normative Multi-Agent Systems*
- 13 Nana Baah Gyan (VU) *The Web, Speech Technologies and Rural Development in West Africa—An ICT4D Approach*
- 14 Ravi Khadka (UU) *Revisiting Legacy Software System Modernization*
- 15 Steffen Michels (RUN) *Hybrid Probabilistic Logics—Theoretical Aspects, Algorithms and Experiments*
- 16 Guangliang Li (UVA) *Socially Intelligent Autonomous Agents that Learn from Human Reward*
- 17 Berend Weel (VU) *Towards Embodied Evolution of Robot Organisms*
- 18 Albert Meroño Peñuela (VU) *Refining Statistical Data on the Web*
- 19 Julia Efremova (Tu/e) *Mining Social Structures from Genealogical Data*
- 20 Daan Odijk (UVA) *Context & Semantics in News & Web Search*
- 21 Alejandro Moreno Célieri (UT) *From Traditional to Interactive Playspaces: Automatic Analysis of Player Behavior in the Interactive Tag Playground*
- 22 Grace Lewis (VU) *Software Architecture Strategies for Cyber-Foraging Systems*
- 23 Fei Cai (UVA) *Query Auto Completion in Information Retrieval*
- 24 Brend Wanders (UT) *Repurposing and Probabilistic Integration of Data; An Iterative and data model independent approach*
- 25 Julia Kiseleva (TU/e) *Using Contextual Information to Understand Searching and Browsing Behavior*
- 26 Dilhan Thilakarathne (VU) *In or Out of Control: Exploring Computational Models to Study the Role of Human Awareness and Control in Behavioural Choices, with Applications in Aviation and Energy Management Domains*
- 27 Wen Li (TUD) *Understanding Geo-spatial Information on Social Media*
- 28 Mingxin Zhang (TUD) *Large-scale Agent-based Social Simulation—A study on epidemic prediction and control*
- 29 Nicolas Höning (TUD) *Peak reduction in decentralised electricity systems -Markets and prices for flexible planning*
- 30 Ruud Mattheij (UvT) *The Eyes Have It*
- 31 Mohammad Khelghati (UT) *Deep web content monitoring*
- 32 Elco Vriezেকolk (UT) *Assessing Telecommunication Service Availability Risks for Crisis Organisations*
- 33 Peter Bloem (UVA) *Single Sample Statistics, exercises in learning from just one example*
- 34 Dennis Schunselaar (TUE) *Title: Configurable Process Trees: Elicitation, Analysis, and Enactment*

- 35 Zhaochun Ren *Monitoring Social Media: Summarization, Classification and Recommendation*
- 36 Daphne Karreman (UT) *Beyond R2D2: The design of nonverbal interaction behavior optimized for robot-specific morphologies*
- 37 Giovanni Sileno (UvA) *Aligning Law and Action—a conceptual and computational inquiry*
- 38 Andrea Minuto (UT) *MATERIALS THAT MATTER — Smart Materials meet Art & Interaction Design*
- 39 Merijn Bruijnes (UT) *Believable Suspect Agents; Response and Interpersonal Style Selection for an Artificial Suspect*
- 40 Christian Detweiler (TUD) *Accounting for Values in Design*
- 41 Thomas King (TUD) *Governing Governance: A Formal Framework for Analysing Institutional Design and Enactment Governance*
- 42 Spyros Martzoukos (UVA) *Combinatorial and Compositional Aspects of Bilingual Aligned Corpora*
- 43 Saskia Koldijk (RUN) *Context-Aware Support for Stress Self-Management: From Theory to Practice*
- 44 Thibault Sellam (UVA) *Automatic Assistants for Database Exploration*
- 45 Bram van de Laar (UT) *Experiencing Brain-Computer Interface Control*
- 46 Jorge Gallego Perez (UT) *Robots to Make you Happy*
- 47 Christina Weber (UL) *Real-time foresight—Preparedness for dynamic innovation networks*
- 48 Tanja Buttler (TUD) *Collecting Lessons Learned*
- 49 Gleb Polevoy (TUD) *Participation and Interaction in Projects: A Game-Theoretic Analysis*