BART DE KEIJZER

"COOPERATION AND EXTERNALITIES IN ALGORITHMIC GAME THEORY"

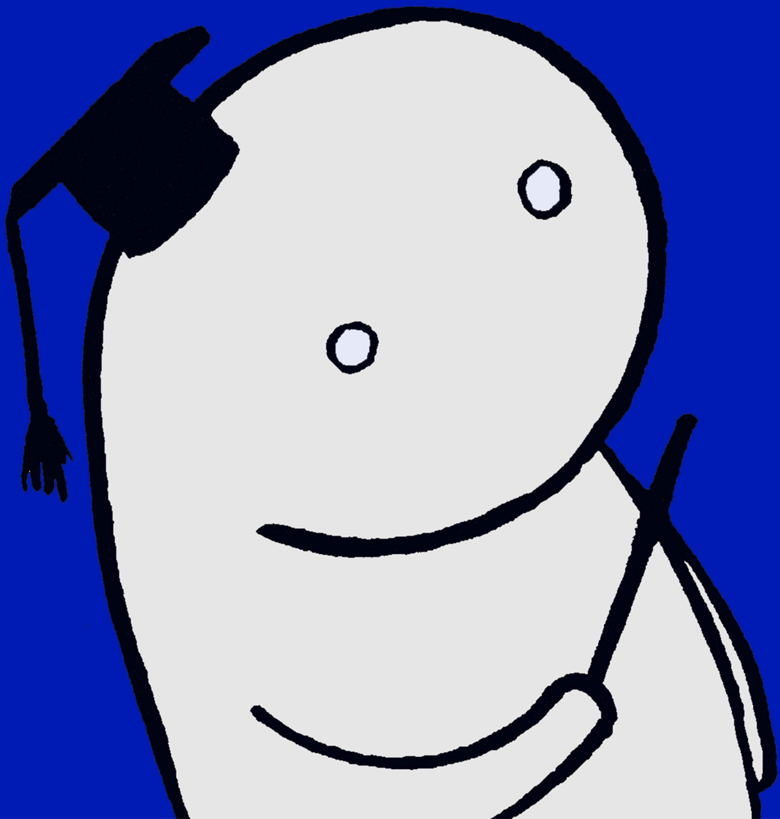

Cover page: A mysterious scholarly creature on an International Klein Blue background. Designed by Bram Gollin.

ISBN 978-94-6259-211-7

This research has been carried out at the Centrum Wiskunde & Informatica (CWI) in Amsterdam.

VRIJE UNIVERSITEIT

Externalities and Cooperation in Algorithmic Game Theory

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad Doctor aan
de Vrije Universiteit Amsterdam,
op gezag van de rector magnificus
prof.dr. F.A. van der Duyn Schouten,
in het openbaar te verdedigen
ten overstaan van de promotiecommissie
van de Faculteit der Economische Wetenschappen en Bedrijfskunde
op maandag 16 juni 2014 om 15.45 uur
in de aula van de universiteit,
De Boelelaan 1105

door

Bart de Keijzer

geboren te Amstelveen

promotor: prof.dr. G. Schäfer

promotor:	prof.dr. G. Schäfer
overige leden:	prof.dr. K.R. Apt prof.dr. E. Elkind prof.dr. S. Leonardi prof.dr. L. Stougie prof.dr. É. Tardos

Acknowledgements

First, I thank my supervisor, Guido Schäfer. Thanks to him, I enjoyed tremendously my time as a Ph.D. student. This experience could not be more pleasant and I thank him for the consistently great advice, support, collaboration, and supervision in general. I think we did some really interesting research together and hope that we will again do so in the future.

I thank Monique Laurent, the leader of the Networks and Optimization research group at CWI, for letting me be a group member and letting me do my research in her group. I also thank her for all her support and for the great policy to have lunch at 12:00 every day with the entire group. The lunches were very nice.

I thank my thesis committee: Krysztof Apt, Edith Elkind, Stefano Leonardi, Leen Stougie, and Éva Tardos. It is an honor for me that they are in my committee, and without them it would obviously not be possible for me to obtain my Ph.D. degree. I also thank them for reading my thesis and providing very useful suggestions and comments that helped me further improve the thesis.

I want to thank my co-authors with whom I did the research that is the subject of this thesis: Aris Anagnostopoulos, Haris Aziz, Luca Becchetti, Po-An Chen, David Kempe, Vangelis Markakis, and Orestis Telelis. I thank my other co-authors as well: Krzysztof Apt, Sylvain Bouveret, Irving van Heuven van Staereling, Tomas Klos, and Yingqian Zhang. It has been a pleasure doing reasearch with all of you. May we collaborate again soon.

I thank the people with whom I shared an office in the course of the past four years: Félix Carvalho Rodrigues, Po-An Chen, Riccardo Colini Baldeschi, Vincent Conitzer, Amy Greenwald, Irving van Heuven van Staereling, Gauthier van den Hove, Teun Janssen, Mona Rahn, and Orestis Telelis. I thank them for a nice time and for all the good conversations.

I will proceed to thank a bunch of colleagues now, for simply being nice colleagues, and for too many other things as well. From the Networks and Optimization research group, these are (besides some of the people that I already mentioned): Tony Huynh, Ross Kang, Anargyros Katsampekis, Teresa Piovesan, Guus Regts, Matteo Seminaroti, Sunil Simon, Tobias Mü, and Antonis Varvitsiotis. Colleagues that I met elsewhere: Aida Abiad, Simina Brânzei, Csapo Gergely, Ruben Hoeksma, Jasper de Jong, Maria Kyropoulou, Manolis Pountourakis, and Joris Scharpff. Moreover, I thank all other people that I met at conferences (and other scientific events) for the nice times.

I also want to thank a lot of good friends for their emotional support. I actually consider many of the people mentioned above to be friends. But I also need to thank many of my non-science friends: Frits van Campen, Simon Coppoolse, Rene Corbet, Phu Do, Yoram Duijn, Daan Eijkel, Bram Gollin, Lex Groot, Robin van Halem, Kasper van Heerden, Menno den Hollander, Jan Hommes, Hannah Jalink, Remco de Lange, Daan Stet, Youri van Veen, Bart Witteveen, Matthijs Wolzak, Jochum Zijlstra.

Additionally I thank many other people that I met in the past four years. (I admit: this is a clause that I include just because I am afraid that I forgot to mention someone important. Taking into account my somewhat disorganized mind, it is quite probable that I did forget someone, and I am sorry for that.)

Lastly, but certainly not least importantly, I thank my family. Especially my close family: my parents Peter and Heleen, and my two sisters Marissa and Rosanne, for their unconditional love and support.

Contents

Chapter 1

Introduction

This thesis covers some of the research in *algorithmic game theory* that I have carried out during the four years that I was a Ph.D. student. Algorithmic game theory is an area of research that lies in the intersection of economics, mathematics and computer science. The field essentially came into existence in the 1990's, as a result of the Internet developing into the most important medium for communication and information. An Internet application typically features many users that are, to some extent, self-interested. We can therefore expect such users to *strategize* over the input that they communicate to such applications. Often, the goal or task that the application needs to carry out, is not perfectly (or even not at all) aligned with the goals of the users application, and therefore such strategic behavior may be harmful to the performance of the application. Hence, certain questions become relevant in this setting, and deal for example with:

- the harm that the strategic behavior does to the performance of the application;
- the best way for a user to choose his or her strategy;
- how to design the application such that the users can not abuse it;
- whether there is any structure in the way the application and users will behave;
- which useful or desirable properties the application may or may not possess;
- how the users will cooperate, or what the best way of cooperation would be.

To answer such questions, it is natural to make use of the tools that have been developed in the field of *game theory*.

Game theory encompasses the mathematical study of strategic decision making, and models of conflict. The central objects of study in game theory are *games*, which model abstract settings in which multiple participating entities, called *players*, must choose among a set of *strategies* so as to maximize their own *utility* function. A player's utility function in turn depends on the strategies adopted by all players in the game.

One of the first basic insights that game theory has contributed is that *selfish individual behavior may be bad for the set of players as a whole*. This is demonstrated by the famous *Prisoner's Dilemma*: a canonical example of a game in game theory.

Example 1 (The Prisoner's Dilemma)**.** The *Prisoner's Dilemma* is formulated as follows. Two men are suspected of a crime and get arrested. They get separated and get offered a deal:

- If both stay silent, then both spend 1 year in prison.
- If both betray each other, then both spend 2 years in prison.
- If one betrays and the other stays silent, then the betrayee gets 3 years in prison and the betrayer 0 years.

The Prisoner's Dilemma is an example of a two-player game where both players must choose among two strategies (betrayal or staying silent). A choice of strategies where no player has a reason to change strategies is called an *equilibrium*. When we analyze the Prisoner's Dilemma, we see that the only equilibrium is when both players choose to betray. We also see that this particular choice of strategies results in the largest possible total amount of time spent in prison. The choice of strategies that results in the least total amount of time spent in prison, is when both players choose to remain silent. Despite that the latter choice of strategies is optimal for the two players together, it is unlikely that this choice of strategies will form, due to it not being an equilibrium. The Prisoner's Dilemma demonstrates that in a game, the worst possible choice of strategies may be the only choice of strategies that is "stable" with respect to the selfish behavior of the players. In the sections below, we provide a introduction and formal exposition of basic game theory (including formal definitions of games and equilibria) and algorithmic game theory.

The Internet application setting sketched above (i.e., an application with self-interested users) can be modeled as a game if we let the players be the users and if we let the strategies be the set of inputs that a user can communicate to the application. Game theory may subsequently provide us with insights into the type of questions stated above.

As soon as game theory became relevant to computer science, computer scientists started posing computationally-themed questions about the various notions and results of game theory. The addition of this computational component to game theory resulted in the field of algorithmic game theory, and has yielded various interesting results and lines of research. We will encounter many of these in this thesis and contribute new results and insights to this field as well.

1.1 This Thesis

The algorithmic game theory research discussed in the thesis is diverse, but nonetheless, the thesis has a central theme, and many parts of the thesis relate to each other in one or more ways. The central theme of the thesis is that of *externalities* in games, and the related notion of *cooperation* among the players.

By *externalities*, we mean that the players in the game take into account, in one way or the other, the well-being of the other players. This other-regarding behavior may manifest itself in a positive way, but also in a negative way: classically, (strategic) game theory assumes that players are completely selfish, and that players maximize their own utility at all costs. In this thesis, we will often be interested in the consequences of changing this assumption in various ways.

We also consider the ability to interact with other players to be a type of externality. This gives rise to the related notion of cooperation among the players. The theme of cooperation will be present in several chapters of the thesis. Indeed, the final chapters will purely deal with algorithmic cooperative game theory.

Although externalities and cooperation will frequently make an appearance, not *all* of this thesis will have to do with these topics. A notable example of this is the second chapter, which is about the inefficiency of multi-unit auctions and in which neither externalities nor cooperation play a role. The thesis has been organized such that each chapter shares one or more topics with its preceding chapter. The chapters of this thesis can therefore be classified into overlapping blocks that each share a particular sub-theme. In the final section of this introductory chapter, we elaborate on this, and provide a short abstract of each of the chapters.

Each chapter only assumes the preliminary knowledge covered in this introduction chapter, as well as a basic knowledge of combinatorial optimization, algorithms, and computational complexity. Apart from these preliminaries, each chapter is essentially self-contained. All these chapters are based on one or more (published or unpublished) papers that I co-authored. The respective paper on which a chapter is based is indicated in a footnote on the first page of that chapter.

1.2 Preliminary Remarks on Notation and Terminology

Before proceeding, it is necessary to establish some prelimininary notation.

We use the symbols $\subseteq$ and $\supseteq$ for set inclusion, and we use $\subset$ and $\supset$ for *strict* set inclusion. As usual, $\varnothing$ denotes the empty set.

The symbols $\mathbb{N}, \mathbb{Q}, \mathbb{R}$ denote the sets of natural numbers, rational numbers, and real numbers, respectively. The set of natural numbers is assumed to contain 0.

When A is a set of numbers, we may subscript it with "> 0" and "≥ 0" to denote respectively the positive and non-negative subset of A, i.e., $A_{>0} = \{a \in A \mid a > 0\}$ and $A_{\geq 0} = \{a \in A \mid a \geq 0\}$. In particular we use this notation in combination with $\mathbb{N}, \mathbb{Q}, \mathbb{R}$. E.g., $\mathbb{N}_{>0}$ denotes the set of natural numbers excluding zero, and $\mathbb{N}_{\geq 0} = \mathbb{N}$. For $a \in \mathbb{N}_{>0}$, we write $[a]$ to denote the set $\{b \in \mathbb{N}_{>0} \mid b \leq a\}$.

For a real number $a \in \mathbb{R}_{\geq 0}$, we define $\lceil a \rceil = \min\{b \in \mathbb{Z} \mid b \geq a\}$ and $\lfloor a \rfloor = \max\{b \in \mathbb{Z} \mid b \leq a\}$.

When A is a set of numbers, we write 2^A to denote its power set: $2^A = \{A' \mid A' \subseteq A\}$. We use $\times$ for the Cartesian product operation: $A \times B = \{(a, b) \mid a \in A, b \in B\}$.

Let $n \in \mathbb{N}$, let $i \in [n]$, and let $s = (s_1, \ldots, s_n)$ be an n-dimensional vector, we write (s'_i, s_{-i}) to denote the vector $(s_1, \ldots, s_{i-1}, s'_i, s_{i+1}, \ldots s_n)$ obtained from s by changing its ith coordinate to s'_i. We note that the notation (s'_i, s_{-i}) is formally ambiguous and overloads standard notation, since the index i is crucial in determining the vector that (s'_i, s_{-i}) denotes. However, i will always be clear from context, as it is always included in the subscripts used, and confusion will not arise.

Some basic probability theory will be used throughout this thesis. $\mathbf{Pr}$ will be used to refer to a probability measure, and $\mathbf{E}$ will be used for the expectation operator. Therefore, when we use the term *probability distribution*, we mean formally a *probability mass function* in case the distribution is discrete and a *probability density function* in case of a continuous distribution. When a random variable s has probability distribution σ, we write this as $s \sim \sigma$. We will oftentimes use this notation as a subscript in the symbols $\mathbf{Pr}$ and $\mathbf{E}$ in order to provide clarity about the probability space we work with.

Let σ be a probability distribution on $A_1 \times \cdots \times A_n$ (for $n \in \mathbb{N}_{>0}$ an arbitrary number, and $A_1, \ldots, A_n$ arbitrary sets). σ is said to be the *product distribution of* $(\sigma_1, \ldots, \sigma_n)$, where σ_i is a probability distribution on A_i, $i \in [n]$, iff[1] $\sigma(s) = \prod_{i=1}^{n} \sigma(s_i)$ for all $s = (s_1, \ldots, s_n) \in A_1, \times \cdots \times A_n$. When there exist distributions σ_i such that σ is the product distribution of $(\sigma_1, \ldots, \sigma_n)$, we say that σ *is a product distribution*. Moreover, for $i \in [n]$ we write σ_{-i} to denote the *projection of σ on*

[1] Throughout this thesis we use the word "*iff*" to signify a definition, as opposed to "*if*" or "*if and only if*" which will be used in the context of non-definitional statements (i.e., those that may turn out to be true or false).

$A_{-i} = A_1 \times \cdots \times A_{i-1} \times A_{i+1} \times \cdots \times A_n$, i.e., σ_{-i} is the probability distribution on A_{-i} so that $\sigma_{-i}(s_{-i}) = \sum_{s_i \in A_i} \sigma(s_1, \ldots, s_n)$ for all $s_{-i} \in A_{-i}$.

1.3 Introduction to Game Theory

The field of game theory is essentially subdivided into two subfields: *non-cooperative game theory* and *cooperative game theory*. These two disciplines are disjoint to a significant extent. In this thesis, we nonetheless study topics from both non-cooperative and cooperative game theory, and therefore provide the reader with an introduction to both disciplines (although for cooperative game theory we will be less elaborate). Unless explicit references are given, the material introduced here can be found in any introductory book on game theory. We start with non-cooperative game theory.

1.3.1 Non-Cooperative Games

The central object of study in strategic game theory is the following:

Definition 2 (Game)**.** A *(finite) (strategic* or *non-cooperative) (full information) game* is a triple $\Gamma = (n, \Sigma, u)$ where $n \in \mathbb{N}_{>0}$ and the set $[n]$ is referred to as the set of *players* of the game. The set Σ is a Cartesian product of n finite sets $\Sigma_1, \ldots, \Sigma_n$. For $i \in [n]$, the set Σ_i is referred to as the set of *strategies* of player i. The elements in Σ are called the *strategy profiles* of Γ. u is a vector of n *utility functions*. The ith component of u, for $i \in [n]$, is the function $u_i : \times_i \Sigma_i \to \mathbb{R}$ which we call the *utility function* of player i.

Example 3. The prisoner's dilemma mentioned at the start of this chapter is the game $(2, \Sigma, u)$, $\Sigma = \Sigma_1 \times \Sigma_2$, $u = (u_1, u_2)$, $\Sigma_1 = \Sigma_2 = \{b, c\}$, where b is the strategy where the player betrays the other player, and c is the strategy where the player remains silent. The utility functions are valued as the negation of how many years the player would have to serve in prison, and are thus defined by $u_1(b, c) = u_2(c, b) = 0$, $u_1(c, b) = u_2(b, c) = 3$, $u_1(b, b) = u_2(b, b) = 2$, $u_1(c, c) = u_2(c, c) = 1$.

Games are studied in the following context: Players in a game are interpreted as rational autonomous entities that each want to optimize their utility function. They each pick a strategy from their strategy set, so that a strategy profile arises. A player wants to maximize the value that its utility function maps this strategy profile to, i.e., it wants to maximize its utility.

1.3.1.1 Equilibria

This interpretation of a game gives rise to various notions of “stable” strategy profiles, also called *equilibria*. The most straightforward such notion is that of a strategy profile for which it would not be benificial for any player to unilaterally change its strategy.

Definition 4 (Pure equilibrium)**.** A *pure (Nash) equilibrium* of a game is a strategy profile $s \in \Sigma$ such that for all $i \in [n]$, and $s'_i \in \Sigma_i$,

$$u_i(s'_i, s_{-i}) \leq u_i(s). \tag{1.1}$$

The set of pure equilibria of Γ is denoted by PE_Γ.

Notions of stable strategy profiles in games, such as the notion of a pure equilibrium, are often referred to as *solution concepts*. This is because such a notion can be regarded as a rule for predicting which strategy profiles in a game will naturally arise (in a very broad sense) or “end up in”. Besides this descriptive interpretation of solution concepts, sometimes a solution concept may be interpreted in a prescriptive way as well: one could in some cases argue that the players in a game *should* play some strategy profile that conforms to a particular solution concept, as this results in a stable situation where no player wants to change its strategy profile.

Solution concepts take a central role in the study of game theory. There are various other popular solution concepts. We explain some basic game-theoretical concepts first, and use them to define the other solution concepts that we encounter in this thesis. We explain the motivation behind their definition subsequently.

Definition 5 (Mixed strategy, best response)**.** A *mixed strategy* of player $i \in [n]$ is a probability distribution on Σ_i. A *best response* of player $i \in [n]$ to a probability distribution σ on Σ is a strategy $s'_i \in \Sigma_i$ that maximizes $\mathbf{E}_{s_{-i}\sim\sigma_{-i}}[u_i(s'_i, s_{-i})]$.

Definition 6 (Nash equilibrium, correlated equilibrium, coarse equilibrium)**.**

- A *coarse (correlated) equilibrium* of a game is a probability distribution σ on Σ with the following property: For each player $i \in [n]$, and all $s'_i \in \Sigma_i$:

$$\mathbf{E}_{s\sim\sigma}[u_i(s)] \geq \mathbf{E}_{s_{-i}\sim\sigma_{-i}}[u_i(s'_i, s_{-i})]. \tag{1.2}$$

- A *(mixed) Nash equilibrium* of a game is a coarse equilibrium σ that is a product distribution of $(\sigma_1, \dots, \sigma_n)$ for some probability distributions σ_i on Σ_i, $i \in [n]$. We will sometimes abuse this definition and refer to the vector $(\sigma_1, \dots, \sigma_n)$ as a mixed Nash equilibrium.

- A *correlated equilibrium* is a probability distribution σ on Σ with the following property: For each player $i \in [n]$ and strategies $s_i^*, s_i' \in \Sigma_i$ with $\mathbf{Pr}_{s\sim\sigma}[s_i = s_i^*] > 0$:

$$\mathbf{E}_{s\sim\sigma}[u_i(s) \mid s_i = s_i^*] \geq \mathbf{E}_{s\sim\sigma}[u_i(s_i', s_{-i}) \mid s_i = s_i^*]. \tag{1.3}$$

For a game Γ, we refer to its sets of coarse, mixed Nash, and correlated equilibria as respectively CsE_Γ, NE_Γ, and ClE_Γ.

When we view the pure equilibria as probability distributions with singleton support, it holds that:

$$CsE_\Gamma \supseteq ClE_\Gamma \supseteq NE_\Gamma \supseteq PE_\Gamma.$$

The first of these containments holds because (1.3) is a more restrictive condition than (1.2). The last of these three containments holds because a pure equilibrium (viewed as a distribution with singleton support) is a product distribution. For the middle containment we have to give a proof. We do this along with another basic fact:

Proposition 7. *The mixed strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a Nash equilibrium if and only if for all $i \in [n]$, the support of σ_i is a subset of the strategies in Σ_i that are best responses to σ. Moreover, $ClE_\Gamma \supseteq NE_\Gamma$.*

Proof. Let us write down the expected utility of player i.

$$\begin{aligned}\mathbf{E}_{s\sim\sigma}[u_i(s)] &= \sum_{s\in\Sigma} u_i(s) \prod_{i\in[n]} \sigma_i(s_i) \\ &= \sum_{s_i\in\Sigma_i} \sigma_i(s_i)\mathbf{E}_{s_{-i}\sim\sigma_{-i}}[u_i(s_i, s_{-i})].\end{aligned}$$

This shows that the expected utility player i receives for playing σ_i against the opposing mixed strategies σ_{-i} is a convex combination of the terms $\mathbf{E}_{s_{-i}\sim\sigma_{-i}}[u_i(s_i, s_{-i})]$, $s_i \in \Sigma_i$. If i plays with non-zero probability a strategy s_i' that is not a best response to σ_{-i}, then it can improve its expected utility by playing a best response with probability 1. This proves the first claim.

The second claim follows because the expectation on the left hand side of (1.2) is equal to $\mathbf{E}_{s\sim\sigma}[u_i(s)] \mid s_i = s_i^*]$ for all $s_i^* \in \Sigma_i$ such that $\mathbf{Pr}_{s\sim\sigma}[s_i = s_i^*] > 0$ when σ is a Nash equilibrium, from what we just established. The expression on the right hand side of (1.2) is equal to $\mathbf{E}_{s\sim\sigma}[u_i(s_i', s_{-i})] = \mathbf{E}_{s\sim\sigma}[u_i(s_i', s_{-i})|s_i = s_i^*]$ for $s_i^* \in \Sigma_i, \mathbf{Pr}_{s\sim\sigma}[s_i = s_i^*] > 0$ when σ is a Nash equilibrium. I.e., (1.2) and (1.3) are equivalent in case σ is a Nash equilibrium. □

1.3.1.2 Nash Equilibria

The motivation behind Nash equilibria is apparent when one allows the players in a game Γ to play a strategy according to an arbitrary probability distribution. One can view this setting as a new "game" $\Gamma' = (n, \Sigma', u')$, where $\Sigma' = \Sigma_1 \times \cdots \times \Sigma_n$, and Σ_i is the set of mixed strategies of player i.[2] u'_i is then defined as $u'_i(\sigma) = \mathbf{E}_{s\sim\sigma}[u_i(s)]$, and the pure equilibria of Γ' are the Nash equilibria of Γ. As we conclude from the above proposition, in a mixed equilibrium every player only has best responses in its support, so the pure equilibrium condition (1.1) for Γ' is equivalent to the Nash equilibrium condition (1.2) for Γ. The Nash equilibria are thus the stable strategies when we allow the players to choose their strategies at random.

A fundamental result on Nash equilibria is that they are guaranteed to exist (if the strategy sets are finite and the player set is finite as well, but this is inherent in the definition of a game that we have given here).

Theorem 8 (Existence of Nash equilibria [Nash, 1950]). *For all games* Γ*,* $NE_\Gamma \neq \varnothing$.

As a consequence of 1.3.1.1, ClE_Γ and CsE_Γ are also guaranteed to be non-empty.

1.3.1.3 Correlated Equilibria

Correlated equilibria can be interpreted as there being an "advisor", telling the players what would be the best thing to do. In real life, there are settings where the presence of such an advisor is natural. Think for example of a road junction with traffic lights. The traffic lights here tell the drivers what to do. The drivers (approximately) adhere to the advice of the traffic lights, so one could argue that the strategy distribution of the drivers that the traffic lights induce, is a correlated equilibrium (in a loose sense).

Correlated equilibria are interesting because they tell us which stable strategy profiles might form when we have the ability to advise the players. These equilibria may be much nicer than regular Nash equilibria, as the following example shows.

Example 9 (Game of Chicken). The *game of chicken* is a game with 2 players. Both players have strategy set $[2]$. If both players play strategy 1, the utility of both is 0. If one player plays strategy 1 and the other plays strategy 2, then the player that plays strategy 2 gets utility 2, and the other gets utility 7. If both players play 2, then they both get utility 6.

---

[2] We write "game" between quotation marks because we define here a game with strategy sets of infinite cardinality, which fall outside Definition 2. We do however occasionally study games with infinite strategy sets; as we discuss later on in this chapter.

There are two pure equilibria: $(1, 2)$ and $(2, 1)$. There is one mixed equilibrium, where both players play strategy 1 with probability $1/3$. However, it is straightforward to verify that the distribution σ, defined by

$$\sigma(1, 2) = 1/3, \sigma(2, 1) = 1/3, \sigma(2, 2) = 1/3, \sigma(1, 1) = 0,$$

is a correlated equilibrium. This correlated equilibrium gives the players a higher average expected utility than any of the three Nash equilibria.

1.3.1.4 Coarse Equilibria

The *coarse equilibrium* is a solution concept that can be motivated in the context of *no-regret sequences*.

Definition 10 (Regret, no-regret sequence)**.** Suppose players have to play a single utility maximization game Γ repeatedly a number of times, say T times. Let the strategy profiles of these T plays be $s^1, \ldots s^T$. The *regret* $r_\Gamma(i, (s^1, \ldots, s^T))$ of a player $i \in [n]$ for the sequence $(s^1, \ldots, s^T)$ is the maximum amount of additional utility that i could have received if it would have played a fixed strategy each of the T plays:

$$r_\Gamma(i, (s^1, \ldots, s^T)) = \max \left\{ 0, \sum_{t \in [T]} (u_i(s'_i, s^t_{-i}) - u_i(s^t)) \;\middle|\; s'_i \in \Sigma_i \right\}.$$

If Γ is instead a game where each player $i \in [n]$ wants to minimize a cost function c_i, rather than to maximize a utility function,[3] then the regret of $(s^1, \ldots s^T)$ is defined as follows.

$$r_\Gamma(i, (s^1, \ldots, s^T)) = \max \left\{ 0, \sum_{t \in [T]} (c_i(s^t) - c_i(s'_i, s^t_{-i})) \;\middle|\; s'_i \in \Sigma_i \right\}.$$

The sequence $(s^1, \ldots s^T)$ is *no-regret* for a player $i \in [n]$ if $r_\Gamma(i, (s^1, \ldots, s^T)) = 0$. This means that in hindsight, player i would not have wanted to play any fixed strategy in Σ_i all T times.

Given $s_1, \ldots, s_T$, let σ be a distribution on Σ such that $\sigma(s) = |\{i \mid s^i = s\}|/T$. Then we see that $s_1, \ldots, s_T$ is no-regret for each player if and only if σ is a coarse equilibrium.

[3] Such games are called *cost minimization games*; see Section 1.3.1.7 below.

Definition 11 (Vanishing regret)**.** Let $(s^1, s^2, \ldots)$ be an infinite sequence of strategy profiles in Σ. This sequence has *vanishing regret* iff for all $i \in [n]$,

$$\lim_{T \to \infty} r_\Gamma(i, (s^1, \ldots, s^T)) = 0.$$

Coarse equilibria and some theorems about vanishing regret sequences appear in Chapters 4 and 5.

1.3.1.5 Games with Infinite Strategy Sets

Sometimes in this thesis, we study games with infinite strategy sets. Formally, Definition 2 is merely generalized by removing the requirement that $\Sigma_i, i \in [n]$ is finite. One should note that for these games, Theorem 8 is false. Games with infinite strategy sets will be studied in Chapters 2, 5, and 6.

1.3.1.6 Incomplete Information Games

In many economic situations, players have no common and full knowledge of the game that is being played. A player might for example have no complete knowledge of the preferences and knowledge of the other players of the game. These situations can not satisfactorily be modeled by means of a game as defined above. Instead, we model these as *incomplete information games*. Incomplete information games are relevant to Chapter 2.

Definition 12 (Incomplete information game)**.** An *incomplete information game*, alternatively called a *Bayesian game*, is a quintuple $\Gamma = (n, \Sigma, u, V, \pi)$ where n and Σ are defined in the same way as for a full information game. The set V is the Cartesian product of the sets $V_1, \ldots, V_n$. The set V_i is called the set of *types* of player $i \in [n]$, and is assumed to be finite. The set V is called the set of *type profiles*. Vector u is the vector of utility functions $(u_1, \ldots, u_n)$ of the players. The utility functions are defined in the same way as for full information games, but the domain of u_i is now $\Sigma \times V_i$. The object π is called the *type distribution* and is a probability distribution on V. We assume that π is a product distribution, and therefore (using mild notational abuse) we sometimes identify π with the vector $(\pi_1, \ldots, \pi_n)$ of which π is the product distribution.

Extensions of this definition are encountered in the literature: Sometimes, the domain of u_i is $\Sigma \times V$. For this reason the type of incomplete information game we just defined, is actually usually called a *private value* complete information game. We will not encounter non-private value incomplete information games in this thesis. A second generalization (that will also not be dealt with in this thesis) is when π may be arbitrary (i.e., is not required to be a product distribution). We speak in that case of *correlated types*.

The interpretation of an incomplete information game is that the players do not have complete knowledge over the types of the other players: player i only knows the distribution π and the valuation v_i that is drawn for it from π_i. They will therefore want to optimize their expected utility, based on only this knowledge: Each player $i \in [n]$ wants to choose for each type $v_i \in V_i$ that it may get assigned, a strategy $s(v_i) \in \Sigma_i$. Strategy s can thus be regarded as a function from V_i to Σ_i. We call such a function a *Bayesian strategy*.

Definition 13 (Pure and mixed Bayesian strategies)**.** A *(pure) Bayesian strategy* is a function from V_i to Σ_i. We denote by $\Sigma_i^{V_i}$ the set of (pure) Bayesian strategies of player $i \in [n]$. A *mixed Bayesian strategy* is a function from V_i to the set $\Delta(\Sigma_i)$ of probability distributions on Σ_i. We denote by $\Delta(\Sigma_i)^{V_i}$ the set of mixed Bayesian strategies of player $i \in [n]$.

The vectors in the set Σ^V, which we define as $\times_{i\in[n]}\Sigma_i^{V_i}$, are called *(pure) Bayesian strategy profiles*. The vectors in the set $\Delta(\Sigma)^V$, which we define as $\times_{i\in[n]}\Delta(\Sigma_i)^{V_i}$, are called *mixed Bayesian strategy profiles*.

When $v \in V$ and $s \in \Sigma^V$ we use the notation $s(v)$ for $(s_1(v_1), \ldots, s_n(v_n))$. Moreover, when $s'_i \in \Sigma_i^{V_i}$, we write $(s'_i(v_i), s_{-i}(v_{-i}))$ for $(s_1(v), \ldots, s_{i-1}(v_{i-1})\ , s'_i(v_i), s_{i+1}(v_{i+1}), \ldots, s_n(v_n))$.

Each player choosing a Bayesian strategy gives rise to a Bayesian strategy profile $s \in \Sigma^V$. A player $i \in [n]$ thus wants to choose in this profile s its own Bayesian strategy s_i such that for all $v_i \in V_i$

$$\mathbf{E}_{v_{-i}\sim\pi_{-i}}[u_i(s(v), v_i)]$$

is optimized.

A natural solution concept for an incomplete information game is therefore the following.

Definition 14 (pure and mixed Bayes-Nash equilibrium)**.** A *pure Bayes-Nash equilibrium* is a Bayesian strategy profile $s \in \Sigma^V$ such that for every player $i \in [n]$, every Bayesian strategy $s'_i \in \Sigma_i^{V_i}$, and every type $v_i \in V_i$ it holds that

$$\mathbf{E}_{v_{-i}\sim\pi_{-i}}[u_i(s(v), v_i)] \geq \mathbf{E}_{v_{-i}\sim\pi_{-i}}[u_i((s'_i(v_i), s_{-i}(v_{-i})), v_i)].$$

A *mixed Bayes-Nash equilibrium* is a mixed Bayesian strategy profile σ in $\Delta(\Sigma)^V$ such that for every player $i \in [n]$, every Bayesian strategy $s'_i \in \Sigma_i^{V_i}$, and every type $v_i \in V_i$ it holds that

$$\mathbf{E}_{\substack{v_{-i}\sim\pi_{-i}\\ s\sim\sigma}}[u_i(s(v), v_i)] \geq \mathbf{E}_{\substack{v_{-i}\sim\pi_{-i}\\ s\sim\sigma}}[u_i((s'_i(v_i), s_{-i}(v_{-i})), v_i)].$$

$PBNE_\Gamma$ and $MBNE_\Gamma$ refer respectively to the set of all pure Bayes-Nash equilibria and all mixed Bayes-Nash equilibria of an incomplete information game Γ.

Pure and mixed Bayes-Nash equilibria for incomplete information games are the natural analogues of the pure and mixed Nash equilibria that we defined for full information games.

Just like for full information games, the existence of a mixed Bayes-Nash equilibrium is guaranteed when the strategy sets are finite.

Theorem 15. *Let Γ be an incomplete information game where $|\Sigma|$ is finite. Then $MBNE_\Gamma \neq \varnothing$.*

The above theorem is a corollary of e.g. the main result of Milgrom and Weber [1985].

1.3.1.7 Cost Minimization Games

We discuss in Chapters 3 to 5 games where players are assumed to want to minimize their function u instead of maximizing it. We refer to such games as *cost minimization games*, as opposed to *utility maximization games*. In this case, the functions u are referred to as *cost* functions instead of *utility functions*. Moreover, in the setting of cost minimization games, our convention will be to rename the functions u to $c = (c_1, \ldots, c_n)$. When discussing a cost minimization game $\Gamma = (n, \Sigma, c)$, the pure, Nash, correlated, and coarse equilibria of Γ are defined as the pure, Nash, correlated, and coarse equilibria of the utility maximization game (n, Σ, u), where $u = -c$. We will only study cost minimization games in the context of full information, and not in the context of incomplete information.

1.3.1.8 The Price of Anarchy and Price of Stability

It is well known that equilibria are often suboptimal for the set of players as a whole. The need to gain an accurate understanding of the extent of suboptimality caused by selfish behavior has led to the study of the *inefficiency of equilibria* in algorithmic game theory. In this context, a common inefficiency measure is the *price of anarchy*, which relates the worst equilibrium of a game to the optimal strategy profile.

In order to make this formal, we need to study a game $\Gamma = (n, \Sigma, u)$ in combination with a *social welfare function* $U_\Gamma : \Sigma \to \mathbb{R}$ of which the role is to quantify the quality of the strategy profiles in Σ.

Definition 16 (Social welfare function)**.** A *social welfare function* for a full information utility maximization game $\Gamma = (n, \Sigma, u)$ is a function $U_\Gamma : \Sigma \to \mathbb{R}$. When Γ is an incomplete information utility maximization game, a social welfare function for Γ is a function $U_\Gamma : \Sigma \times V \to \mathbb{R}$.

In case of cost minimization games, by convention we rename U to C, and C is referred to as a *social cost function* rather than a social welfare function.

The subscript Γ in U_Γ and C_Γ will henceforth be omitted, since it will always be clear from the context.

The choice of U is usually a quite straightforward and sensible one, such as:

$$U(s) = \sum_{i \in [n]} u_i(s), \tag{1.4}$$

for full information games, or

$$U(s, v) = \sum_{i \in [n]} u_i(s, v_i),$$

for incomplete information games. Likewise, for full information cost minimization games, the above becomes

$$C(s) = \sum_{i \in [n]} c_i(s).$$

Another popular one (for full information games) is

$$U(s) = \min\{u_i(s) \mid i \in [n]\},$$

which measures the utility of the player that is least well off, under a given strategy profile. In the context of cost minimization games, the latter social welfare function becomes:

$$C(s) = \max\{c_i(s) \mid i \in [n]\}. \tag{1.5}$$

In Chapters 2 to 6, we will be concerned with the quality of equilibria with respect to the optimal value of a social welfare function. The following notions are therefore central to those chapters.

Definition 17 (price of anarchy and price of stability for full information games)**.**

- The *price of anarchy* of a full information utility maximization game Γ for a set S of distributions on the set Σ of strategy profiles of Γ, with respect to social welfare function U, is defined by

$$PoA(\Gamma, S) = \max\left\{ \frac{U(s^*)}{\mathbf{E}_{s \sim \sigma}[U(s)]} \;\middle|\; \sigma \in S, s^* \in \Sigma \right\}.$$

- The *price of anarchy* of a full information cost minimization game Γ for a set S of distributions on the set Σ of strategy profiles of Γ, with respect to social cost function C, is defined by

$$PoA(\Gamma, S) = \max\left\{ \frac{\mathbf{E}_{s\sim\sigma}[C(s)]}{C(s^*)} \;\middle|\; \sigma \in S, s^* \in \Sigma \right\}.$$

- The *price of stability* of a full information utility maximization game Γ for a set S of distributions on the set Σ of strategy profiles of Γ, with respect to social welfare function U, is defined by

$$PoS(\Gamma, S) = \min\left\{ \frac{\max\{U(s^*) \mid s^* \in \Sigma\}}{\mathbf{E}_{s\sim\sigma}[U(s)]} \;\middle|\; \sigma \in S \right\}.$$

- The *price of stability* of a full information cost minimization game Γ for a set S of distributions on the set Σ of strategy profiles of Γ, with respect to social cost function C, is defined by

$$PoS(\Gamma, S) = \min\left\{ \frac{\mathbf{E}_{s\sim\sigma}[C(s)]}{\min\{C(s^*) \mid s^* \in \Sigma\}} \;\middle|\; \sigma \in S \right\}.$$

We obtain the definitions of *coarse price of anarchy/stability*, *correlated price of anarchy/stability*, *mixed price of anarchy/stability*, and *pure price of anarchy/stability* (all with respect to U) when in the above definitions we take for S respectively the sets CsE_Γ, ClE_Γ, NE_Γ, and PE_Γ (when we regard the pure equilibria in PE_Γ as probability distributions with singleton support).

We define the coarse, correlated, mixed, and pure price of anarchy and stability for a *class* of games $\mathcal{G}$ as, respectively, the supremum of the coarse, correlated, mixed, and pure price of anarchy and stability of the games in $\mathcal{G}$.

The price of anarchy was defined first in Koutsoupias and Papadimitriou [1999, 2009], and the price of stability was first defined in Anshelevich et al. [2004].

The price of anarchy notion has spawned a large amount of research in algorithmic game theory. Some of the price of anarchy literature will be discussed in Chapters 2 to 6.

For incomplete information games, we need to define the price of anarchy slightly differently.

Definition 18 (Price of anarchy for incomplete information games)**.** The *mixed Bayes-Nash price of anarchy* of an incomplete information game Γ, with respect to social welfare function U is given by:

$$PoA(\Gamma, MBNE_\Gamma) = \max \left\{ \frac{\mathbf{E}_{v\sim\pi}[U(s^v, v)]}{\mathbf{E}_{\substack{v\sim\pi \\ s\sim\sigma}}[U(\sigma, v)]} \;\middle|\; \sigma \in MBNE_\Gamma \right\},$$

where $s^v \in \Sigma$ maximizes $U(\cdot, v)$, for $v \in V$.

The definition of the *pure Bayes-Nash price of anarchy* is obtained from the above definition by replacing $MBNE_\Gamma$ by $PBNE_\Gamma$ (when again regarding the elements of $PBNE_\Gamma$ as functions in $\Delta(\Sigma)^V$ that map to distributions with singleton support).

We do not define the price of stability for incomplete information games, as we do not study the price of stability in an incomplete information setting in this thesis.

1.3.1.9 Smoothness

Smoothness is a technique introduced by Roughgarden [2009] in order to prove upper bounds on the price of anarchy of classes of games. It plays an important role in Chapters 2, 4 and 5.

Definition 19 ((λ, μ)-smooth game)**.** Let Γ be a (full information) utility maximization game and let U be a social welfare function for Γ. Game Γ is said to be (λ, μ)-*smooth* with respect to U iff for every pair $s^*, s \in \Sigma$:

$$\sum_{i\in[n]} u_i(s_i^*, s_{-i}) \geq \lambda U(s^*) - \mu U(s). \tag{1.6}$$

If Γ is instead a (full information) cost minimization game and if C is a social cost function for Γ, then Γ is said to be (λ, μ)-*smooth* with respect to C iff for every pair $s^*, s \in \Sigma$:

$$\sum_{i\in[n]} c_i(s_i^*, s_{-i}) \leq \lambda C(s^*) + \mu C(s). \tag{1.7}$$

A class of games is said to be (λ, μ)-smooth iff every game in that class is (λ, μ)-smooth.

The concept of (λ, μ)-smoothness derives its importance from the following fact. When a game or class of games is shown to be (λ, μ)-smooth (and μ satisfies a technical condition), then under a mild assumption on the social welfare or social cost function, a bound on the coarse price of anarchy is immediate.

Definition 20 (Sum-bounded social welfare/cost function)**.** A social welfare function U is *sum-bounded* with respect to a game $\Gamma = (n, \Sigma, u)$ iff $U(s) \geq \sum_{i\in[n]} u_i(s)$ for all $s \in \Sigma$. A social cost function C is sum-bounded with respect to a game $\Gamma = (n, \Sigma, c)$ iff $C(s) \leq \sum_{i\in[n]} c_i(s)$ for all $s \in \Sigma$.

Theorem 21. *Let $\lambda, \mu \in \mathbb{R}$. If a utility maximization game is (λ, μ)-smooth with respect to a non-negative sum-bounded social welfare function U and $\mu > -1$, then the coarse price of anarchy of that game is at most $(1 + \mu)/\lambda$. If a cost minimization game Γ with non-negative sum-bounded social cost function C is (λ, μ)-smooth and $\mu < 1$, then the coarse price of anarchy of that game is at most $\lambda/(1 - \mu)$.*

Proof. Let Γ be a (λ, μ)-smooth utility maximization game, let U be a sum-bounded social welfare function for Γ, let $\sigma \in CsE_\Gamma$ and let $s^* \in \Sigma$ be strategy profile that maximizes U. The inequalities in the following derivation follow from respectively sum-boundedness of U, the coarse equilibrium condition (1.2), the smoothness condition (1.6), and linearity of expectation.

$$\begin{aligned}
\mathbf{E}_{s\sim\sigma}[U(s)] &\geq \mathbf{E}_{s\sim\sigma}\left[\sum_{i\in[n]} u_i(s)\right] \\
&\geq \mathbf{E}_{s\sim\sigma}\left[\sum_{i\in[n]} u_i(s_i^*, s_{-i})\right] \\
&\geq \mathbf{E}_{s\sim\sigma}[\lambda U(s^*) - \mu U(s)] \\
&\geq \lambda U(s^*) - \mu \mathbf{E}_{s\sim\sigma}[U(s)].
\end{aligned}$$

Rearranging terms, we see that this is equivalent to $(1 + \mu)\mathbf{E}_{s\sim\sigma}[U(s)]/U(s^*) \geq \lambda$. When we divide by $(1 + \mu)$ and assume that $\mu > -1$ (so that $(1 + \mu)$ is positive), we see that the latter is in turn equivalent to $U(s^*)/\mathbf{E}_{s\sim\sigma}[U(s)] \leq (1 + \mu)/\lambda$. This establishes that the coarse price of anarchy is at most $(1 + \mu)/\lambda$.

For cost minimization games, the proof is entirely analogous. □

We note that the sum-boundedness condition is actually not required: the above is merely the theorem and proof as given by Roughgarden [2009]. In Chapter 5, we show how to get rid of the sum-boundedness requirement, even under a strictly less restrictive and more general smoothness condition.

Roughgarden [2009] defined smoothness as a consequence of the observation that many proofs of upper bounds on the price of anarchy for various classes of games that were given by that time, follow the same pattern: in essence, they come down to proving that a game is (λ, μ)-smooth for some particular $(\lambda, \mu) \in \mathbb{R}^2$. In many of those cases, the upper bound was proved only for the pure or mixed price of anarchy. The smoothness notion made it clear that such existing upper bounds generalize automatically to upper bounds on the coarse price of anarchy.

After its introduction, the smoothness concept has been extended and adapted in various ways. For example, in Lucier and Paes-Leme [2011], the notion of *semi-smoothness* is used, and Roughgarden and Schoppmann [2011] define a smoothness variation called *local smoothness*. Both these notions are used on certain classes of games to establish bounds on the price of anarchy that cannot be proved through regular smoothness. Nadav and Roughgarden [2010] showed that smoothness bounds apply all the way to solution concepts called "average coarse correlated equilibrium," but not beyond.

The basic idea behind the smoothness technique is to bound the sum of cost increases of individual players switching strategies by a combination of the costs of two states. Because these types of bounds capture local improvement dynamics, they bound the price of anarchy not only for Nash equilibria, but also more general solution concepts, including coarse correlated equilibria. Informally, the local smoothness notion of Roughgarden and Schoppmann [2011] require the types of bounds described above only for nearby states, thus obtaining tighter bounds on the price of anarchy, albeit only for more restrictive solution concepts and convex strategy sets. Using the local smoothness framework, Roughgarden and Schoppmann [2011] obtain optimal upper bounds for atomic splittable congestion games (i.e., a variation on congestion games where the players have to split their impact among multiple sets of facilities).

Smoothness has also been extended for use with incomplete information games by Syrgkanis [2012], and independently by Roughgarden [2012]. Furthermore, by Syrgkanis and Tardos [2013] a variation of smoothness is proposed that is specifically tailored to auction settings with incomplete information. This variation of smoothness will be introduced, used and further discussed in Chapter 2.

In Chapters 4 and 5, we will propose two generalizations of smoothness as well. These generalizations are suited for analyzing games with *altruistic players*, and players for which their behavior is influenced by an underlying *social context*. These concepts will be defined and further explained in the respective chapters.

1.3.1.10 Congestion Games

Congestion games form an important class of games that occur in multiple chapters of the thesis: They are studied in Chapters 4 and 5, and variations of congestion games are studied in Chapters 3 and 7. The definition of a congestion game is as follows:

Definition 22 (Congestion Game)**.** A *congestion game* is a cost minimization game

(n, Σ, c) for which there exists an $m \in \mathbb{N}_{>0}$ such that $\Sigma_i \subseteq 2^{[m]}$. Moreover, for each $j \in [m]$ there exists a function $d_j : [n] \to \mathbb{R}_{\geq 0}$, and for each player $i \in [n]$, its cost function c_i is given by $c_i(s) = \sum_{j \in s_i} d_j(|\{i' \in [n] \mid j \in s_{i'}\}|)$ for $s \in \Sigma$. In the context of a congestion game, the elements of $[m]$ are called *facilities*, and d_j is referred to as the *delay function* of facility $j \in [m]$.

A congestion game may thus be represented as a triple (n, m, Σ, d) where $d = (d_1, \ldots, d_m)$ is the vector of delay functions.

There are some important special subclasses of congestion games:

Definition 23 (Linear, symmetric, and singleton congestion games, fair cost sharing games)**.**

- A congestion game is said to be *linear* iff for each facility $j \in [m]$ there are numbers $a_j, b_j \in \mathbb{Q}_{\geq 0}$ such that the delay function d_j of facility j satisfies $d_j(x) = a_j x + b_j$ for all $x \in [n]$.
- A congestion game is said to be *symmetric* iff $\Sigma_i = \Sigma_j$ for all $i, j \in [n]$.
- A congestion game is said to be *singleton* iff for all $i \in [n]$, Σ_i only contains sets of cardinality 1.
- A congestion game is said to be a *fair cost sharing game* if for each facility $j \in [n]$ there is a *cost* $a_j \in \mathbb{R}_{\geq 0}$ and its delay function d_j is given by $d_j(x) = a_j/x$ for all $x \in [n]$.

Congestion games were given their name through the following observation: If the function d is increasing (unlike the special case of fair cost sharing games), then the delay of a facility increases as more player include the facility in their strategy. The facility hence gets *congested*, and the players choosing this facility will experience an increase in their cost.

Congestion games can be used to model a variety of situations, of which the most natural ones are that of *traffic congestion* or *network congestion*: Consider a traffic network in which multiple players must choose a route between two points in the network. Model the edges of this network as the facilities of a congestion game, and let the set of strategies of a player be the set of all routes from its source to its destination. Finally, model the delay functions such that a facility gets more congested as more players choose it. Using this, one may study the phenomenon of traffic congestion as a result of selfish behavior of the drivers, by means of a congestion game. More general network congestion phenomena may be studied in a similar fashion.

Another scenario is one where the players have an interest in *building* infrastructure, which can be modeled as a cost sharing game. In a fair cost sharing game, individual players have an interest in building infrastructure (i.e., the facilities), and share the cost of building a facility equally among the players that choose to build it.

Congestion games were introduced in Rosenthal [1973]. Moreover, in the same paper, congestion games were shown to always possess pure equilibria.

Theorem 24 (Rosenthal [1973])**.** *Every congestion game has a pure equilibrium.*

Monderer and Shapley [1996] prove that congestion games coincide with the class of *exact potential games*.

Definition 25. *Potential games* are cost minimization games $\Gamma = (n, \Sigma, c)$ for which there exists a function $\Phi : \Sigma \to \mathbb{R}$ such that $\Phi(s'_i, s_{-i}) \leq \Phi(s)$ if $c_i(s'_i, s_{-i}) \leq c_i(s)$ for all $s, s' \in \Sigma, i \in [n]$. Γ is an *exact potential game* iff there exists a function $\Phi : \Sigma \to \mathbb{R}$ such that $\Phi(s) - \Phi(s'_i, s_{-i}) = u_i(s) - u_i(s'_i, s_{-i})$ for all $s, s' \in \Sigma, i \in [n]$.

Since Φ cannot decrease indefinitely, every potential game trivially has a pure equilibrium. Of course, for utility maximization games, one can come up with an analogous definition of potential games, but we will not encounter such games in this thesis.

Congestion games have received a great deal of attention. They were generalized by Milchtaich [1996] to congestion games with player-specific delay functions, of which the equilibria were further studied by Ackermann et al. [2006]. In the area of algorithmic game theory, the price of anarchy of congestion games was studied in Awerbuch et al. [2005], Christodoulou and Koutsoupias [2005a,b], Aland et al. [2006], Fotakis [2007], Bilò et al. [2011b], Fotakis et al. [2009]. For two related classes of games, called *splittable congestion games* and *non-atomic (Wardrop) congestion games*, the price of anarchy was studied by Roughgarden and Schoppmann [2011] and Roughgarden and Tardos [2002], respectively. The problems of finding social-welfare optimizing and fair allocations were studied in Chakrabarty et al. [2005], Blumrosen and Dobzinski [2006], Meyers and Schulz [2012].

For the special case of fair cost sharing games, the price of stability was studied in Anshelevich et al. [2004, 2008].

With respect to cooperation and externalities, congestion games and related classes of games have been a popular target of study as well, see e.g., Hayrapetyan et al. [2006], Babaioff et al. [2007, 2009], Roth [2008], Bilò et al. [2011a,b], Babaioff et al. [2007], Hoefer and Skopalik [2009b], Caragiannis et al. [2010], Chen and Kempe [2008].

The literature just mentioned will be discussed in more detail in the chapters of this thesis that deal with congestion games. We note that the complete body of literature that deals with congestion games is too vast to discuss exhaustively, and the papers mentioned above form only a tiny fraction it.

1.3.1.11 Mechanism Design

Mechanism design is a subfield of game theory that is concerned with designing games such that a set of desirable criteria are met. Some mechanism design problems will be studied in Chapters 6, 7, and 8. This section deliberately describes mechanism design from a high level. Formal details can be found in the three chapters just mentioned.

In a mechanism design problem in its broadest sense, there is a set of *outcomes*, and a game needs to be specified for any set of players, where the players express preferences over these outcomes. Thus, this game is not a regular game in the sense of Definition 2. Instead, the game is a mapping from the set of strategy profiles (which is part of what needs to be designed) to the set of outcomes, and the players in turn receive a utility based on the outcome.

In classical mechanism design problems, we usually assume that the preferences of the players are private information. I.e., we need to design the game (including the strategy sets of the players) in such a way that the outcome does not depend on the preferences of the players, but only on the given strategy profile. Such a mechanism design problem will be considered in Chapter 7.

Chapters 6 and 8, however, discuss problems that can be considered to fall within mechanism design in the broader sense. In Chapter 6, we study mechanism design *in a restricted domain*, which means that we are not completely free in the design of our game, and instead need to pick a game from a given class of games. Moreover, we do not adopt a model of private preferences. In Chapter 8, we study mechanism design in the context of housing markets. The study of the latter such markets borders on the areas of social choice theory and borrows some concepts from cooperative game theory (see the section below). As a consequence, in mechanism design problems for this setting we consider a different set of objectives than usual: first, it does not involve money (i.e., charging payments to the players), and secondly we want our mechanism to satisfy different properties, such as the *core selection* property, which essentially means that our mechanism should be resistant to sets of players *cooperating* and reallocating their houses among each other.

1.3.2 Cooperative Games

This section is relevant to the final two chapters of the thesis, which are about *cooperative games*. There is a weak relation of this section to Chapter 8 as well.

A cooperative game is an abstract object that one defines in order to study in isolation the cooperative aspects of any setting in which multiple parties, players, or entities participate. Cooperative games are fundamentally different from the strategic games that we discuss in the sections above. The type of questions one is interested in, when dealing with cooperative games, are usually about how to divide among the players of a coalition P the utility that P can collectively generate.

Definition 26 (Cooperative game, characteristic function)**.** A *cooperative game* is a pair $\Gamma = (n, v)$ where $n \in \mathbb{N}_{\geq 1}$, and $v : 2^{[n]} \to \mathbb{R}$. The *players* of Γ are the elements in $[n]$. The function v is referred to as the *characteristic function*, and $v(P), P \subseteq [n]$ is referred to as the *value* of P: it reflects the value (e.g. the amount of money) that coalition P can generate when the players in P cooperate with each other. A subset of $[n]$ is referred to as a *coalition* in the context of cooperative games. When the number of players n is clear from context or arbitrary, we will sometimes refer to v as a cooperative game, instead of (n, v).

Two important properties that a cooperative game may possess are *non-negativity*, *monotonicity* and *superadditivity*.

Definition 27 (Cooperative game, characteristic function)**.** A cooperative game (n, v) is

- *non-negative* iff $v(S) \geq 0$ for all $P \subseteq [n]$.
- *monotone* iff $v(P_1) \leq v(P_2)$ for all $P_1 \subseteq P_2 \subseteq [n]$.
- *superadditive* iff $v(P_2 \cup P_2) \geq v(P_1) + v(P_2)$ for all $P_1, P_2 \subseteq [n], P_1 \cap P_2 = \varnothing$.

Most important classes of cooperative games are monotone and non-negative, and indeed the games studied in the final two chapters of this thesis are all non-negative and monotone.

A central question in cooperative game theory is how to divide among the players of a given cooperative game (n, v) the maximum value that the players can generate. In case v is superadditive, this maximum value is $v([n])$; attained when the players all cooperate. In case v is not superadditive, this maximum value is reached by having the players cooperate according to a partition $\mathcal{P}$ of $[n]$ that maximizes $\sum_{P \in \mathcal{P}} v(P)$. In the context of cooperative games, $\mathcal{P}$ is referred to as a *coalition structure*.

Definition 28 (Coalition structure)**.** Let $\Gamma = (n, v)$ be a cooperative game. A partition of $[n]$ is a *coalition structure* of Γ.

Suppose that (n, v) is a superadditive game. Let $x \in \mathbb{R}^n_{\geq 0}$ be a way of dividing $v([n])$ among the players: For $i \in [n]$, the value x_i is player i's share, and $\sum_{i \in [n]} x_i = v([n])$. A first condition that any reasonable x needs to satisfy is that a player does not generate more value than x_i on its own, because otherwise player i is not incentivized to work together with the other players. This motivates the notion of an *imputation*.

Definition 29 (Imputation)**.** An imputation of a cooperative game (n, v) is a vector $x \in \mathbb{R}^n_{\geq 0}$ such that $\sum_{i \in [n]} x_i = v([n])$ and $x_i \geq v(\{i\})$ for all $i \in [n]$.

Within the context of cooperative game theory, a set of vectors $x \in \mathbb{R}^n_{\geq 0}$ that specify for each player how much it receives, is referred to as a *solution concept*. Imputations form a very weak solution concept, as it imposes only very minimal requirements on what a good way of allocating value to the players should look like: we do not want to allocate more than what the players can collectively generate (i.e., $v[n]$), and we do not want individual players to refrain from cooperating. There have been proposed many solution concepts for cooperative games, and indeed, most of them form a subset of the set of imputations.

By far the most popular solution concept is *the core*. This solution concept is more restrictive than the set of all imputations, since it requires that *coalitions* of players must be incentivized to work together with the other players, in addition to just individual players.

Definition 30 (Core)**.** The *core* of a cooperative game $\Gamma = (n, v)$ is the set of imputations $x \in \mathbb{R}^n_{\geq 0}$ of Γ for which it holds that $\sum_{i \in P} x_i \geq v(P)$ for all $P \subseteq [n]$.

A necessary condition for the core of a cooperative game Γ to be non-empty, is that Γ is superadditive. However, as mentioned above, if Γ is not superadditive, then it is not a reasonable assumption that players will cooperate as the single coalition $[n]$, so $v([n])$ is not a reasonable value to divide among the players. Instead, there is a coalition structure that the players will form, which generates more total value, and we should instead focus on dividing the values of these coalitions among the players in the coalition. The problem of finding the optimal coalition structure is the central topic of Chapter 9.

A final solution concept that we mention is the *Shapley value*. Again, this is a solution concept suitable for superadditive games, since it divides $v([n])$ among the players. Unlike the core and the broad set of imputations, the Shapley value is only a single vector, instead of a set.

Definition 31 (Shapley value)**.** Let (n, v) be a cooperative game. The *Shapley value of* (n, v) is an n-dimensional vector $\varphi(n, v) = (\varphi_1(n, v), \ldots, \varphi_n(n, v))$, where for $i \in [n]$

$$\varphi_i(n, v) = \kappa_i(n, v)/n!, \tag{1.8}$$

and

$$\kappa_i(n, v) = \sum_{P \subseteq [n] \setminus \{i\}} (|P|!)(n - |P| - 1)!(v(P \cup \{i\}) - v(P)). \tag{1.9}$$

For $i \in [n]$ the value $\varphi_i(n, v)$ is called the *Shapley value of player* i and the value $\kappa_i(n, v)$ is called the *raw Shapley value of player* i. It is well-known and straightforward to derive that the raw Shapley value can be written as

$$\kappa_i(n, v) = \sum_{\sigma \in S_n} (v(p(i, \sigma) \cup \{i\}) - v(p(i, \sigma))),$$

where S_n denotes the set of permutations on $[n]$ and for $i \in [n], \sigma \in S_n$, $p(i, \sigma)$ denotes the set $\{i' \mid \sigma^{-1}(i') < \sigma^{-1}(i)\}$ of players that appear before i in σ.

The Shapley value is a solution concept that is of central importance in cooperative game theory. It has been shown that it is the sole solution concept that satisfies the following four properties simultaneously [Winter, 2002]:

- *Efficiency*: $\sum_{i \in [n]} \varphi_i(n, v) = v(n)$;
- *Symmetry*: if $i, i' \in [n]$ are symmetric, then $\varphi_i(n, v) = \varphi_{i'}(n, v)$; (Players $i, i' \in [n]$ are called *symmetric* in (n, v) if $v(P \cup \{i\}) = v(P \cup \{i'\})$ for any coalition $P \subseteq [m] \setminus \{i, i'\}$.)
- *Dummy*: if i is a dummy, then $\varphi_i(n, v) = 0$; (A player $i \in [n]$ is a *dummy* if $v(P \cup \{i\}) - v(P) = 0$ for all $P \subseteq [n]$.)
- *Additivity*: For two cooperative games (n, v_1) and (n, v_2) it holds that $\varphi_i(n, v^1 + v^2) = \varphi_i(n, v_1) + \varphi_i(n, v_2)$ for all $i \in [n]$;[4]
- *Anonymity*: permuting the players does not affect their Shapley value.

Various other such *axiomatic characterizations* of the Shapley value exist. We refer the reader to any introductory text on cooperative game theory.

Chapter 10 of this thesis deals with computing the Shapley value in a well-known class of games called *matching games*.

1.4 Summary of Notational Conventions

The symbols introduduced above will be used consistently throughout this thesis. We summarize these here. For convenience, a list of symbols can be found at the end of the thesis.

[4] The sum of two characteristic functions v_1 and v_2 on the same player set is defined in the standard way: as $v_1(P) + v_2(P)$ for all $P \subseteq [n]$.

The standard symbol that we use to denote a game is Γ. The standard symbol that we use to denote a (hyper)graph is G. All games and graphs that we study or discuss throughout this thesis are assumed to have *finite* player sets and vertex sets, respectively.

In the context of games, n will be used to denote the number of players in the game, and $[n]$ will usually be identified with the player set of the game. Symbols Σ and Σ^V will denote the set of strategy profiles and Bayesian strategy profiles, respectively (and we use Σ_i and $\Sigma_i^{V_i}$ to refer to the strategy set and Bayesian strategy set of player $i \in [n]$, respectively). We use u to denote the vector of utility functions of a game. The symbol i will be used to denote a player in $[n]$, and P will be used to refer to a subset of $[n]$. Symbol s will be used to refer to strategy profiles in Σ (for full information games) or Bayesian strategy profiles in Σ^V (for incomplete information games). Symbol σ will be used to refer to probability distributions on Σ (for full information games) or elements in $\Delta(\Sigma)^V$ (for incomplete information games). We use p_i to denote the type distribution of player $i \in [n]$, and we use π to refer to the vector $(\pi_1, \ldots, \pi_n)$ or to the product distribution of $(\pi_1, \ldots, \pi_n)$, interchangeably. Symbol U will be used to refer to a social welfare function. In case we discuss a cost minimization game, the convention will be to name the cost functions c instead of u and C instead of U. Symbol CsE_Γ, ClE_Γ, NE_Γ, and PE_Γ will be used to refer to the sets of coarse, correlated, Nash, and pure equilibria of game Γ, respectively. MBNE_Γ and PBNE_Γ will be used to refer to the set of mixed and pure Bayes-Nash equilibria of incomplete information game Γ, respectively.

We may sometimes discuss multiple games simultaneously. In cases where confusion can arise, we subscript all relevant symbols with the appropriate game (e.g., in such a situation, $u_{i,\Gamma}$ denotes the utility function of player i in game Γ).

In many games that we will discuss throughout this thesis, there will be a notion of *facility* of *machine* present in the definition of these games, such as in the congestion games we introduced above. We will denote in these settings the number of facilities or machines by m, and the set of facilities or machines will be identified with $[m]$. We use j to denote particular facilities or machines. The symbol S will be used to refer to a subset of $[m]$.

In the context of a graph or hypergraph $G = (V, E)$, n will denote the number of vertices in the graph, and the vertex set V is usually identified with $[n]$. Symbol E is the standard symbol that we use for denoting the edge set of the graph, and m will normally denote the number of edges of the graph, i.e., $m = |E|$. We normally use v when we want to refer to a particular vertex in $[n]$ and e to refer to a particular edge in E. When the vertices of a graph correspond to the players of a game (which will sometimes be the case), we may also use i to refer to a particular vertex of the graph.

1.5 Outline of the Thesis

The subsequent chapters cover a large part of the research in algorithmic game theory that I and various collaborators carried out while I was a Ph.D. student. The topics studied are somewhat diverse, but a large part of them has in common that it deals with algorithmic problems in games in which there are externalities present among the players, or the players may behave cooperatively in some way. Multiple sub-themes can be identified throughout the thesis, and each of the chapters deals with one or more of these. The chapters have been grouped together according to these themes in the following way:

- Chapters 2 to 6 concern the study of the *price of anarchy*.
- Chapters 4 to 7 deal with various aspects of *externalities*.
- Chapters 6 to 8 cover various *mechanism design* problems.
- In Chapters 7 to 10, the *algorithmic* theme is most prominently present.
- In Chapters 9 and 10, various algorithmic problems in *cooperative game theory* are studied.

In more detail, the contents of the chapters are as follows.

Chapter 2. The first chapter studies the price of anarchy of some classes of auction games that fall under the family of *multi-unit auctions*. These are auctions in which there are multiple copies of a single good, and the auction allocates these to the players (possibly allocating multiple copies to the same player). A player has a valuation for each number of copies of the good that it may get allocated. A player therefore submits a bid for each number of copies that they may recieve. Items are allocated to the players corresponding to the highest marginal bids. In the *discriminatory* variant of the auction, the player pays a price equal to the bid corresponding to the number of bids it receives. In the *uniform price* variant of the auction, the player pays a price equal to the highest marginal bid for which no item is allocated. For both variants, we derive upper and lower bounds on the price of anarchy. We do this for two different bidding interfaces, two different classes of valuation functions, under both a full information setting and an incomplete information setting. Some of the bounds obtained are the first known bounds, while the others are significant improvements over the previously best known bounds.

Chapter 3. Here, we are concerned with the price of anarchy when deviations of *sets* of players can take place, rather than deviations of individual players. A strategy profile that for which no set of player can deviate such that each player in the set improves its utility, is called a *strong equilibrium*. We study the *strong* price of anarchy, i.e., price of anarchy for strong equilibria, for a variation of congestion games called *bottleneck congestion games*, in which strong equilibria are guaranteed to exist. In bottleneck congestion games, the cost function of a player is defined as the maximum delay among the facilities that it chooses. The social cost function we consider is the one that returns the maximum cost among the players. We restrict our studies to linear latency functions, and derive various upper and asymptotically matching lower bounds on the strong price of anarchy.

Chapter 4. This chapter is concerned with the price of anarchy as well. This time, we consider games with (partially) *altruistic players*. for various classes of games when players are (partially) altruistic. We model altruistic behavior by associating a number $\alpha_i \in [0, 1]$ to each player i, and assuming that player i's perceived cost is a convex combination of $1 - \alpha_i$ times his direct cost and α_i times the social cost. Within this framework, we study altruistic extensions of various well-known game classes: linear congestion games, fair cost-sharing games and valid utility games. We derive (tight) bounds on the price of anarchy of these games for several solution concepts. Thereto, we adapt the *smoothness* notion and show that often this variation on smoothness captures the essential properties to determine the price of anarchy of these games.

Chapter 5. We extend here the ideas introduced in the previous chapter, and consider a more fine-grained model of social player behavior: for each pair of players (i, j), we introduce a parameter $\alpha_{i,j}$ that reflects the attitude of player i towards player j. A positive value of $\alpha_{i,j}$ reflects a friendly attitude, while a negative value reflects a spiteful attitude. The perceived utility of a player i is then a linear combination of the direct utilities of all the players, where the values $\alpha_{i,j}, j \in [n]$ are the coefficients. We refer to the matrix α as a *social context*, and our interest is in the price of anarchy for classes of games in which the players are embedded into such a social context. We extend the smoothness notion further, in order to make it suitable studying the price of anarchy of games with social contexts. We apply this smoothness framework to three classes of games: linear congestion games, minsum scheduling games, and generalized second price auction. For these classes of games, we establish upper bounds on the price of anarchy for a broad class of altruistic social contexts. For linear congestion games and minsum scheduling games, the price of anarchy remains bounded by a constant.

Chapter 6. We study the phenomenon of *spite* in procurement auctions (i.e., auctions where the auctioneer buys an item from the bidders). As a result of spite, bidders that do not win the auction themselves might derive negative utility from the auction when an enemy or competitor wins, altering the players' bidding strategies. We consider generalizations of first-price and second-price procurement auctions with spiteful players under a full information model. The class of auctions studied is more general than the standard first-price and second-price procurement auction due to the fact that for each bidder a penalty multiplier μ may be set by the auction designer, such that the bid considered in the auction is the bid of the bidder multiplied by μ. For two bidders, we characterize all ϵ-equilibria (i.e., bidding profiles such that no bidder can change strategy and improve by more than ϵ) for $\epsilon > 0$, and quantify the extent to which the spite levels of the bidders impact the quality of the equilibria, by deriving the price of anarchy and price of stability relative to the spiteless optimum. Moreover, we find closed form expressions for setting the penalty multipliers so as to minimize the price of anarchy and price of stability of the auction, for a setting where no assumptions can be made on the valuations of the two bidders, but the spite levels are known. For $n > 2$ bidders, we characterize the set of ϵ-equilibria when $\epsilon \to 0$, give a polynomial time algorithm to compute them all, and derive the price of anarchy and price of stability.

Chapter 7. In this chapter we depart from the price of anarchy theme that has been dominant in the previous chapters, but retains the theme of externalities among the players. We consider a variant of congestion games where every player i expresses for each resource e and player j a positive *externality*, i.e., a value for being on e together with player j. We adopt an optimization point of view and consider the problem of optimizing the social welfare, i.e., sum of all players' utilities.

We show that this problem is NP-hard even for very special cases, notably also for the case where the players' utility functions for each resource are affine (contrasting with the tractable case of linear functions Blumrosen and Dobzinski [2006]). We derive a 2-approximation algorithm by rounding an optimal solution of a natural LP formulation of the problem. Our rounding procedure is sophisticated because it needs to take care of the dependencies between the players resulting from the pairwise externalities. We also show that this is essentially best possible by showing that the integrality gap of the LP is close to 2.

Small adaptations of our rounding approach enable us to derive approximation algorithms for several generalizations of the problem. Most notably, we obtain an $(r+1)$-approximation when every player may express for each resource externalities on player sets of size r. Further, we derive a 2-approximation when the strategy sets of the players are restricted and a $\frac{3}{2}$-approximation when these sets are of size 2.

Finally, we consider the associated mechanism design problem where the players may misreport the values of their externalities and the mechanism may charge payments to the bidders. We show that in this setting, the social welfare can be $(r+1)$-approximated in expectation, within expected polynomial time, while satisfying some desirable mechanism-design-theoretic properties: *truthfulness*, *individual rationality*, and *non-negative payments*. On the negative side, this mechanism requires the extension of the solution set where the mechanism is allowed to "disable" a set of externalties, preventing the players from obtaining the utility associated with these externalities.

Chapter 8. We study in this chapter another mechanism design problem. The problem we study appears in a social setting, in *housing markets*. The (Shapley-Scarf) housing market is a well-studied and fundamental model of an exchange economy. Each agent owns a single house and the goal is to reallocate the houses to the agents in a mutually beneficial and stable manner. The focus of this chapter of the thesis, is on the case where agents express indifferences among houses. In Alcalde-Unzu and Molis [2011] and Jaramillo and Manjunath [2011], Shapley-Scarf housing markets with such indifferences were independently examined. These papers proposed two important families of mechanisms, known as TTAS and TCR respectively. We formulate in this chapter a family of mechanisms which not only includes TTAS and TCR but also satisfies many desirable properties of both families. As a corollary, we show that TCR satisfies the desirable property that it always outputs an allocation that is in the *strict core*, in case such an allocation exists. Finally, we settle an open question regarding the computational complexity of the TTAS mechanism.

Chapter 9. We move in this setting to an abstract cooperative problem and consider the *coalition structure generation problem*, in which the goal is to partition the players into exhaustive and disjoint coalitions so as to maximize the social welfare. The optimization problem considered in Chapter 7 is a special case of a coalition structure generation problem, and in this chapter we study it from a broader perspective. One of our results is a polynomial time algorithm to solve the problem for all coalitional games provided that *player types* are known and the number of player types is bounded by a constant. As a corollary, we obtain a polynomial-time algorithm to compute an optimal partition for weighted voting games with a constant number of weight values, for *coalitional skill games* with a constant number of *skills*, and for *linear games* with a constant number of *desirability classes*. We also consider the coalition structure generation problem for various coalitional games defined compactly on combinatorial domains. For these games, we characterize the complexity of computing an optimal coalition structure by presenting polynomial time algorithms, approximation algorithms, or NP-hardness and inapproximability lower bounds.

Chapter 10. This chapter studies in more depth one of the classes of cooperative games introduced in Chapter 9, concerns the analysis of the *Shapley value*. Matching games form a fundamental class of cooperative games that help understand and model auctions and assignments. In a matching game, the players correspond to vertices of a graph, and the value of a coalition of vertices is the weight of the maximum size matching in the subgraph induced by the coalition. After establishing some general insights, we show that the Shapley value of matching games can be computed in polynomial time for some special cases: graphs with maximum degree two, and graphs that have a small modular decomposition into cliques or cocliques (complete k-partite graphs are a notable special case of this). The latter result extends to various other well-known classes of graph-based cooperative games. We continue by showing that computing the Shapley value of unweighted matching games is #P-complete in general. Finally, a fully polynomial-time randomized approximation scheme (FPRAS) is presented. This FPRAS can be considered the best positive result conceivable, in view of the #P-completeness result.

Chapter 2

On the Inefficiency of Standard Multi-Unit Auctions*

In this chapter, we analyze the price of anarchy for a family of incomplete-information utility maximization games: *multi-unit auctions*. Among the material presented in Chapter 1, Section 1.3.1 up to Section 1.3.1.9 is important to the present chapter. Particularly relevant are Sections 1.3.1.6 and the material about incomplete information games in Section 1.3.1.8.

The solution concept we will focus on is the mixed Bayes-Nash equilibrium.

To define properly a multi-unit auction, we first need the notion of a *submodular vector*.

Definition 32 (Submodular vector, submodular function, subadditive function)**.** A *submodular vector* $v \in \mathbb{R}^m_{\geq 0}$ is a vector such that $v_{i+1} - v_i \leq v_i - v_{i-1}$ and $v_{i+1} - v_i \geq 0$ for all $i \in [m-1]$, where we define $v_0 = 0$.

When f is a function from $\{0\} \cup [m]$ to $\mathbb{R}_{\geq 0}$, we define it to be:

- *submodular* iff $(f(0), f(1), \ldots, f(m))$ is a submodular vector.
- *subadditive* iff $f(a+b) \leq f(a) + f(b)$ for all $a, b \in \mathbb{N}_{\geq 0}$ such that $a + b \leq [m]$.

We note that the class of submodular functions is strictly contained in the class of subadditive ones [Lehmann et al., 2006].

In this chapter, the symbols we use for vectors are often subscripted (for example, s_i will be used to denote a vector of more than one element). For convenience, when

*The contents of this chapter have been published as De Keijzer et al. [2013a,b].

we want to denote a particular element of such a vector, we will write its index between suffixed parentheses (for example, $s_i(1)$ is the first element of s_i).

Definition 33 (Multi-unit auction)**.** A multi-unit auction is a game where there are $m \in \mathbb{N}_{>0}$ indivisible units of a single good that are to be allocated to the players $[n]$. Each player i submits a bid $\hat{s}_i(j)$ for each $j \in [m]$, reflecting the amount of money that i is willing to pay for getting j items. A player does this by specifying a *marginal bid vector*: a vector s_i of m *marginal bids* (i.e., numbers in $\mathbb{R}$), such that $\sum_{k=1}^{j} s_i(j) = \hat{s}_i(j)$ for $j \in [m]$. The marginal bid vector must satisfy a set of requirements, called a *bidding interface*, to be specified later. The subset of $\mathbb{R}^m$ of marginal bid vectors that satisfies the requirements of the bidding interface thus serves as player i's strategy set Σ_i for all $i \in [n]$. The set of n-dimensional vectors of bid vectors that satisfy the requirements of the bidding interface thus form the set Σ of strategy profiles of the game, and the strategy profiles will in this context be referred to as *bid profiles*. We often refer to the marginal bid vectors of a bid profile as simply *bid vectors*, and we often refer to the marginal bids of a bid vector as simply *bids*.

Based on the bid profile, the items are allocated to the players according to an allocation function $x : \Sigma \to \Delta_{m,n}$, where $\Delta_{m,n} = \{a \in \mathbb{N}^n_{\geq 0} \mid \sum_{i=1}^{n} a_i = m\}$. $x_i(s)$ indicates how many items player i receives on bid profile s. This function is defined as follows: $x_i(s)$ items are given to a player iff exactly $x_i(s)$ elements of s_i are among the m highest marginal bids of s, i.e., exactly $x_i(s)$ elements of s_i are among the first m elements of the non-increasingly ordered vector of all nm marginal bids of s. Ties are broken according to an any fixed tie-breaking rule.

Moreover, each player $i \in [n]$ has to pay a certain price $p_i(s)$, where $p : \Sigma \to \mathbb{R}^n_{\geq 0}$. Associated to a player $i \in [n]$ is a subadditive *valuation function* $v_i : [m] \cup \{0\} \to \mathbb{R}_{\geq 0}$, $v_i(0) = 0$, which expresses for each possible number of items that i can get allocated, how much i values getting that number of items. The utility of a player is then given by $u_i(s, v_i) = v_i(x_i(s)) - p_i(s)$.

In the full information setting, the valuation function is fixed, so that the utility function boils down to a function from Σ to $\mathbb{R}$. In the incomplete information setting, the valuation functions take the role of the types, and are drawn from a given type distribution π. Note that, in accordance to Definition 12, we assume that π has finite support, and that π is the product distribution of vector $(\pi_1, \ldots, \pi_n)$ of type distributions, where $\pi_i, i \in [n]$ is on a given finite set V_i of valuation functions for player i. The set $V = \times_{i \in [n]} V_i$ therefore takes the role of the set of type profiles and will also be called the set of *valuation profiles* in this setting.

We study two standard multi-unit auction formats:

Definition 34 (Discriminatory auction, uniform price auction)**.**

- A *discriminatory auction* is a multi-unit auction where $p_i(s) = \sum_{j \in [x_i(s)]} s_i(j)$
- A *uniform price auction* is a multi-unit auction where $p_i(s) = x_i(s)\ell(s)$, where $\ell(s)$ is defined as the $(m+1)$th element in the non-increasingly ordered vector of all mn marginal bids. I.e., $\ell(s)$ is the "highest losing marginal bid".

Note that for the special case of a single item, the discriminatory auction is equal to a *first price auction*, where the highest bidding player gets the item, and pays its bid. Likewise, the uniform price auction is for a single item equal to a *second price auction*, where the highest player gets the item, and pays the second-highest bid. Discriminatory and uniform price auctions can thus be seen as a natural generalization of first price and second price auctions to multiple items.

Both auction formats are popular in practice. Among the applications of these auctions are [Ausubel and Cramton, 2002]: the allocation of state bonds to investors, spectrum auctions, the Eurosystem, to online sales over the internet, facilitated by popular online brokers.

We consider two bidding interfaces:

Definition 35 (Standard bidding interface, uniform bid vector, uniform bidding interface)**.**

A multi-unit auction is said to have the *standard bidding interface* iff Σ consists of all vectors of nondecreasing and non-negative marginal bid vectors. I.e., the vectors s such that for all $i \in [n]$ the vector $\hat{s}_i$ (defined as $\hat{s}_i(j) = \sum_{k \in [j]} s_i(k)$ for $j \in [m]$) is submodular, non-negative and non-decreasing.

A *uniform bid vector* is a non-decreasing and non-negative marginal bid vector (i.e., conform the standard bidding interface) that additionally satisfies the requirement that there are numbers $a \in [k] \cup \{0\}$ and $b \in \mathbb{R}_{\geq 0}$ such that the marginal bid $s_i(j) = b$ when $j \leq a$ and $s_i(j) = 0$ when $j > a$. A multi-unit auction is said to have the *uniform bidding interface* iff Σ consists only of the vectors of uniform bid vectors.

The standard bidding interface is most prevalent in the scientific literature, and the uniform bidding interface is more popular in practice. We evaluate the economic inefficiency of both multi-unit auction formats for both bidding interfaces, by means of upper and lower bounds on the price of anarchy for pure equilibria and mixed Bayes-Nash equilibria.

The usual definition of a multi-unit auction is more restrictive than the one we give here: it is usually required that the valuation functions of the players are submodular, instead of subadditive. The bounds on the price of anarchy that were obtained prior to this work [Markakis and Telelis, 2012, Syrgkanis and Tardos, 2013] indeed assume submodular valuation functions.

We consider here the more general subadditive valuation functions, and we study submodular valuation functions separately as a special case. Our results signify that these auctions are nearly efficient, which provides further justification for their use in practice.

Our developments improve the bounds that have been obtained in Markakis and Telelis [2012] and Syrgkanis and Tardos [2013] for submodular valuation functions.

2.1 Background

Multi-unit auctions are one of the most widespread and popular tools for selling identical units of a good with a single auction process. In practice, they have been in use for a long time, one of their most prominent applications being the auctions offered by the U.S. and U.K. Treasuries for selling bonds to investors, see e.g., of Treasury [1998]. In more recent years, they are also implemented by various online brokers: eBid [2013], Ockenfels et al. [2006]. In the literature, multi-unit auctions have been a subject of study ever since the seminal work of Vickrey [1961] (although the need for such a market enabler was conceived even earlier, in Friedman [1960]) and the success of these formats has led to a resurgence of interest in auction design.

There are three simple standard multi-unit auction formats that have prevailed and are being implemented; these are the two auction formats defined above (i.e., the discriminatory and uniform price auctions) and the *Vickrey multi-unit auction*, which charges prices according to an instance of the *Clarke payment rule*, which is a standard payment rule in the mechanism design literature. All three formats have seen extensive study in auction theory [Krishna, 2002, Milgrom, 2004].

Except for the Vickrey auction, which is truthful (i.e., incentivizes players to bid equal to their valuation) and efficient (i.e., attains maximum social welfare when players behave rationally), the others suffer from a *demand reduction* effect [Ausubel and Cramton, 2002], whereby players may have incentives to submit a bid vector that is less than their valuation, in order to receive less units at a better price. This effect is amplified when players have non-submodular valuation functions, since the bidding interface forces them to encode their value within a submodular bid vector. Even worse, in many practical occasions players need to submit bids according to the uniform bidding interface. In such a setting, each player is required to "compress" its valuation function into a bid that scales linearly with the number of units. The mentioned allocation and pricing rules apply also in this *uniform bidding* setting, thus yielding different versions of discriminatory and price auctions. Despite the volume of research from the economics community [Ausubel and Cramton, 2002, Noussair, 1995, Engelbrecht-Wiggans and Kahn, 1998, Binmore and Swierzbinski, 2000, Reny, 1999, Bresky, 2008] and the widespread popularity of these auction formats, the first attempts of quantifying their economic efficiency are [Markakis and Telelis, 2012, Syrgkanis and Tardos, 2013]. There has also been no study of these auction formats for non-submodular valuations, as noted by Milgrom [2004].

The multi-unit auction formats that we examine here present technical and conceptual resemblance to the *simultaneous auctions* format that has received attention in [Feldman et al., 2013, Christodoulou et al., 2008, Bhawalkar and Roughgarden, 2011, Hassidim et al., 2011, Syrgkanis and Tardos, 2013]. However, upper bounds in this setting do not carry over to our format. Simultaneous auctions were first studied by Christodoulou et al. [2008]. The authors proposed that each of a collection of distinct goods, with one unit available for each of them, is sold in a distinct *second price auction*, simultaneously and independently of the other goods. Bidders in this setting may have combinatorial valuation functions over the subsets of goods, but they are forced to bid separately for each good. For players with fractionally subadditive valuation functions, they proved a tight upper bound of 2 on the mixed Bayes-Nash price of anarchy of the Simultaneous Second Price Auction. Bhawalkar and Roughgarden [2011] extended the study of inefficiency for players with subadditive valuations functions and showed an upper bound of $O(\log(m))$ which was reduced to 4 by Feldman et al. [2013]. For arbitrary valuation functions, Fu et al. [2012] proved an upper bound of 2 on the inefficiency of pure Nash equilibria, when they exist.

Hassidim et al. [2011] studied *simultaneous first price auctions*. They showed that pure equilibria in this format are always efficient, when they exist. They proved constant upper bounds on the inefficiency of Nash equilibria for (fractionally) subadditive valuation functions and $O(\log(m))$ and $O(m)$ for the inefficiency of mixed Bayes-Nash equilibria for subadditive and arbitrary valuation functions. Syrgkanis [2012] showed that this format has mixed Bayes-Nash price of anarchy $e/(e-1)$ for fractionally subadditive valuation functions. Feldman et al. [2013] proved an upper bound of 2 for subadditive ones.

Syrgkanis and Tardos [2013] and Roughgarden [2012] independently developed extensions of the smoothness technique for games of incomplete information. In Syrgkanis and Tardos [2013], these ideas are further developed for analyzing the inefficiency of simultaneous and sequential compositions of simple auction mechanisms. They demonstrate extensive applications of their techniques on welfare analysis of standard multi-unit auction formats and their compositions. For submodular valuation functions, they prove inefficiency upper bounds of $2e/(e-1)$ and $4e/(e-1)$ for the discriminatory and uniform price auction, respectively. Here, we improve upon these results; our improvements carry over to simultaneous and sequential compositions as well.

2.2 Contributions and Outline

The main results we give in this chapter are bounds on the pure price of anarchy and mixed Bayes-Nash price of anarchy of multi-unit auctions. For the incomplete information setting, our bounds are improvements over the previously best known bounds for the case that the players have submodular valuation functions, and the first known bounds on the price of anarchy for the case that the players have subadditive valuation functions. The price of anarchy studied here is with respect to the *sum of valuations* as a choice of social welfare function. This is a standard choice of social welfare function in the case of auction games, and all the price of anarchy bounds in the literature we mentioned above are with respect to this social welfare function. Note that this is equal to the "standard" social welfare function (1.4) when we consider the auctioneer as a player of the game, for which its utility is the total amount of money that the other players have to pay.

The pure price of anarchy, in the full information setting, is studied in Section 2.4. For uniform price auctions, we show that the *undominated* pure price of anarchy (explained in the next section) is at least $(1 - 1/e - 2/m)^{-1}$ when the valuation functions of the players are submodular and the uniform bidding interface is used. (For the case of the standard bidding interface, an almost matching upper bound of $(1 - 1/e)^{-1}$ was shown to hold in Markakis and Telelis [2012].) In the same section, we additionally point out for uniform price auctions some of the consequences that our bounds on the mixed Bayes-Nash price of anarchy (proved later in the chapter) have for the pure price of anarchy. For the discriminatory auction, we show that the existence of pure equilibria is strongly dependent on the tie-breaking rule that the auction employs. We prove that there is always a tie breaking rule in which pure equilibria are guaranteed to exist (although this rule needs to be tailored to the valuation functions of the players), and when they exist, we prove that they always attain the optimal social welfare. I.e., the pure price of anarchy of discriminatory auctions is 1, even for arbitrary valuation functions under the uniform bidding interface.

We focus in Section 2.5 on the efficiency of mixed Bayes-Nash equilibria. Our results for the mixed Bayes-Nash price of anarchy are summarized in Table 2.1. For uniform-price auctions, our bounds assume a restriction on the bid profiles that may form, known as *no-overbidding*. This restriction is explained in the next section.

For submodular valuation functions, we derive upper bounds of $e/(e-1)$ and $3.1462 = |W_{-1}(-1/e^2)| < 2e/(e-1)$ for the discriminatory and the uniform price auctions, respectively. These improve upon the previously best known bounds of $2e/(e-1)$ and $4e/(e-1)$ [Syrgkanis and Tardos, 2013]. For the uniform price auction, our bound is less than a factor 2 away from the known lower bound of $e/(e-1)$ [Markakis and Telelis, 2012]. We also prove lower bounds of $e/(e-1)$ and 2 for

Valuation Functions	**Auction Format** (`bidding: standard | uniform`) *discriminatory auction*	*uniform price auction*
Submodular	$\frac{e}{e-1}$	$\lvert W_{-1}(-1/e^2)\rvert \approx 3.1462$
Subadditive	$2 \,\Big\vert\, \frac{2e}{e-1}$	$4 \,\Big\vert\, 2\lvert W_{-1}(-1/e^2)\rvert \approx 6.2924$

Table 2.1: Upper bounds on the Bayes-Nash price of anarchy of multi-unit auctions. The upper bounds for uniform price auctions are under the no-overbidding assumption (see Section 2.3). W_{-1} denotes the lower branch of the Lambert W function.

the discriminatory price auction and uniform price auction, with respect to the known proof techniques of Syrgkanis and Tardos [2013], Feldman et al. [2013], Christodoulou et al. [2008], Bhawalkar and Roughgarden [2011], Hassidim et al. [2011]. As a consequence, unless the upper bound of $e/(e-1)$ for the discriminatory auction is tight, its improvement requires the development of novel tools; the same holds for reducing the uniform price auction upper bound below 2 (if $e/(e-1)$ from Markakis and Telelis [2012] is indeed the worst-case). For subadditive valuations, we obtain bounds of $2e/(e-1)$ and $6.2924 < 4e/(e-1)$ for discriminatory and uniform price auctions respectively, independent of which of the two bidding interfaces is used.

Further, for the standard bidding interface we derive improved bounds of 2 and 4, respectively, by adapting a technique from Feldman et al. [2013]. We also give a lower bound of almost 2 for uniform pricing and subadditive valuations. In Section 2.6 we discuss further applications of our results in connection with the smoothness framework of Syrgkanis and Tardos [2013]. In particular, some of these bounds carry over to simultaneous and sequential compositions of such auctions (see Table 2.2 in Section 2.6).

The outline of this chapter is as follows. After introducing some preliminary knowledge in Section 2.3, we present in Sections 2.4 and 2.5 our bounds results for the full information setting and incomplete information setting, respectively. Finally in Section 2.6, the connections between our result and the *smoothness*-framework of Syrgkanis and Tardos [2013] is explored. Some concluding remarks are given in Section 2.7.

2.3 Definitions and Preliminaries

Proposition 36. *For any non-negative non-decreasing function $f : \{0\} \cup [m] \to \mathbb{R}_{\geq 0}$ and any integers $x, y \in [m], x < y$, the following hold:*

- *If f is submodular, then $f(x)/x \geq f(y)/y$.*
- *If f is subadditive, then $f(x)/x \geq f(y)/(x+y)$.*

Proof. We prove $f(a)/a \geq f(a+1)/(a+1)$ for all $a \in [m]$. Note that the value $f(a)/a$ is the average increase of f on the interval $[0, a]$. $f(a+1)$ is the sum of $f(a)$ and $f(a+1) - f(a)$. The latter, $f(a+1) - f(a)$, is the increase on $[a, a+1]$. This is at most $f(a) - f(a-1)$ by submodularity, and the latter is in turn at most $f(a)/a$, also by submodularity.

It remains to prove the second claim. Note that subadditivity implies that for all $a \in \mathbb{Q}_{\geq 0}$ and $b \in \mathbb{N}_{\geq 0}$ such that $ab \in \mathbb{N}$, $f(ab) \leq f(\lceil a \rceil b) \leq \lceil a \rceil f(b)$. We write y as $(1+a)x$, for $a \in \mathbb{Q}_{\geq 0}$. We see that $f(y)/(x+y) = f((1+a)x)/(2+a)x \leq \lceil 1+a \rceil f(x)/(2+a)x \leq f(x)/x$, because $2 + a \geq \lceil 1 + a \rceil$. □

Marginal valuations. Just like for bid vectors, we can specify a valuation function v_i by its marginal vector of valuations. We will denote this vector by $\check{v}_i$, where $\check{v}_i(j) = v_i(j) - v_i(j-1)$ for $j \in [m]$. If v_i is submodular it holds that $\check{v}_i$ is non-increasing.

Winning bids. For a vector of vectors of (marginal) bids s, and $i \in [m]$ we write $\beta_i(s)$ to denote the mth element of the non-increasingly ordered vector of all bids of s. We refer to $\beta(s) = (\beta_1(s), \ldots, \beta_m(s))$ as the vector of *winning bids of* s, as the number of winning bids that a player submits, is the number of items it wins.

Distributions on vectors of bid vectors. When σ is a mixed Bayesian strategy and $v \in V$ and $i \in [n]$, we write $\bar{\sigma}(v)$ to refer to the distribution of $(\sigma_1(v_1), \ldots, \sigma_n(v_n))$ on $\times_{i \in [n]} \Sigma_i$. In general, when we use the symbol σ, we will write a bar above σ to indicate that $\bar{\sigma}$ is a distribution on bid profiles, rather than a mixed Bayesian strategy profile.

Lambert W. The Lambert W function defined as: For all $x \in \mathbb{R}$, $W(x)$ is the number satisfying the solution $x = W(x)e^{W(x)}$. This function is two-valued on a certain interval, and hence we discriminate between the lower branch W_{-1} and the upper branch W_0 of the Lambert W function.

Social welfare. The social welfare function we use for studying the price of anarchy is the following: In case of incomplete information, for a bid profile $s \in \Sigma$ and a valuation profile $v \in V$:

$$U(s, v) = \sum_{i \in [n]} v_i(x_i(s)).$$

For $v \in V$, we will denote the bid profile maximizing $U(\cdot, v)$ by s^v (i.e., as in Definition 18).

In case of complete information, the vector v is fixed, and the social welfare function is:

$$U(s) = \sum_{i \in [n]} v_i(x_i(s)).$$

As we mentioned earlier in the background discussion above: this is a standard social welfare function in the study of auction games. This function is equal to the "standard" social welfare function (1.4) when we consider the auctioneer as a player of the game, for which its utility is the total amount of money that the other players have to pay.

No-overbidding. For the case of the uniform price auction, an important note is in order: We assume that the players *do not bid more* than their valuation.

Definition 37 (No-overbidding bid vector, no-overbidding Bayesian strategy)**.** For a player i and a valuation $v_i \in V_i$, a *no-overbidding* bid vector is a bid vector $s_i \in \Sigma_i$ such that $\hat{s}_i(j) \leq v_i(j)$ for all $j \in [m]$ (where $\hat{s}$ is the vector defined by $\hat{s}_i(j) = \sum_{k \in [j]} s_i(k)$ for $j \in [m]$. A mixed Bayesian strategy $\sigma \in \Delta(\Sigma_i)^{V_i}$ of player $i \in [n]$ is said to be no-overbidding iff for all $v_i \in V_i$, a bid vector drawn from $\sigma(v_i)$ is a no-overbidding bid vector with probability 1.

This *no-overbidding assumption* implies, for uniform price auctions, that our analysis of the pure price of anarchy is actually with respect to only the subset of the pure equilibria that contain only no-overbidding bid vectors. Likewise, our analysis of the mixed Bayes-Nash price of anarchy for uniform price auctions is only with respect to the subset of mixed Bayes-Nash equilibria where it is almost surely the case that no player overbids. We refer to these equilibria as *no-overbidding pure equilibria* and *no-overbidding mixed Bayes-Nash equilibria* respectively, and to its corresponding prices of anarchy as the *no-overbidding pure price of anarchy* and *no-overbidding mixed Bayes-Nash price of anarchy* respectively. We emphasize that the bounds given for uniform price auctions in Table 2.1 and discussed in Section 2.2 are on the no-overbidding mixed Bayes-Nash price of anarchy. The bounds for discriminatory auctions do allow overbidding.

The assumption of no-overbidding that we make on uniform price auctions is standard, and is made due to the observation that without this assumption, unnatural equilibria can emerge. These equilibria may have bad social welfare and as a consequence they should not be part of a reasonable solution concept. More specifically: strategies in which players overbid are *dominated*, meaning that it is never beneficial for a player to bid more than its valuation; instead it may only harm the player. It is easy to see that if a player overbids in some given strategy profile, then it will always get at least as much utility if it instead switches to bidding precisely its valuation on the number of items that it wins.

Existence of equilibria. As the auctions we study are not full-information games in which the players have finite strategy sets, equilibria are not guaranteed to exist in multi-unit auctions. The bounds on the mixed Bayes-Nash price of anarchy that we present in this chapter are thus preconditioned on the assertion that there is at least one Bayes-Nash equilibrium in the auction. We omit stating this precondition in the theorems below. In case the reader finds this assertion unsatisfying, we note that our results can be easily seen to continue to hold in case we impose on our multi-unit auction the restriction that for each player there is a maximum amount of money that it may bid, and moreover it may only bid multiples of a (sufficiently small) prespecified least unit of money. This is a valid assumption for every realistic setting that one may think of. Moreover, if we impose this restriction on the strategy sets of the players, then the strategy sets become finite and consequently a mixed Bayes-Nash equilibrium is guaranteed to exist (by Theorem 15 in Section 1.3.1.6).

2.4 The Price of Anarchy of Full Information Multi-Unit Auctions

In this section we discuss the properties of pure equilibria of the two multi-unit auction formats, under both the standard and uniform bidding interfaces. As we show, pure equilibria are always efficient under the discriminatory auction, unlike the uniform price auction.

2.4.1 Uniform Pricing

Pure equilibria of the uniform price auction have been analyzed in Markakis and Telelis [2012]. It is well-known that for uniform price auctions, a pure Nash equilibrium that attains optimal social welfare is guaranteed to exist. Markakis and Telelis [2012] showed that the pure price of anarchy is $e/(e-1)$ under the standard bidding interface, when only the subset of *undominated* pure equilibria is considered.[1]

The following simple example shows that the inefficiency of pure equilibria of the uniform price auction for no-overbidding strategies is slightly higher than the bound of Markakis and Telelis [2012] (under the standard or uniform bidding interface):

Theorem 38. *The no-overbidding price of anarchy of the uniform price auction is at least* $(2m-1)/m$ *when all players have submodular valuation functions. This holds for both the standard and uniform bidding interface.*

Proof. Consider two players with submodular valuation functions, as follows:

$$v_1(j) = \begin{cases} m & \text{if } j \geq 1\,, \\ 0 & \text{otherwise}\,. \end{cases} \quad \text{and} \quad v_2(j) = j$$

The socially optimal allocation achieves social welfare $2m-1$, by allocating one item to player 1 and $m-1$ items to player 2. For a pure equilibrium, consider the uniform bid profile s where $s_1(j) = 1$ and $s_2(j) = 0$ for all $j \in [m]$. Clearly, $u_1(s) = m$, and player 2 cannot improve its utility by deviating. Therefore, s is an equilibrium with $U(s) = m$, and the no-overbidding pure price of anarchy is $(2m-1)/m$. □

For both bidding interfaces, upper bounds on the no-overbidding pure price of anarchy for uniform price auctions emerge from our results for the more general incomplete information model (Section 2.5), combined with the following lemma:

Lemma 39. *For arbitrary valuation functions, every no-overbidding pure equilibrium of the uniform price auction with uniform bidding is also a pure equilibrium of the auction under standard bidding.*

Proof. Let $\bar{s}$ be a no-overbidding pure equilibrium of the uniform price auction under the uniform bidding interface, for arbitrary valuation functions. We argue that it remains a pure equilibrium under the standard bidding interface. If any player i that does not win an item has an incentive to deviate using a (standard) bid s_i, in order to win at least one unit, it may do so also by using a uniform bid obtained from s_i by setting $s_i(j)$ to 0 for $j \geq 2$, a contradiction. If a player that wins an item under $\bar{s}$ has an incentive to deviate using a standard bid s_i, in order to win q_i items, then it may as well

[1] A strategy s_i of a player $i \in [n]$ is said to be *dominated* if for any choice of strategies of the other players, there is a different strategy s_i' that will guarantee the player achieve a utility that is at least as high as the utility it would get when it plays s (and moreover, there is at least one strategy profile s_{-i} of the opposing players such that $u_i(s_i, s_{-i}) < u_i(s_i', s_{-i})$). An undominated pure equilibrium is a pure equilibrium where everyone plays a strategy that is not dominated

do so with a uniform bid obtained from s_i by bidding $\sum_{j\in[q_i]} s_i(j)/q_i$ as the first q_i of its bid vector and 0 on the remaining elements, because $\sum_{j\in[q_i]} s_i(j)/q_i \geq s_i(q_i)$ due to submodularity of s_i. Thus, the assumed uniform bid should grant i at least as many units as s_i against a price of at most $p_i(\bar{s})$. □

The above lemma, along with Theorem 45 in Section 2.5, leads to the following:

Corollary 40. *The no-overbidding pure price of anarchy of the uniform price auction with submodular players under the standard or the uniform bidding interface is at most* $|W_{-1}(-1/e^2)|) \approx 3.1462$.

For classes of valuation functions that are more general than the submodular one, we do not know whether the uniform price auction generally has pure Nash equilibria. To the best of our knowledge, and as mentioned by Milgrom [2004], the standard multi-unit auction formats have not been studied before for any larger class of valuation functions. In Section 2.5 we give upper bounds of 4 and $2|W_{-1}(-1/e^2)|) \approx 6.2924$ on the no-overbidding mixed Bayes-Nash price of anarchy of uniform price auctions for subadditive valuation functions, under the standard and uniform bidding interfaces, respectively. By Lemma 39 however, the former bound of 4 is valid for the no-overbidding pure price of anarchy also for uniform bidding.

Corollary 41. *The no-overbidding pure price of anarchy of the uniform price auction with players with subadditive valuation functions under the standard or the uniform bidding interface is at least* 2 *and at most* 4.

The upper bound of this corollary follows by Lemma 39 and Theorem 46, discussed in Section 2.5. The lower bound is constructed in Section 2.5.3.

2.4.2 Discriminatory Pricing

The discriminatory auction is not guaranteed to possess pure equilibria; their existence depends heavily on the choice of a tie-breaking rule, as is often the case for games where players have a continuum of strategies. For example, consider a two player discriminatory auction with one item (i.e., first-price auction) where the valuation of player 1 is 1, the valuation of player 2 is $\epsilon < 1$, and the tie-breaking rule always favors player 2. Obviously there can be no equilibrium where player 2 bids above 1. Furthermore, if player 2 bids some value $\delta < 1$, then player 1 does not have a best response in $(\delta, 1)$; no matter what it bids to win the unit, it always has an incentive to lower its bid while still being above δ. Hence there is no pure equilibrium for this auction. We show here that, as with first-price auctions, an appropriate choice of a tie-breaking rule induces a *uniform bidding* profile that is a pure equilibrium, even under the standard bidding interface. Additionally, we show that we can always obtain close approximations to pure equilibria, i.e., pure ϵ-equilibria, for every possible tie breaking rule.

Proposition 42. *The following two statements hold:*

- *For every discriminatory auction there is a tie-breaking rule inducing a uniform bid profile that is a pure Nash equilibrium under that tie-breaking rule.*
- *For every $\epsilon > 0$, the discriminatory auction has a pure ϵ-equilibrium.*

Proof. We start with proving the first claim. Let $\check{v}$ be the nm-dimensional vector obtained by non-increasingly ordering all marginal valuations $\check{v}_i(j), i \in [n], j \in [m]$. We show that the set of bid profiles s where every player sets all its marginal bids to $\check{v}(m)$ is a pure equilibrium, if ties are broken according to any tie-breaking rule that satisfies $\check{v}_i(x_i(s)) \geq \check{v}(m)$.

Assume without loss of generality that there are at least two players. Let s be the bid profile where all players bid $\check{v}(m)$ on all items, and break ties in any way that satisfies that $\check{v}_i(x_i(s)) \geq \check{v}(m)$. To see why this is a pure equilibrium, consider the player deviating to bid vector s'_i. Note that $s'_i - s_i$ is non-increasing. Define ℓ as the lowest index such that $s_i(\ell) - s'_i(\ell)$ is negative (and define ℓ as $m+1$ if there is no such index). If $\ell \leq x_i(s)$, then the utility of player i will certainly not increase by deviating to s'_i, as it will lose utility from the fact that $x_i(s) - \ell$ less items are now allocated to it under (s'_i, s_{-i}), compared to s. As player i used to derive non-negative utility from these items under s, this removal of items accounts for a non-negative decrease in utility. Moreover, player i increases its bid on its first ℓ items, so this accounts for a non-negative decrease in utility as well. His total utility will therefore decrease in this case.

In case $\ell > x_i(s)$, we are in a situation where player i increases its bids (under s'_i, compared to s_i) on some of the first $x_i(s)$ items by at least 0, so it will win these items under (s'_i, s_{-i}) but spend more money on it, leading to a decrease in utility. On any remaining items that player i wins under (s'_i, s_{-i}), it overbids. This also accounts for a non-negative decrease in utility. The total decrease in utility is thus non-negative in this case.

We proceed with the proof of the second claim. Let s^* be a social welfare maximizing bid profile. Consider the uniform bid profile s as defined in the proof of the first claim. Let $\xi_1, \ldots, \xi_n$ be vectors that indicate an optimal allocation $x(s^*)$, i.e., ξ_i is the $(0, 1)$-vector of which the first $x_i(s^*)$ entries are 1, and the remaining entries are 0. We show that $\tilde{s} = s + \epsilon\xi/k$ is a pure ϵ-equilibrium. The reasoning we apply is largely analogous to the proof of the first claim.

First, observe that there are no ties that need to be broken under s, and that the allocation $x(\tilde{s})$ satisfies $\check{v}_i(x_i(s)) \geq \check{v}(m)$.

Consider the player deviating to a bid vector s'_i. If player i wins less items under $(s'_i, \tilde{s}_{-i})$ than under $\tilde{s}$, it will experience an increase in utility of at most $(x_i(\tilde{s}) - x_i(s'_i, \tilde{s}_{-i}))\epsilon/k$ due to losing items, because the utility that player i derived under $\tilde{s}$ from each of these lost items was at least $-\epsilon/k$. On the remaining items that player i still wins, the player decreases its bid by at most ϵ/k, and this accounts for an increase in utility of at most $x_i(s'_i, \tilde{s}_{-i})\epsilon/k$. The total increase in utility is thus at most $x_i(\tilde{s})\epsilon/k \leq \epsilon$.

If player i wins at least as much items under $(s'_i, \tilde{s}_{-i})$ than under $\tilde{s}$, the player will have decreased its bids on the first $x_i(\tilde{s})$ items by at most ϵ/k, and by at most 0 on the remaining items. For these remaining items, the player experiences non-positive utility under $(s'_i, \tilde{s}_{-i})$, whereas it experienced 0 utility under $\tilde{s}$. Therefore, the total increase in utility is in this case at most $x_i(\tilde{s})\epsilon/k \leq \epsilon$. $\square$

We show next that, whenever pure equilibria exist in a discriminatory auction, then they are socially optimal, even with arbitrary valuation functions (even valuations functions that are not-necessarily non-decreasing). This is in analogy with other results on mechanisms with "first price" rules [Hassidim et al., 2011].

Theorem 43. *The pure price of anarchy of discriminatory auctions is* 1*. This holds for both the standard and uniform bidding interface, and players with arbitrary valuation functions.*

The proof of Theorem 43 is based on the following lemma, that captures the main structural properties of pure equilibria in discriminatory auctions.

Lemma 44. *Let s be a pure Nash equilibrium in a given discriminatory auction where the players have general valuation functions. Let $d = \max\{s_i(j) \mid i \in [n], j \in [m], j > x_i(s)\}$. Then:*

1. *For each player $i \in [n]$ that wins at least one item under s, and for all $j \in [x_i(s)]$, $s_i(j) = d$,*

2. *$\ell d \leq \sum_{j=x_i(s)-\ell+1}^{x_i(s)} s_i(j)$ for all $i \in [n]$, and $\ell \in [x_i(s)]$,*

3. *$\sum_{j=x_i(s)+1}^{x_i(s)+\ell} \check{v}_i(j) \leq \ell d$ for all $i \in [n]$, and $\ell \in [m - x_i(s)]$.*

Proof. Let c be the smallest value in $\{s_i(j) \mid i \in [n], j \in [k], j \leq x_i(s)\}$, i.e., the smallest winning marginal bid. Observe that $c = d$: Otherwise, a player i that bids $s_i(x_i(s)) = c$ could change $s_i(x_i(s))$ to a lower bid in order to obtain more utility.

For the same reasons, we conclude that any winning marginal bid $s_i(j)$ is equal to the largest marginal bid that is smaller than $s_i(j)$. It follows inductively that all winning marginal bids are equal to d. This establishes point 1 of the claim.

Suppose that for some $i \in [n]$ and $\ell \in [x_i(s)]$, it holds that

$$\ell d = \sum_{j=x_i(s)-\ell+1}^{x_i(s)} s_i(j) > \sum_{j=x_i(s)-\ell+1}^{x_i(s)} \check{v}_i(j).$$

Then if player i changes all marginal bids $s_i(j)$ for $j \in \{j' \mid \ell \leq j'\}$ to 0, it would increase its utility. This is not possible since s is a pure equilibrium, so we conclude that for all $i \in [n]$ and $\ell \in [x_i(s)]$, it holds that

$$\sum_{j=x_i(s)-\ell+1}^{x_i(s)} s_i(j) \leq \sum_{j=x_i(s)-\ell+1}^{x_i(s)} \check{v}_i(j).$$

This establishes point 2 of the claim.

Note that there is no $i \in [n]$ and $\ell \in [m - x_i(s)]$ such that $\sum_{j=x_i(s)+1}^{x_i(s)+\ell} \check{v}_i(j) > \ell d$: Otherwise, if player i would change its marginal bids $s_i(j), j \in \{j' \mid 1 \leq j' \leq x_i(s) + \ell\}$ to $d+\epsilon$ for some $\epsilon > 0$, then player i's utility increases by $\sum_{j=x_i(s)+1}^{x_i(s)+\ell} \check{v}_i(j) - \ell(d+\epsilon) - x_i(s)\epsilon$. Because $\sum_{j=x_i(s)+1}^{x_i(s)+\ell} \check{v}_i(j) - \ell d$ is positive, this total increase is positive when we take for ϵ a sufficiently small value. This is in contradiction with the fact that s is a pure equilibrium, and this establishes point 3 of the claim. □

Proof of Theorem 43. Let s^* be a bid vector that attains the optimum social welfare. Denote by A the set of players that get more items under s than under s^*. For a player $i \in A$, define ℓ_i as the number of extra items that i gets under s, when compared to s^*; i.e., $\ell_i = x_i(s) - x_i(s^*)$. Denote by B the set of players that get more items under s^* than under s. For a player $i \in B$, define ℓ_i as the number of extra items that i gets under s^*, when compared to s; i.e., $\ell_i = x_i(s^*) - x_i(s)$. Then,

$$\begin{aligned}
\sum_{i\in[n]} v_i(x_i(s)) - \sum_{i\in[n]} v_i(x_i(s^*)) &= \sum_{i\in[n]} \left(\sum_{j\in[x_i(s)]} \check{v}_i(j) - \sum_{j\in[x_i(b^*)]} \check{v}_i(j) \right) \\
&= \sum_{i\in A} \sum_{j=x_i(s)-\ell_i+1}^{x_i(s)} \check{v}_i(j) - \sum_{i\in B} \sum_{j=x_i(s)+1}^{x_i(s)+\ell_i} \check{v}_i(j) \\
&\geq \sum_{i\in A} \ell_i d - \sum_{i\in B} \ell_i d = 0.
\end{aligned}$$

The inequality in the derivation above follows from points 2 and 3 of Lemma 44, and the final equality holds because $\sum_{i \in A} \ell_i = \sum_{i \in B} \ell_i$. Thus, the social welfare of the pure equilibrium s is optimal. □

2.5 The Price of Anarchy for Incomplete Information Multi-Unit Auctions

We derive bounds on the mixed Bayes-Nash price of anarchy for the discriminatory auctions and the no-overbidding mixed Bayes-Nash price of anarchy for the uniform price auctions, with submodular and subadditive valuation functions. The class of subadditive valuation functions has not been studied before in the literature of standard multi-unit auctions (see also the commentary in [Milgrom, 2004, Chapter 7]).

Theorem 45. *The mixed Bayes-Nash price of anarchy (under the standard or uniform bidding interface) is at most* $e/(e-1)$ *and* $2e/(e-1)$ *for the discriminatory price auction when the valuation functions are submodular and subadditive valuation functions, respectively. The no-overbidding mixed Bayes-Nash price of anarchy (under the standard or uniform bidding interface) is at most* $|W_{-1}(-1/e^2)| \approx 3.1462 < 2e/(e-1)$ *and* $2|W_{-1}(-1/e^2)| \approx 6.2924 < 4e/(e-1)$ *for the uniform price auction with submodular and subadditive valuation functions, respectively (*W_{-1} *being the lower branch of the Lambert W function).*

This theorem improves on the bounds of $2e/(e-1)$ and $4e/(e-1)$ for the discriminatory auction and the uniform price auction, respectively, with submodular valuation functions due to Syrgkanis and Tardos [2013]. For the uniform price auction, this further reduces the gap from the known lower bound of $e/(e-1)$ [Markakis and Telelis, 2012]. Syrgkanis and Tardos [2013] obtained their bounds through an adaptation of the *smoothness framework* for games with incomplete information [Roughgarden, 2012, Syrgkanis, 2012]. The bounds of Theorem 45 and some additional results can also be obtained through this framework. We comment on this in more detail in Section 2.6.

For subadditive valuation functions and the standard bidding interface, however, better bounds can be obtained by adapting a technique introduced by Feldman et al. [2013], which does not fall within the smoothness framework. We were unable to derive these bounds via a smoothness argument and believe that this is due to the additional flexibility provided by this technique.

Theorem 46. *The mixed Bayes-Nash price of anarchy is at most* 2 *and* 4 *for the discriminatory price auction with subadditive valuation functions under the standard bidding interface. The no-overbidding mixed Bayes-Nash price of anarchy is at most* 4

for the uniform price auction with subadditive valuation functions under the standard bidding interface.

2.5.1 Proof Template for Mixed Bayes-Nash Price of Anarchy

In order to present all our bounds from Theorem 45 and Theorem 46 in a self-contained and unified manner, we make use of a proof template which is formalized in Theorem 47 below. Variants of this approach have been used in several other works (e.g., Markakis and Telelis [2012], Christodoulou et al. [2008], Bhawalkar and Roughgarden [2011]).

Theorem 47. *Suppose that for every valuation profile $v \in V$, for every player $i \in [n]$, and for every probability distribution $\bar{\sigma}_{-i}$ on profiles of bid vectors in $\Sigma_{-i} = \times_{i' \in [n] \setminus \{i\}} \Sigma'_i$, there is a bid vector $s'_i \in \Sigma_i$ such that the following inequality holds for some $\lambda \in \mathbb{R}_{>0}$ and $\mu \in \mathbb{R}_{\geq 0}$:*

$$\mathbf{E}_{s_{-i} \sim \bar{\sigma}_{-i}} \left[u_i((s'_i, s_{-i}), v_i) \right] \geq \lambda v_i(x_i(s^v)) - \mu \mathbf{E}_{s_{-i} \sim \bar{\sigma}_{-i}} \left[\sum_{j \in [x_i(s^v)]} \beta_{m-j+1}(s_{-i}) \right]. \tag{2.1}$$

Then:

- *for the discriminatory auction, the mixed Bayes-Nash price of anarchy is at most $1/\lambda$ if $\mu \leq 1$,*
- *If $\bar{\sigma}_{-i}$ has only no-overbidding bid vectors in its support, then the no-overbidding mixed Bayes-Nash price of anarchy of discriminatory auctions is at most $(\mu - 1)/\lambda$,*
- *If $\bar{\sigma}_{-i}$ has only no-overbidding bid vectors in its support, then the no-overbidding mixed Bayes-Nash price of anarchy for uniform price auctions is at most $(\mu + 1)/\lambda$.*

Note that in this theorem we make no assumptions regarding which of the two bidding interfaces is used. Proving a bound for the uniform bidding interface only requires that we provide a uniform bid vector s'_i for each player $i \in [n]$ and for each distribution $\bar{\sigma}_{-i}$ on profiles of uniform non-overbidding vectors s_{-i}.

Proof of Theorem 47. Consider a mixed Bayes-Nash equilibrium σ. Fix any valuation profile $v = (v_i, v_{-i}) \in V$ and a player $i \in [n]$. Assume that player i deviates according to a bid vector s'_i satisfying (2.1), where $\bar{\sigma}_{-i} = \bar{\sigma}_{-i}(w_{-i}), w_{-i} \sim \pi_{-i}$. Taking expectation over all valuation profiles $w_{-i} \in V_{-i} = \times_{j\in[n]\setminus\{i\}} V_j$, we obtain

$$\begin{aligned}
&\mathbf{E}_{w_{-i}\sim\pi_{-i}}\left[\mathbf{E}_{s_{-i}\sim\bar{\sigma}_{-i}(w_{-i})}\left[u_i((s'_i, s_{-i}), v_i)\right]\right] \\
&\quad\geq \lambda v_i(x_i(s^v)) - \mu \mathbf{E}_{w_{-i}\sim\pi_{-i}}\left[\mathbf{E}_{s_{-i}\sim\bar{\sigma}_{-i}(w_{-i})}\left[\sum_{j\in[x_i(s^v)]} \beta_{m-j+1}(s_{-i})\right]\right] \\
&\quad\geq \lambda v_i(x_i(s^v)) - \mu \mathbf{E}_{w\sim\pi}\left[\mathbf{E}_{s\sim\bar{\sigma}(w)}\left[\sum_{j\in[x_i(s^v)]} \beta_{m-j+1}(s)\right]\right],
\end{aligned}$$

where the last inequality holds because $\beta_{m-j+1}(s_{-i}) \leq \beta_{m-j+1}(s)$ for every $j \in [m]$. Because σ is a mixed Bayes-Nash equilibrium, player i does not have an incentive to deviate and thus

$$\mathbf{E}_{w_{-i}\sim\pi_{-i}}\left[\mathbf{E}_{s\sim\bar{\sigma}(v_i,w_{-i})}\left[u_i(s, v_i)\right]\right] \geq \mathbf{E}_{w_{-i}\sim\pi_{-i}}\left[\mathbf{E}_{s_{-i}\sim\bar{\sigma}_{-i}(w_{-i})}\left[u_i(b'_i, b_{-i}), v_i)\right]\right].$$

We conclude that

$$\begin{aligned}
&\mathbf{E}_{w_{-i}\sim\pi_{-i}}\left[\mathbf{E}_{s\sim\bar{\sigma}(v_i,w_{-i})}\left[u_i(s, v_i)\right]\right] + \mu \mathbf{E}_{w\sim\pi}\left[\mathbf{E}_{s\sim\bar{\sigma}(w)}\left[\sum_{j\in[x_i(s^v)]} \beta_{m-j+1}(s)\right]\right] \\
&\quad\geq \lambda v_i(x_i(s^v)).
\end{aligned}$$

Taking the sum over all players and all valuation functions on both sides gives us

$$\begin{aligned}
&\sum_{i\in[n]}\sum_{v\in V}\pi(v)\Bigg(\mathbf{E}_{w_{-i}\sim\pi_{-i}}\left[\mathbf{E}_{s\sim\bar{\sigma}(v_i,w_{-i})}\left[u_i(s, v_i)\right]\right] \\
&\qquad + \mu \mathbf{E}_{w\sim\pi}\left[\mathbf{E}_{s\sim\bar{\sigma}(w)}\left[\sum_{j\in[x_i(s^v)]} \beta_{m-j+1}(s)\right]\right]\Bigg) \\
&\qquad\geq \sum_{i\in[n]}\sum_{v\in V}\pi(v)\lambda v_i(x_i(s^v)) \\
&\qquad= \sum_{v}\pi(v)\sum_{i\in[n]}\lambda v_i(x_i(s^v)) \\
&\qquad= \lambda \mathbf{E}_{v\sim\pi}[U(s^v, v)].
\end{aligned}$$

Simplifying the left hand side of the above inequality, we obtain

$$\mathbf{E}_{v\sim\pi}\left[\mathbf{E}_{s\sim\bar{\sigma}(v)}\left[\sum_{i\in[n]} u_i(s, v_i)\right]\right] + \mu \mathbf{E}_{w\sim\pi}\left[\mathbf{E}_{s\sim\bar{\sigma}(w)}\left[\sum_{i\in[n]}\sum_{j\in[x_i(s^v)]} \beta_{m-j+1}(s)\right]\right]$$

$$\geq \lambda \mathbf{E}_{v\sim\pi}[U(s^v, v)].$$

Note that $\sum_{i\in[n]} x_i(s^v) = m$ and that $\beta_{m-j+1}(s)$ is non-decreasing in j. We can therefore bound

$$\sum_{i\in[n]} \sum_{j\in[x_i(s^v)]} \beta_{m-j+1}(s) \leq \sum_{j\in[m]} \beta_{m-j+1}(s)$$

and obtain

$$\mathbf{E}_{v\sim\pi}\left[\mathbf{E}_{s\sim\bar{\sigma}(v)}\left[\sum_{i\in[n]} u_i(s, v_i)\right]\right] + \mu \mathbf{E}_{w\sim\pi}\left[\mathbf{E}_{s\sim\bar{\sigma}(w)}\left[\sum_{j\in[m]} \beta_{m-j+1}(s)\right]\right]$$
$$\geq \lambda \mathbf{E}_{v\sim\pi}[U(s^v, v)]. \tag{2.2}$$

Note that for the discriminatory pricing rule the total payments under $s \in \Sigma$ are equal to $\sum_{j\in[m]} \beta_{m-j+1}(s)$. Thus (2.2) yields

$$\mathbf{E}_{v\sim\pi}\left[\mathbf{E}_{s\sim\bar{\sigma}(v)}\left[U(s, v)\right]\right] + (\mu - 1)\mathbf{E}_{w\sim\pi}\left[\mathbf{E}_{s\sim\bar{\sigma}(w)}\left[\sum_{j\in[m]} \beta_{m-j+1}(s)\right]\right]$$
$$\geq \lambda \mathbf{E}_{v\sim\pi}[U(s^v, v)].$$

If $\mu \leq 1$, the first statement of the theorem holds. If $\mu > 1$, then we exploit that the total payments satisfy $\sum_{j\in[m]} \beta_{m-j+1}(s) \leq \sum_{i\in[n]} v_i(x_i(s)) = U(s, v)$ because players never overbid. Dividing both sides by $\mu > 0$ proves the first statement of the theorem in this case.

For the uniform pricing rule we use that $\sum_{j\in[m]} \beta_{m-j+1}(s) \leq \sum_{i\in[n]} v_i(x_i(s)) = U(s, v)$ because of the no-overbidding assumption, and that

$$\sum_{i\in[n]} u_i(s, v_i) \leq \sum_{i\in[n]} v_i(x_i(s)) = U(s, v).$$

Thus (2.2) yields

$$(\mu + 1)\mathbf{E}_{v\sim\pi}\left[\mathbf{E}_{s\sim\bar{\sigma}(v)}\left[U(s, v)\right]\right] \geq \lambda \mathbf{E}_{v\sim\pi}[U(s^v, v)].$$

Dividing both sides by $\mu + 1 > 0$ proves the claim. □

In Section 2.5.3 we show that our bound of $e/(e-1)$ for the discriminatory price auction is essentially best possible if one sticks to the proof template of Theorem 47. This rules out that better bounds can be obtained via the techniques in Syrgkanis and Tardos [2013], Feldman et al. [2013].

2.5.2 Key Lemma and Proofs of Theorem 45 and Theorem 46

The following is our key lemma to prove Theorem 45. We point out that it applies to arbitrary valuation functions (even those that are not subadditive) and to any multi-unit auction that is *discriminatory price dominated*, i.e., the total payment $p_i(s)$ of player $i \in [n]$ under profile $s \in \Sigma$ satisfies $p_i(s) \leq \sum_{j \in [x_i(s)]} s_i(j)$.

Lemma 48 (Key Lemma)**.** *Let v be a valuation profile and suppose that the pricing rule is discriminatory price dominated. Define $\tau_i = \arg_j \min\{v_i(j)/j \mid j \in [x_i(s^v)]\}$ for every $i \in [n]$. Then for every player $i \in [n]$ there exists a no-overbidding randomized uniform bid strategy $\bar{\sigma}'_i(v_i)$ (i.e., a probability distribution on the subset of uniform bid vectors of Σ_i) such that for every bid profile s_{-i}:*

$$\mathbf{E}_{s'_i \sim \bar{\sigma}'_i(v_i)}[u_i((s'_i, s_{-i}), v_i)] \geq \alpha \left(1 - \frac{1}{e^{1/\alpha}}\right) x_i(s^v) \frac{v_i(\tau_i)}{\tau_i} - \alpha \sum_{j \in [x_i(s^v)]} \beta_{m-j+1}(s_{-i}). \tag{2.3}$$

Proof. Define $B = (1 - e^{-1/\alpha})$ and let c_i be the vector of which the first $x_i(s^v)$ elements are $v_i(\tau_i)/\tau_i$, and the other elements are 0. Let t be a random variable with probability distribution $f(t) = \alpha/(1-t)$, on $[0, B]$. Define $\bar{\sigma}'_i$ as the probability distribution of $s'_i = tc_i$. Note that s'_i is indeed always a uniform bid vector.

Let m^* be the number of items that player i would win under profile (Bc_i, s_{-i}), i.e., the number of items won by i, when i would deviate to bid vector Bc_i. For $j \in [m^*] \cup \{0\}$, let γ_j refer to the infimum value in $[0, B]$ such that player i would win j items if it would deviate to bid vector $\gamma_j c_i$. Note that this definition is equivalent to defining γ_j as the least value in $[0, B]$ that satisfies $\gamma_j v_i(\tau_i)/\tau_i = \beta_{m-j+1}(b_{-i})$. For notational convenience, we define $\gamma_{m^*+1} = B$.

Let $x_i(s'_i, s_{-i})$ be the random variable that denotes the number of units allocated to player $i \in [n]$ under (s'_i, s_{-i}). It always holds that $x_i(s'_i, s_{-i}) \leq m^* \leq x_i(s^v)$, because player i bids $s'_i(j) = 0$ for all $j \in \{x_i(s^v) + 1, \ldots, m\}$. More precisely, we have $x_i(s'_i, s_{-i}) = j$ if $t \in (\gamma_j, \gamma_{j+1}]$ for $j \in [k^*] \cup \{0\}$. By assumption, the payment of player i under profile (s'_i, s_{-i}) is at most $tx_i(s'_i, s_{-i})v_i(\tau_i)/\tau_i$. Also note that, by definition of τ_i, it holds that $v_i(j) \geq jv_i(\tau_i)/\tau_i$ for $j \leq x_i(s^v)$. Using these two facts, we can bound the expected utility of player i as follows:

$$\begin{aligned}
\mathbf{E}_{s'_i \sim \bar{\sigma}'_i(v_i)}[u_i((s'_i, s_{-i}), v_i)] &\geq \sum_{j \in [m^*]} \int_{\gamma_j}^{\gamma_{j+1}} \left(v_i(j) - tj\frac{v_i(\tau_i)}{\tau_i}\right) f(t)dt \\
&\geq \sum_{j \in [m^*]} \int_{\gamma_j}^{\gamma_{j+1}} j\frac{v_i(\tau_i)}{\tau_i}(1-t)f(t)dt
\end{aligned}$$

$$
\begin{aligned}
&= \alpha \sum_{j\in[m^*]} j\frac{v_i(\tau_i)}{\tau_i}\int_{\gamma_j}^{\gamma_{j+1}} 1dt \\
&= \alpha \sum_{j\in[m^*]} j\frac{v_i(\tau_i)}{\tau_i}(\gamma_{j+1}-\gamma_j) \\
&= \alpha B m^*\frac{v_i(\tau_i)}{\tau_i} - \alpha\sum_{j\in[m^*]}\gamma_j\frac{v_i(\tau_i)}{\tau_i} \\
&= \alpha B m^*\frac{v_i(\tau_i)}{\tau_i} - \alpha\sum_{j\in[m^*]}\beta_{m-j+1}(s_{-i}) \\
&\geq \alpha B x_i(s^v)\frac{v_i(\tau_i)}{\tau_i} - \alpha\sum_{j\in[x_i(s^v)]}\beta_{m-j+1}(s_{-i}).
\end{aligned}
$$

The last inequality holds because $Bv_i(\tau_i)/\tau_i \leq \beta_{m-j+1}(s_{-i})$, for $j \in \{m^*+1, \ldots, x_i(s^v)\}$, by the definition of m^*. The above derivation implies (2.3). □

The bid $\bar{\sigma}_i'$ defined in Lemma 48 is a distribution on uniform bid strategies. That is, the lemma applies to both the standard and the uniform bidding format. Observe also that $\bar{\sigma}_i'$ satisfies the no-overbidding assumption.

Proof of Theorem 45. First consider the case of submodular valuation functions. In this case, $\tau_i = x_i(s^v)$ for every $i \in [n]$, by Proposition 36. Using our Key Lemma, we conclude that Theorem 47 holds for $(\lambda,\mu) = (\alpha\left(1-e^{-1/\alpha}\right),\alpha)$, for all $\alpha \in \mathbb{R}_{>0}$. The stated bounds are obtained by choosing $\alpha = 1$ for the discriminatory auction and $\alpha = -1/(W_{-1}(-1/e^2)+2) \approx 0.87$ for the uniform price auction.

Next consider the case of subadditive valuation functions. The following lemma shows that subadditive valuation functions can be approximated by uniform ones, only losing at most a factor 2.

Lemma 49. *If v_i is a subadditive valuation function, then*

$$
\frac{v_i(\tau_i)}{\tau_i} \geq \frac{v_i(x_i(s^v))}{2x_i(s^v)},
$$

where $\tau_i = \arg_j \min\{v_i(j)/j \mid j \in [x_i(s^v)]\}$.

Proof. We analyze two cases: $\tau_i \leq x_i(s^v) \leq 2\tau_i$, and $x_i(s^v) > 2\tau_i$. (Note that the case $x_i(s^v) < \tau_i$ is impossible by the definition of τ_i.)

In case $\tau_i \le x_i(s^v) \le 2\tau_i$ it holds that

$$\begin{aligned}
\frac{v_i(x_i(s^v))}{2x_i(s^v)} &\le \frac{1}{2}\left(\frac{v_i(x_i(s^v - \tau_i))}{x_i(s^v)} + \frac{v_i(\tau_i)}{x_i(s^v)}\right) \\
&\le \frac{1}{2}\left(\frac{v_i(x_i^v - \tau_i)}{\tau_i} + \frac{v_i(\tau_i)}{\tau_i}\right) \\
&\le \frac{v_i(\tau_i)}{\tau_i},
\end{aligned}$$

where we use subadditivity for the first inequality and monotonicity for the third inequality (i.e., $v_i(x_i(s^v) - \tau_i) \le v_i(\tau_i)$).

In case $x_i(s^v) > 2\tau_i$ it holds that

$$\begin{aligned}
\frac{v_i(x_i(s^v))}{2x_i(s^v)} &\le \frac{1}{2}\left(\frac{v_i(x_i(s^v - \tau_i))}{x_i(s^v)} + \frac{v_i(\tau_i)}{x_i(s^v)}\right) \\
&\le \frac{1}{2}\left(\frac{v_i(x_i(s^v - \tau_i))}{x_i(s^v)} + \frac{v_i(\tau_i)}{\tau_i}\right) \\
&\le \frac{v_i(\tau_i)}{\tau_i},
\end{aligned}$$

where the first inequality holds by subadditivity, and the second inequality holds because $(v_i(x_i(s^v) - \tau_i))/x_i(s^v) = (v_i(x_i(s^v) - \tau_i))/(x_i(s^v) - \tau_i + \tau_i) \le v_i(\tau_i)/\tau_i$, which follows from the second point of Proposition 36. □

By combining Lemma 49 with our Key Lemma, it follows that Theorem 47 holds for

$$(\lambda, \mu) = \left(\frac{\alpha}{2}\left(1 - e^{-1/\alpha}\right), \alpha\right).$$

The bounds stated in Theorem 45 are obtained by the same choices of α as for submodular valuation functions. □

Next, consider subadditive valuations under the *standard* bidding interface. We derive improved bounds of 2 and 4 for the discriminatory and uniform price auction, respectively. To this end, we adapt a technique by Feldman et al. [2013] to establish an analog of Lemma 48. The main idea is to construct the bid $s_i' \in \Sigma_i$ by using the distribution $\bar{\sigma}_{-i}(v), v \sim \pi$ on Σ_{-i}. Theorem 46 then follows from Theorem 47 in combination with Lemma 50 below.

Lemma 50. *The conditions of Theorem 47 hold for $(\lambda, \mu) = (1/2, 1)$ for discriminatory auctions with subadditive valuation functions (both under the standard and* *uniform bidding interface). The conditions of Theorem 47 hold for* $(\lambda, \mu) = (1/2, 1)$ *for uniform price auctions with subadditive valuation functions, when there are only no-overbidding bid-vectors in the support of* $\bar{\sigma}_{-i}$ *(both under the standard and uniform bidding interface).*

Proof. Let V be the class of subadditive valuation functions. We consider first the discriminatory auction. We shall re-use certain arguments from the proof of this case, for the proof of the case of uniform price auctions. Fix any player $i \in [n]$, number $x \in [m]$, valuation $v \in V$, and let $s_i' \in \Sigma$ be a bid vector, satisfying the requirements of the standard bidding interface, and having only its first x components equal to a non-zero value, for some $x \in [m]$. Given any bid profile $s_{-i} \in \Sigma_{-i}$:

$$u_i((s_i', s_{-i}), v_i) \geq v_i(x_i(s_i', s_{-i})) - \sum_{j \in [x]} s_i'(j),$$

because i may pay at most $\sum_{j \in [m]} s_i'(j) = \sum_{j \in [x]} s_i'(j)$, by the definition of s_i' and the auction's payment rule. Taking expectation over the distribution $\bar{\sigma}_{-i}$ defined as $\bar{\sigma}_{-i}(v), v \sim V$, we have:

$$\mathbf{E}_{s_{-i} \sim \bar{\sigma}_{-i}} \left[u_i((s_i', s_{-i}), v_i) \right] \geq \mathbf{E}_{s_{-i} \sim \bar{\sigma}_{-i}} \left[v_i(x_i(s_i', s_{-i})) \right] - \sum_{j \in [x]} s_i'(j). \tag{2.4}$$

From this point on we analyze the right-hand side of (2.4). Given the distribution $\bar{\sigma}_{-i}$ of $s_{-i}(v), v \sim \pi, s_{-i} \sim \bar{\sigma}_{-i}(v)$, let D denote the distribution of the vector $(\beta_1(s_{-i}), \ldots, \beta_m(s_{-i})), s_{-i} \sim \bar{\sigma}_{-i}$. For every fixed bid vector of player i, the expected utility of i when the other players bid according to $\bar{\sigma}_{-i}$, is equal to the expected utility of player i in the two-player auction, where the other player bids according to D. We can thus assume that i competes only against $\beta \sim D$.

We consider what happens when i responds to $\bar{\sigma}_{-i}$ (i.e., D, in the two-players auction), by bidding a s_i' that is drawn from a random distribution constructed as follows: the player first samples a vector β from D and *zeroes out* the $k - x$ *highest* values in β and reorders the vector such that it is non-increasing. Subsequently, the player adds to all non-zero components of the "truncated" vector a sufficiently small $\epsilon > 0$. Let $\tilde{D}$ denote the distribution of s_i'. Moreover, for a vector β in the support of D, we denote by $\tilde{\beta}$ the vector obtained from β by executing on it the process just described (i.e., zeroing out the $k - x$ highest values, reordering non-increasingly, and adding ϵ to each of its elements). Continuing from (2.4), the expected utility of i over $s_i' \sim \tilde{D}$ is:

$$\mathbf{E}_{s_i' \sim \tilde{D}} \left[\mathbf{E}_{s_{-i} \sim \bar{\sigma}_{-i}} \left[u_i((s_i', s_{-i}), v_i) \right] \right]$$

$$
\begin{aligned}
&\geq \mathbf{E}_{s'_i \sim \tilde{D}} \left[\mathbf{E}_{s_{-i} \sim \bar{\sigma}_{-i}} \left[v_i(x_i(s'_i, s_{-i})) \right] - \sum_{j \in [x]} s'_i(j) \right] \\
&\geq \mathbf{E}_{\substack{\beta' \sim D \\ s_{-i} \sim \bar{\sigma}_{-i}}} \left[v_i(x_i(\tilde{\beta}', s_{-i})) - \sum_{j \in [m] \setminus [m-x]} \beta'_j \right] \\
&= \mathbf{E}_{\substack{\beta' \sim D \\ \beta \sim D}} \left[v_i(x_i(\tilde{\beta}', \beta)) \right] - \mathbf{E}_{\substack{\beta' \sim D \\ \beta \sim D}} \left[\sum_{j \in [m] \setminus [m-x]} \beta'_j \right] \\
&= \frac{1}{2} 2\mathbf{E}_{\substack{\beta' \sim D \\ \beta \sim D}} \left[v_i(x_i(\tilde{\beta}', \beta)) \right] - \mathbf{E}_{\beta \sim D} \left[\sum_{j \in [m] \setminus [m-x]} \beta_j \right] \\
&= \frac{1}{2} \mathbf{E}_{\substack{\beta' \sim D \\ \beta \sim D}} \left[v_i(x_i(\tilde{\beta}', \beta)) + v_i(x_i(\tilde{\beta}, \beta')) \right] - \mathbf{E}_{\beta \sim D} \left[\sum_{j \in [m] \setminus [m-x]} \beta_j \right] \\
&\geq \frac{1}{2} \mathbf{E}_{\substack{\beta' \sim D \\ \beta \sim D}} \left[v_i(x) \right] - \mathbf{E}_{\beta \sim \mathcal{D}} \left[\sum_{j \in [m] \setminus [m-x]} \beta_j \right] \\
&= \frac{1}{2} v_i(x) - \mathbf{E}_{s_{-i} \sim \bar{\sigma}_{-i}} \left[\sum_{j \in [m] \setminus [m-x]} \beta_j(s_{-i}) \right],
\end{aligned}
$$

where the last inequality holds by subadditivity of v_i, particularly because $x_i(\tilde{\beta}', \beta) + x_i(\tilde{\beta}, \beta') \geq x$. Taking for x the value $x_i(s^v)$ and applying Theorem 47 proves the claim for discriminatory auctions.

We now move on to uniform price auctions. Fix any player $i \in [n]$ and $v \in V$, and let $s'_i \in \Sigma_i$ be a bid vector with non-zero value only on its first x elements. Notice that, given any bid profile $s_{-i} \in \Sigma_{-i}$:

$$
\begin{aligned}
u_i((s'_i, s_{-i}), v_i) &= v_i(x_i(s'_i, s_{-i})) - x_i(s'_i, s_{-i}) p(s'_i, s_{-i}) \\
&\geq v_i(x_i(s'_i, s_{-i})) - \sum_{j \in [x]} s'_i(j).
\end{aligned}
$$

Taking expectation over the distribution $\bar{\sigma}_{-i}$ of $s_{-i}(v), v \sim \pi, s_{-i} \sim \bar{\sigma}_{-i}(v)$ we have:

$$
\mathbf{E}_{s_{-i} \sim \bar{\sigma}_{-i}} \left[u_i((s'_i, s_{-i}), v_i) \right] \geq \mathbf{E}_{s_{-i} \sim \bar{\sigma}_{-i}} \left[v_i(x_i(s'_i, s_{-i})) \right] - \sum_{j \in [m]} s'_i(j). \tag{2.5}
$$

Our analysis from this point on will focus on identifying an appropriate bid s'_i for player i that satisfies no-overbidding. Subsequently, we shall return to (2.5) and rewrite it further. As previously, given the distribution $\bar{\sigma}_{-i}$ of s_{-i}, we let D denote the distribution of the vector $(\beta_1(s_{-i}), \ldots, \beta_m(s_{-i})), s_{-i} \sim \bar{\sigma}_{-i}$.

First, let T be a function from the support of D to subsets of $[m] \setminus [m-x]$, where for all β in D's support, $T(\beta) \subseteq [m] \setminus [m-x]$ is a maximal subset of indices such that: $\sum_{j \in T(\beta)} \beta_j > v_i(|T(\beta)|)$. Let $\bar{T}(\beta) = ([m] \setminus [m-x]) \setminus T(\beta)$. Then, we claim that:

$$\sum_{j \in \bar{T}_\beta} \beta_j \leq v_i(|\bar{T}(\beta)|).$$

Indeed, if there exists $R \subseteq \bar{T}(\beta)$ with $\sum_{j \in R} \beta_j > v_i(|R|)$, then, by subadditivity of v_i and monotonicity:

$$v_i(|R \cup T(\beta)|) \leq v_i(|R|) + v_i(|T(\beta)|) < \sum_{j \in R} \beta_j + \sum_{j \in T(\beta)} \beta_j = \sum_{j \in R \cup T(\beta)} \beta_j.$$

which contradicts the maximality of $T(\beta)$. Next, using T we define $\tilde{D}$ to be a distribution defined as follows. For every vector β in the support of D that occurs with a certain probability, a vector $\tilde{\beta}$ exists in the support of $\tilde{D}$, that occurs with the same probability and is obtained from β by setting to 0 the elements of β at indices $\{x+1, \ldots, m\} \cup T(\beta)$, adding a sufficiently small number $\epsilon > 0$ to each of the element at indices $\bar{T}(\beta)$, and finally reordering the resulting vector non-increasingly.

Sampling a vector from $\tilde{D}$ is equivalent to sampling a vector β from D and constructing $\tilde{\beta}$ as just prescribed. For a vector β in the support of D, we denote by $\tilde{\beta}$ the vector obtained from β by executing on it the process just described, and we denote by $\bar{\beta}$ the vector obtain from β by executing on it the process just described, except for the last step (adding the ϵ). For any $\beta \sim D$ and for any arbitrary bid s'_i, we observe that $x_i(s'_i, \bar{\beta}) \leq x_i(s'_i, \beta) + |T_\beta|$. Thus, $v_i(x_i(s'_i, \bar{\beta})) \leq v_i(x_i(s'_i, \beta)) + v_i(|T(\beta)|)$, by subadditivity and monotonicity of v_i. Using $\sum_{j \in [m] \setminus [m-x]} \beta_j - \sum_{j \in \bar{T}_\beta} \beta_j = \sum_{j \in T_\beta} \beta_j \geq v_i(|T_\beta|)$, we obtain:

$$v_i(x_i(s'_i, \beta)) - \sum_{j \in \bar{T}(\beta)} \beta_j \geq v_i(x_i(s'_i, \tilde{\beta})) - \sum_{j \in [m] \setminus [m-x]} \beta_j.$$

Thus:

$$\mathbf{E}_{\beta \sim D} \left[v_i(x_i(s'_i, \beta)) - \sum_{j \in \bar{T}(\beta)} \beta_j \right] \geq \mathbf{E}_{\beta \sim D} \left[v_i(x_i(s'_i, \tilde{\beta})) - \sum_{j \in [m] \setminus [m-x]} \beta_j \right]. \tag{2.6}$$

Now consider s'_i being drawn from the distribution $\tilde{D}$. Notice that, by their construction, all the bid vectors in the support of $\tilde{D}$ satisfy no-overbidding. We take expectation of the left-hand side of (2.6) over $s'_i \sim \tilde{D}$, and rewrite:

$$\mathbf{E}_{\substack{s'_i \sim \tilde{D} \\ \beta \sim D}} \left[v_i(x_i(s'_i, \beta)) - \sum_{j \in \bar{T}(\beta)} \beta_j \right] = \mathbf{E}_{\substack{\beta' \sim D \\ \beta \sim D}} \left[v_i(x_i(\tilde{\beta}', \beta)) - \sum_{j \in \bar{T}(\beta)} \beta_j \right]$$

$$= \mathbf{E}_{\substack{\beta' \sim D \\ \beta \sim D}} \left[v_i(x_i(\tilde{\beta}', \beta)) \right] - \mathbf{E}_{\beta' \sim D} \left[\sum_{j \in \bar{T}(\beta')} \beta'_j \right]. \tag{2.7}$$

Accordingly, we take expectation of the right-hand side of (2.6) over $\tilde{D}$, to obtain:

$$\mathbf{E}_{\substack{s'_i \sim \bar{D} \\ \beta \sim D}} \left[v_i(x_i(s'_i, \tilde{\beta})) - \sum_{j \in [m] \setminus [m-x]} \beta_j \right] \tag{2.8}$$

$$= \mathbf{E}_{\substack{\beta' \sim D \\ \beta \sim D}} \left[v_i(x_i(\tilde{\beta}', \bar{\beta})) \right] - \mathbf{E}_{\beta' \sim D} \left[\sum_{j \in [m] \setminus [m-x]} \beta'_j \right]. \tag{2.9}$$

Finally, we take expectation of (2.5) over $s'_i \sim \tilde{D}$, to obtain:

$$\mathbf{E}_{\substack{s'_i \sim \tilde{D} \\ s_{-i} \sim \bar{\sigma}_{-i}}} [u_i((s'_i, s_{-i}), v_i)] \geq \mathbf{E}_{\substack{s'_i \sim \tilde{D} \\ s_{-i} \sim \bar{\sigma}_{-i}}} \left[v_i(x_i(s'_i, s_{-i})) - \sum_{j \in [m]} s'_i(j) \right]$$

$$= \mathbf{E}_{\substack{\beta' \sim D \\ \beta \sim D}} \left[v_i(x_i(\tilde{\beta}', \beta)) - \sum_{j \in \bar{T}(\beta')} \beta'_j \right]$$

$$= \mathbf{E}_{\substack{\beta' \sim D \\ \beta \sim D}} \left[v_i(x_i(\tilde{\beta}', \beta)) \right] - \mathbf{E}_{\beta' \sim D} \left[\sum_{j \in \bar{T}(\beta')} \beta'_j \right]. \tag{2.10}$$

By (2.10), (2.9), (2.7), and (2.6), we derive:

$$\mathbf{E}_{\substack{s'_i \sim \bar{D} \\ s_{-i} \sim \bar{\sigma}_{-i}}} [u_i((s'_i, s_{-i}), v_i)] \geq \mathbf{E}_{\substack{\beta' \sim D \\ \beta \sim D}} \left[v_i(x_i(\tilde{\beta}', \bar{\beta})) \right] - \mathbf{E}_{\beta' \sim D} \left[\sum_{j \in [m] \setminus [m-x]} \beta'_j \right].$$

The lower bounding by $v_i(x)/2$ of the first term of the right-hand side of this expression is done completely analogously to what we did for the discriminatory auction.

$$
\begin{aligned}
&\mathbf{E}_{\substack{s_i'\sim\bar{D}\\ s_{-i}\sim\bar{\sigma}_{-i}}}[u_i((s_i', s_{-i}), v_i)]\\
&\quad\geq \frac{1}{2}2\mathbf{E}_{\substack{\beta'\sim D\\ \beta\sim D}}\left[v_i(x_i(\tilde{\beta}', \bar{\beta}))\right] - \mathbf{E}_{\beta\sim D}\left[\sum_{j\in[m]\setminus[m-x]}\beta_j\right]\\
&\quad= \frac{1}{2}\mathbf{E}_{\substack{\beta'\sim D\\ \beta\sim D}}\left[v_i(x_i(\tilde{\beta}', \bar{\beta}) + v_i(x_i(\tilde{\beta}, \bar{\beta}'))\right] - \mathbf{E}_{\beta\sim D}\left[\sum_{j\in[m]\setminus[m-x]}\beta_j\right]\\
&\quad\geq \frac{1}{2}\mathbf{E}_{\substack{\beta'\sim D\\ \beta\sim D}}[v_i(x)] - \mathbf{E}_{\beta\sim D}\left[\sum_{j\in[m]\setminus[m-x]}\beta_j\right]\\
&\quad\geq \frac{1}{2}v_i(x) - \mathbf{E}_{s_{-i}\sim\bar{\sigma}_{-i}}\left[\sum_{j\in[m]\setminus[m-x]}\beta_j(s_{-i})\right].
\end{aligned}
$$

This establishes our claim for the uniform price auction. □

2.5.3 Lower Bounds

We provided in in Section 2.4 a lower bound of 2 for the no-overbidding price of anarchy for pure equilibria, for the case of uniform price auctions with submodular valuation functions. This lower bound trivially carries over to the mixed Bayes-Nash price of anarchy of uniform price auctions with submodular valuation functions.

No lower bounds are known for the price of anarchy of discriminatory price auctions, although *demand reduction* (which is responsible for welfare loss in this format) has been observed previously [Krishna, 2002, Ausubel and Cramton, 2002]. In light of this, we prove here an *impossibility result* showing that for the discriminatory price auction, no bound better than $e/(e-1)$ on the price of anarchy can be achieved via the proof template given in Theorem 47.

Theorem 51. *For players with submodular valuation functions, the proof template given in Theorem 47 cannot be used to obtain an upper bound on the mixed Bayes-Nash price of anarchy that is less than $e/(e-1)$ for the discriminatory price auction.*

Before giving the formal proof, we provide some intuition for how this proof works. Our goal is to find a instance of a discriminatory price auction such that (2.1) can only be satisfied with an as large as possible value for $\max\{1, \mu-1\}/\lambda$. To make this search easier, we restrict ourselves to two players and assume that $v_1(x_1(s^v)) = 1$, v_2 is identically 0, and $\mu = 1$. Therefore, we aim to construct a two-player instance with (deterministic) bid vector s such that

$$\sup \left\{ u_1((s'_1, s_{-1}), v_1) + \sum_{j \in [x_1(s^v)]} \beta_{m-j+1}(s) \;\middle|\; s' \in \Sigma \right\}$$

is as small as possible. We simplify the problem further by restricting ourselves to constant marginal valuations, i.e., there is a constant c such that $\check{v}_1(j) = c$ for all j, and we assume that $x_1(s^v) = m$. Then the summation $\sum_{j \in [x_1(s^v)]} \beta_{m-j+1}(s)$ in the above equation can be regarded as a Riemann sum of a non-decreasing non-negative function f that nowhere exceeds c. We can assume without loss of generality that the non-zero bids in s' are equal (as the supremum in the above equation is certainly attained for such a bid vector) and let $\bar{z}(s')$ denote the number of non-zero bids in s'. The term $u_1((b'_1, b_{-1}), v_1)$ can in turn be interpreted as the surface of a rectangle with dimensions $c - f(\bar{z}(s'))$ and $\bar{z}(s')$.

The problem of constructing our instance therefore roughly reduces to the geometric problem of finding the right non-negative non-decreasing curve f such that the surface under the curve, plus the maximum surface of a rectangle with dimensions $c - f(j)$ and j, for $j \in [0, m]$, is minimized. The proof of Theorem 51, gives essentially a (discretized) description of this curve f, and shows subsequently that with this choice of f the supremum above is upper bounded by $(1 - 1/e)U(s^v, v)$ if we take an arbitrarily fine Riemann sum of f (i.e., the Riemann integral), which can be realized by taking m to infinity. Moreover, it coincidentally turns out that if we take for μ a value other than 1, then this particular instance remains "optimal" in the sense that there exists no λ that satisfies (2.1) and simultaneously yields a better value for $\max\{1, \mu - 1\}/\lambda$ than $e/(e-1)$.

Proof. Fix $m \in \mathbb{N}$ and $\mu \in \mathbb{R}_{\geq 0}$ arbitrarily. We construct an instance of the discriminatory price auction with 2 players, submodular valuation functions v_1, v_2, and bid vectors s, such that for every possible strategy s'_i of player $i \in [2]$ we have

$$\sum_{i=1}^{2} u_i((s'_i, s_{-i}), v_i) \leq \mu \left(1 - \frac{1}{e^{1/\mu}} + \frac{1}{m}\left(1 - \frac{1}{e}\right)\right) U(s^v, v) - \mu \sum_{j \in [m]} \beta_{m-j+1}(s).$$

By taking m to infinity, we see that for any fixed value of μ, the best price of anarchy that we can obtain using Theorem 47 is $\max\{1, \mu\}/(\mu(1 - e^{-1/\mu}))$. The latter expression is minimized by taking $\mu = 1$, and from this the claim follows.

The construction of our instance is as follows: Let the valuation functions be defined as $v_1(j) = j$ and $v_2(j) = 0$ for every $j \in [m]$. Then $x_1(s^v) = m$ and $x_2(s^v) = 0$ and the optimal social welfare is $U(s^v, v) = m$. Define the bid vector s_1 of player 1 to be the zero vector and the bid vector s_2 of player 2 as

$$b_2(j) = \begin{cases} 1 - \frac{m}{e^{1/\mu}(m-j+1)} & \text{if } j \in [\lceil m\left(1 - \frac{1}{e^{1/\mu}}\right)\rceil], \\ 0 & \text{otherwise .} \end{cases}$$

We assume that the tie-breaking rule of the auction always assigns a unit to player 1 when there is a tie. Then, if player 2 bids s_2, there is a $j \in [\lceil m\left(1 - e^{-1/\mu}\right)\rceil]$ such that player 1 maximizes its utility when it sets all its bids equal to $s_2(j)$. Let $s_1' = s_2(j)\mathbf{1}$ for some j in this range (where $\mathbf{1}$ is the m-dimensional all-ones vector). We have

$$\begin{aligned} u_1((s_1', s_2), v_1) &= v_1(m-j+1) - (m-j+1)s_2(j) \\ &= \frac{m}{e^{1/\mu}} + \mu \sum_{\ell \in [m]} s_2(j) - \mu \sum_{\ell \in [m]} \beta_j(s) \\ &\leq \frac{m}{e^{1/\mu}} + \mu \int_1^{m(1-e^{-1/\mu})+1} \left(1 - \frac{m}{e^{1/\mu}}(m-t+1)\right) dt \\ &\quad + \mu \sum_{\ell \in [\lceil m(1-\frac{1}{e^{1/\mu}})\rceil]} (s_2(\ell) - s_2(\ell+1)) - \mu \sum_{\ell \in [m]} \beta_j(s) \\ &= m\mu \left(1 - \frac{1}{e^{1/\mu}}\right) + s_2(1) - \mu \sum_{\ell \in [m]} \beta_j(s) \\ &= \mu \left(1 - \frac{1}{e^{1/\mu}} + \frac{1}{m}\left(1 - \frac{1}{e}\right)\right) U(s^v, v) - \mu \sum_{\ell in [k]} \beta_{m-j+1}(s). \end{aligned}$$

For player 2, $u_2((s_2', s_1), v_2) \leq 0$ for every bid vector s_2'. This establishes our claim for the discriminatory price auction.

□

Theorem 51 rules out the possibility of obtaining better upper bounds by means of the smoothness framework of Syrgkanis and Tardos [2013], or by means of *any* approach aiming at identifying the s_i' required by Theorem 47, including Feldman et al. [2013]. These are the only known techniques for obtaining upper bounds on the mixed Bayes-Nash price of anarchy. Thus, any improvement on our upper bound for the discriminatory price auction must use either specific structural properties of the (mixed

Bayes-Nash equilibrium) distribution $\bar{\sigma}$, or a new approach altogether. The same holds for improving the upper bounds for the uniform price auction.

2.6 Smoothness and its Implications

We elaborate in this section on the connections of our results to the smoothness framework that has been developed by Syrgkanis and Tardos [2013].

The results of Syrgkanis and Tardos [2013] concern a class of games that we call *auction mechanisms*. We define auction mechanisms here only informally: An *auction mechanism* is an incomplete information utility maximization game where the type set of each player $i \in [n]$ consists of valuation functions v_i that map some set D_i to $\mathbb{R}$, and a player's utility on a given strategy profile s is defined as $v_i(x_i(s)) - p(s)$, where x_i is a function from strategy profiles to D_i, and p is a function from strategy profiles to $\mathbb{R}$. Clearly, multi-unit auctions are auction mechanisms.

We first review the smoothness definitions introduced in Syrgkanis and Tardos [2013]. As introduced earlier, let $p_i(s)$ refer to the payment of player i under bid profile s.

Definition 52 (Syrgkanis and Tardos [2013])**.** Am auction mechanism is (λ, μ)-*smooth* for $\lambda \in \mathbb{R}_{>0}$ and $\mu\mathbb{R}_{\geq 0}$ if for any valuation profile v and for any strategy profile s there exists a bid profile s'_i for each i such that

$$\sum_{i \in [n]} u_i((s'_i, s_{-i}), v_i) \geq \lambda U(s^v, v) - \mu \sum_{i \in [n]} p_i(s).$$

In Syrgkanis and Tardos [2013], it is shown that if an auction mechanism is (λ, μ)-smooth, then several results related to the price of anarchy follow automatically. One such result concerns an upper bound on the price of anarchy. Another result is that the smoothness property (and therefore the upper bound on the price of anarchy) is roughly retained under *simultaneous and sequential compositions*. In these compositions there is a finite number of auction mechanisms with separate allocation and payment rules. Every player specifies for each auction mechanism a strategy. In a simultaneous composition, these profiles are submitted simultaneously, while in the sequential composition, they are submitted sequentially, i.e., a strategy submitted to an auction mechanism M may depend on the history of strategy profiles submitted to auction mechanisms that occur before M in this sequence. A player expresses its valuation for the m-tuples of outcomes of the auction mechanisms in a restricted way: in the simultaneous composition it is assumed that the valuation function of each player is *fractionally subadditive* across the auction mechanisms (see Syrgkanis and Tardos

[2013] for a definition). In the sequential composition, the valuation function of each player is defined as the maximum of its valuations over these auction mechanisms. We summarize the main composition results of Syrgkanis and Tardos [2013] in the theorem below.

Theorem 53 (Theorems 4.2, 4.3, 5.1, and 5.2 in Syrgkanis and Tardos [2013]). *Let $\lambda \in \mathbb{R}_{>0}$ and $\mu \in \mathbb{R}_{\geq 0}$.*

1. *If an auction mechanism is (λ, μ)-smooth, then the correlated price of anarchy and mixed Bayes-Nash price of anarchy of that auction mechanism is at most $\max\{1, \mu\}/\lambda$.*

2. *Simultaneous compositions of (λ, μ)-smooth auction mechanisms are (λ, μ)-smooth auction mechanisms.*

3. *Sequential compositions of (λ, μ)-smooth auction mechanisms are $(\lambda, \mu + 1)$-smooth auction mechanisms.*

By exploiting our Key Lemma, we can show that the discriminatory price auction is smooth.

Theorem 54. *Let $\alpha \in \mathbb{R}_{\geq 0}$. Every discriminatory auction is (λ, μ)-smooth (under both the standard and uniform bidding interface) with*

1. *$(\lambda, \mu) = (\alpha\left(1 - e^{-1/\alpha}\right), \alpha)$, when the players have submodular valuation functions, and*

2. *$(\lambda, \mu) = ((\alpha/2)\left(1 - e^{-1/\alpha}\right), \alpha)$ when the players have subadditive valuation functions.*

Proof. Note that for the discriminatory price auction we have

$$\sum_{i \in [n]} p_i(s) = \sum_{i \in [n]} \sum_{j \in [x_i(s)]} s_i(j) = \sum_{j \in [m]} \beta_{m-j+1}(s).$$

Note that Lemma 48 holds for every player $i \in [n]$ with $\lambda \in \mathbb{R}_{>0}$ and $\mu \in \mathbb{R}_{\geq 0}$ as stated in Theorem 54. By invoking Lemma 48 and summing (2.3) over all players, we obtain

$$\sum_{i \in [n]} \mathbf{E}_{s_i \sim \bar{\sigma}'_i(v_i)}[u_i((s'_i, s_{-i}), v_i)]$$
$$\geq \lambda \sum_{i \in [n]} v_i(x_i(s^v)) - \mu \sum_{i \in [n]} \sum_{j \in [x_i(s^v)]} \beta_{m-j+1}(s_{-i})$$

	discriminatory auction		**uniform price auction**	
Valuation Functions	*Simultaneous*	*Sequential*	*Simultaneous*	*Sequential*
Submodular	$\frac{e}{(e-1)}$	$\frac{2e}{(e-1)}$	$\left\|W_{-1}\left(\frac{-1}{e^2}\right)\right\| \approx 3.1462$	
Subadditive	$\frac{2e}{(e-1)}$	$\frac{4e}{(e-1)}$	$2\left\|W_{-1}\left(\frac{-1}{e^2}\right)\right\| \approx 6.2924$	

Table 2.2: Upper bounds on the correlated price of anarchy and mixed Bayes-Nash price of anarchy for compositions of discriminatory and uniform price auctions.

$$\geq \lambda U(s^v, v) - \mu \sum_{j\in[m]} \beta_{m-j+1}(s),$$

where the last inequality holds because for every player $i \in [n]$, $\beta_{m-j+1}(s_{-i}) \leq \beta_{m-j+1}(s)$ for every $j \in [m]$, $\sum_{i\in[n]} x_i(s) = m$, and $\beta_{m-j+1}(s)$ is non-decreasing in j. $\square$

Theorem 54 in combination with Theorem 53 leads to the composition results stated in the left column of Table 2.2 (these bounds are achieved for $\alpha = 1$).

For auction mechanisms where one imposes no-overbidding, a different smoothness notion is introduced in Syrgkanis and Tardos [2013]. Given a mechanism, element $j \in D_i$ and strategy s_i for a player $i \in [n]$, define player i's *willingness-to-pay* $B_i(s_i, j)$ as the maximum payment it could ever pay conditional to that the outcome for player i is j, i.e., $B_i(s_i, j) = \sup\{p_i(s_i, s_{-i}) \mid x_i(s_i, s_{-i}) = j\}$. For auction mechanisms, a mixed Bayesian strategy σ_i for a player i is said to satisfy no-overbidding if for all $i \in [n]$ and $v_i \in V_i$ it holds that $\mathbf{E}_{s\sim\sigma_i(v_i)}[B_i(s_i, x_i(s))] \leq \mathbf{E}_{s\sim\sigma}[v_i(x_i(s))]$. This definition of no-overbidding specializes to for multi-unit auctions to the definition of no-overbidding that we provided in Section 2.3. The *no-overbidding mixed Bayes-Nash price of anarchy* and *no-overbidding correlated price of anarchy* are defined in the same way as for the multi-unit auction case, but now with this more general of no-overbidding taken into account. The following definition is taken from Syrgkanis and Tardos [2013]:

Definition 55 (Syrgkanis and Tardos [2013]). An auction mechanism is said to be *weakly* (λ, μ_1, μ_2)-smooth for $\lambda \in \mathbb{R}_{>0}$ and $\mu_1, \mu_2 \in \mathbb{R}_{\geq 0}$ iff for any valuation profile

v and for any bid profile s there exists for each player i a bid profile s'_i such that

$$\sum_{i\in[n]} u_i((s'_i, s_{-i}), v_i) \geq \lambda U(s^v, v) - \mu_1 \sum_{i\in[n]} p_i(s) - \mu_2 \sum_{i\in[n]} B_i(s_i, x_i(s)).$$

Syrgkanis and Tardos [2013] establish the following results.

Theorem 56 (Theorems 7.4, C.4 and C.5 in Syrgkanis and Tardos [2013])**.** *Let $\lambda \in \mathbb{R}_{>0}$ and $\mu_1, \mu_2 \in \mathbb{R}_{\geq 0}$.*

1. *If an auction mechanism is weakly (λ, μ_1, μ_2)-smooth, then the no-overbidding correlated price of anarchy and the no-overbidding mixed Bayes-Nash price of anarchy is at most $(\mu_2 + \max\{1, \mu_1\})/\lambda$.*

2. *Simultaneous compositions of weakly (λ, μ_1, μ_2)-smooth auction mechanisms are weakly (λ, μ_1, μ_2)-smooth mechanisms.*

3. *Sequential compositions of weakly (λ, μ_1, μ_2)-smooth auction mechanisms are weakly $(\lambda, \mu_1 + 1, \mu_2)$-smooth mechanisms.*

Using our Key Lemma, we can show that the uniform price auction is weakly smooth.

Theorem 57. *Every uniform price auction is weakly (λ, μ_1, μ_2)-smooth (both under the standard and uniform bidding interface) with*

1. $(\lambda, \mu_1, \mu_2) = (\alpha\left(1 - e^{-1/\alpha}\right), 0, \alpha)$ *when the players have submodular valuation functions, and*

2. $(\lambda, \mu_1, \mu_2) = ((\alpha/2)\left(1 - e^{-1/\alpha}\right), 0, \alpha)$ *when the players have subadditive valuation functions.*

The following transformation will be helpful in the poof of Theorem 57. Let s be an arbitrary bid profile. We derive a uniform bid profile $\bar{s}$ from s as follows: Let $c_i = s_i(x_i(s))$. For ease of notation, we adopt the convention that $s_i(0) = 0$. Define $\bar{s}_i$ as the vector that is c_i on the first $x_i(s)$ entries and zero everywhere else. Clearly, $\bar{s}$ is a uniform bid profile.

Lemma 58. *Let $\bar{s}$ be the uniform bid profile derived from s as described above. Then for the uniform price auction, the following holds for every player $i \in [n]$:*

1. $x_i(\bar{s}) = x_i(s)$*;*

2. $B_i(\bar{s}_i, x_i(\bar{s})) = B_i(s_i, x_i(s))$.

Proof. Let β_m and β_{m+1} refer to the mth and $(m+1)$th element in the non-increasingly ordered vector of all bids of s, respectively. Note that all of the highest m bids among all bids of s are at least β_m. Also, all of the remaining bids of s are at most β_{m+1}. We conclude that a bid is among the highest m bids of $\bar{s}$ if and only if it is among the highest m bids under s, which proves the first point of the claim.

Observe that $B_i(\bar{s}_i, x_i(\bar{s})) = x_i(\bar{s})\bar{s}_i(x_i(\bar{s}))$. Now, using the first point of the claim (which we just proved), and the definition of $\bar{s}_i$, we obtain:

$$\begin{aligned} B_i(\bar{s}_i, x_i(\bar{s})) &= x_i(\bar{s})\bar{s}_i(x_i(\bar{s})) \\ &= x_i(s)\bar{s}_i(x_i(s)) \\ &= x_i(s)s_i(x_i(s)) \\ &= B_i(s_i, x_i(s)), \end{aligned}$$

which shows the second point of the claim. □

Proof of Theorem 57. We first prove weak smoothness for the uniform bidding interface and then extend this result to the standard bidding interface via a coupling argument.

For every uniform price auction it holds that $B_i(s_i, j) = jb_i(j)$ for all $i \in [n], j \in [m], s_i \in \Sigma_i$. As before, exploiting Lemma 48 and summing inequality (2.3) over all players, we obtain that for every strategy profile s and valuation profile v there is a randomized uniform bid strategy $\bar{\sigma}'_i(v_i)$ such that

$$\sum_{i\in[n]} \mathbf{E}_{s'_i\sim\bar{\sigma}'_i(v_i)}[u_i((s'_i, s_{-i}), v_i)] \quad \geq \lambda U(s^v, v) - \mu \textstyle\sum_{j\in[m]} \beta_{m-j+1}(s).$$

If s is a uniform bid profile then the claim follows because

$$\sum_{j\in[m]} \beta_{m-j+1}(s) = \sum_{i\in[n]} x_i(s)s_i(x_i(s)) = \sum_{i\in[n]} B_i(s_i, x_i(s)). \tag{2.11}$$

Note that for the standard bidding interface, the first equality here would be false, because we can only infer that $\sum_{j\in[m]} \beta_{m-j+1}(s) \geq \sum_{i\in[n]} x_i(s)s_i(x_i(s))$. However, the following work-around establishes weak smoothness for the uniform price auction and the standard bidding interface.

Note that in general $p_i(\bar{s}) \neq p_i(s)$, because all losing bids in $\bar{s}$ are zero. However, the above two properties turn out to be sufficient to prove weak smoothness in the standard bidding interface: Let $\bar{s}$ be the uniform bid profile that we obtain from s as described above. Applying Lemma 48 to the uniform bid profile $\bar{s}$ and pricing rule $p_i(\bar{s}) = B_i(\bar{s}_i, x_i(\bar{s}))$ (which is discriminatory price dominated), we conclude that for every player $i \in [n]$, for every strategy profile $s \in \Sigma$, and for every valuation profile $v \in V$, there exists a no-overbidding randomized uniform bid strategy $\bar{\sigma}'_i(v_i)$ such that

$$\begin{aligned}&\mathbf{E}_{s'_i \sim \bar{\sigma}'_i(v_i)}[v_i(x_i(s'_i, \bar{s}_{-i})) - B_i(s'_i, x_i(s'_i, \bar{s}_{-i}))] \\ &\quad \geq \lambda v_i(x_i(s^v)) - \mu \sum_{j \in [x_i(s^v)]} \beta_{m-j+1}(\bar{s}_{-i}). \end{aligned} \tag{2.12}$$

Note that by Lemma 58:

$$\begin{aligned} u_i(s'_i, s_{-i}), v_i) &= v_i(x_i(s'_i, s_{-i})) - p_i(s'_i, s_{-i}) \\ &\geq v_i(x_i(s'_i, s_{-i})) - B_i(s'_i, x_i(s'_i, s_{-i})) \\ &= v_i(x_i(s'_i, \bar{s}_{-i})) - B_i(s'_i, x_i(s'_i, \bar{s}_{-i})). \end{aligned} \tag{2.13}$$

By summing (2.12) over all players and using (2.13), we obtain:

$$\begin{aligned} \sum_{i \in [n]} \mathbf{E}_{s'_i \sim \bar{\sigma}'_i(v_i)}[u_i((s'_i, s_{-i}), v_i)] &\geq \lambda U(s^v, v) - \mu \sum_{j \in [m]} \beta_{m-j+1}(\bar{s}) \\ &= \lambda U(s^v, v) - \mu \sum_{i \in [n]} B_i(\bar{s}_i, x_i(\bar{s})) \\ &= \lambda U(s^v, v) - \mu \sum_{i \in [n]} B_i(s_i, x_i(s)), \end{aligned}$$

where the first equality follows from (2.11) because $\bar{s}$ is a uniform bid profile and the second equality holds because of Lemma 58. □

As a consequence, we obtain the composition results stated in right column of Table 2.2. (These bounds are achieved by choosing $\alpha = -1/(W_{-1}(-1/e^2) + 2) \approx 0.87$).

2.7 Conclusions

We derived inefficiency upper bounds in the incomplete information model for the widely popular discriminatory and uniform price auctions, when players have submodular or subadditive valuation functions. Notably, our bounds for subadditive valuation functions already improve upon the ones that were known for players with submodular valuation functions [Markakis and Telelis, 2012, Syrgkanis and Tardos, 2013]. Moreover, for each of the two formats and valuation function classes we considered both the *standard* bidding interface [Krishna, 2002, Milgrom, 2004] and a practically motivated *uniform* bidding interface. To derive our results, we elaborated on several techniques from the literature on *simultaneous auctions* [Syrgkanis and Tardos, 2013, Feldman et al., 2013, Christodoulou et al., 2008, Bhawalkar and Roughgarden, 2011]. By the developments of Syrgkanis and Tardos [2013], our bounds for players with submodular valuation functions yield improved inefficiency bounds for *simultaneous* and *sequential* compositions of the considered formats. In absence of an indicative lower bound in the incomplete information model, we showed that our upper bound of $e/(e-1)$ for the discriminatory auction with submodular valuation functions is the best possible with respect to the currently known proof techniques. Additionally, for the uniform price auction (with players with submodular valuation functions), we showed that, proving an upper bound of less than 2, also requires novel techniques; this poses a particularly challenging problem, given the lower bound of $e/(e-1)$ from Markakis and Telelis [2012].

Chapter 3

The Strong Price of Anarchy of Linear Bottleneck Congestion Games*

We study in this chapter the inefficiency of *strong equilibria* of *bottleneck congestion games*. Strong equilibria are strategy profiles for which it is not possible for a *coalition* of players to deviate to alternative strategies such that every player in the coalition decreases its cost (or increases its utility, depending on whether the game considered is a cost minimization game or a utility maximization game). Strong equilibria were first introduced as a solution concept by Aumann [1960]. Bottleneck congestion games are a variation on congestion games, where the cost of a player is the maximum delay among the facilities that it chooses, instead of the sum of the delays. These games therefore model situations in which the performance of the slowest utility is the crucial factor, so that the players want to maximize this slowest performance. Bottleneck congestion games are thus cost minimization games. With respect to a social cost function $C : \Sigma \to \mathbb{R}$, the ratio of the worst strong equilibrium and the minimum in the image of C is referred to as the *strong price of anarchy*, analogous to the price of anarchy and its variations that we introduced in Section 1.3.1.8. Strong equilibria are known to exist in bottleneck congestion games Harks et al. [2009], and this motivates the study of the strong price of anarchy of these games, which is the subject of the present chapter. Our choice of social cost function will be the one that takes the maximum cost among all

*The contents of this chapter have been published as De Keijzer et al. [2010].

the players, i.e., (1.5) in Chapter 1.

The material from Chapter 1 that is relevant to this chapter, consists of Section 1.3.1 up to Section 1.3.1.1, and Sections 1.3.1.7, 1.3.1.8, and 1.3.1.10. We proceed by first defining formally the essential concepts just mentioned.

Definition 59 ((Standard, linear, (a)symmetric, network) bottleneck congestion game)**.** A *bottleneck congestion game* is a full information cost minimization game (n, Σ, c) for which there exists an $m \in \mathbb{N}_{>0}$ such that $\Sigma_i \subseteq 2^{[m]}$. Moreover, for each $j \in [m]$ there exists a function $d_j : 2^{[n]} \to \mathbb{R}_{\geq 0}$, and for each player $i \in [n]$, its cost function c_i is given by $c_i(s) = \max_{j \in \Sigma_i} d_j(P_j(s))$ where $P_j(s) = \{i' \in [n] \mid j \in s_{i'}\}$ is the *set of players i choosing j in their strategy s_i*. In the context of a congestion game, the elements of $[m]$ are called *facilities*, and d_j is referred to as the *delay function* of facility $j \in [m]$. A bottleneck congestion game may thus be represented as a quadruple (n, m, Σ, d) where $d = (d_1, \ldots, d_m)$ is the vector of delay functions.

The following are some properties of relevance, that a bottleneck congestion game may possess.

- A bottleneck congestion game is called *monotone* if $d_e(P) \geq d_e(P')$ for all $P, P' \subseteq [n]$ with $P \supseteq P'$.

- A bottleneck congestion game is called *linear* if for each $j \in [m]$ there exists $a_j \in \mathbb{R}_{\geq 1}$ and for each $i \in [n]$ there exists $w_i \in \mathbb{R}_{\geq 1}$ such that $d_j(P) = a_j \sum_{i \in P} w_i$ for all $j \in [m]$ and $P \subseteq [n]$. The number w_i is in this context referred to as the *weight* of player i, for $i \in [n]$.

- A bottleneck congestion game is called *symmetric* if $\Sigma_i = \Sigma_j$ for all $i, j \in [n]$.

- A bottleneck congestion game is a *network bottleneck congestion game* if there is a directed graph $G = (V, E)$, a bijection $b : E \to [m]$, and two vertices v_i, w_i for all $i \in [n]$ such that $\Sigma_i = \{\{b(e) \mid e \in S\} \mid S \text{ is a } (v_i, w_i)\text{-path in } G\}$ for all $i \in [n]$.

Note that for linear bottleneck congestion games, the assumption that $a_j, w_i \geq 1$ for $i \in [n], j \in [m]$, is without loss of generality, as we can always enforce them by scaling the weights and coefficients appropriately in case one of these numbers lies in $(0, 1]$.

Moreover, note the strongly similar definition of congestion games given in Section 1.3.1.10. Bottleneck congestion games differ from congestion games in two aspects: First, a player's cost is now the maximum delay among the facilities, instead of the sum. Secondly, the delay of a facility is a function of the player set that chooses it, instead of only the cardinality of the player set that chooses it.

Definition 60 (Notation for deviations by coalitions)**.** Given a game $\Gamma = (n, \Sigma, c)$, a strategy profile $s \in \Sigma$, and a player set $P \subseteq [n]$, we define $\Sigma_P = \times_{i \in P} \Sigma_i$. For a strategy vector $s'_P \in \Sigma_P$, we denote by (s'_P, s_{-P}) the strategy profile obtained from s by replacing for all $i \in P$ the strategy s_i by the strategy of i in the vector s'_P. Just as our notation for single-player deviations, this notation overloads standard notation and is formally ambiguous. However, it will always be clear from context what is meant, and no confusion will arise.

Since we study strong equilibria only for bottleneck congestion games, which are cost minimization games, we define the strong equilibria and strong price of anarchy only for cost minimization games.

Definition 61 ((k-)strong equilibrium for cost minimization games)**.** A *k-strong equilibrium* for a cost minimization game $\Gamma = (n, \Sigma, c)$ and a number $k \in [n]$ is a strategy profile $s \in \Sigma$ such that for all $P \subseteq [n], |P| \leq k$ and for all $s'_P \in \Sigma_P$, there exists a player $i \in P$ such that $c_i(s'_P, s_{-P}) \geq c_i(s)$. A *strong equilibrium* is an n-strong equilibrium. We denote the set of strong equilibria of a game Γ by SE_Γ, and the set of k-strong equilibria, for $k \in \mathbb{N}_{>0}$, by SE^k_Γ. The subscript Γ may be omitted in case it is clear from the context.

The intuition behind the above definition is that in a k-strong equilibrium, no player set of at most k players will form a coalition in order to deviate in a coordinated way to alternative strategies: there will always be a player in such a set that will not improve its cost by switching, so this player will not be willing to cooperate.

Observe that a 1-strong equilibrium of a game is a pure equilibrium. Clearly, we have for $k, \ell \in \mathbb{N}_{>0}, k < \ell$, that $SE^k_\Gamma \supseteq SE^\ell_\Gamma$ for all games Γ.

Definition 62 ((k-)strong price of anarchy for cost minimization games)**.** The *k-strong price of anarchy* for a cost minimization game $\Gamma = (n, \Sigma, c)$ is the price of anarchy of cost minimization game Γ for the set SE^k_Γ (see Definition 17). The *strong price of anarchy* is the n-strong price of anarchy.

The *k-strong price of stability* for a cost minimization game $\Gamma = (n, \Sigma, c)$ is the price of stability of cost minimization game Γ for the set SE^k_Γ (see Definition 17). The *strong price of stability* is the n-strong price of stability.

For a class of games $\mathcal{G}$ the *k-strong price of anarchy of $\mathcal{G}$* and *k-strong price of stability of $\mathcal{G}$* are defined as the supremum of respectively the k-strong price of anarchy and k-strong price of stability of the games in $\mathcal{G}$. The *strong price of anarchy of $\mathcal{G}$ and strong price of stability of $\mathcal{G}$* are the n-strong price of anarchy and n-strong price of stability, respectively.

Theorem 63 (Harks et al. [2009])**.** *$SE_\Gamma \neq \varnothing$ if Γ is a monotone bottleneck congestion game.*

3.1 Background

A well-studied subclass of bottleneck congestion games are *load balancing games*. These are linear bottleneck congestion games where the strategy set for each player consists of only singleton sets.

The latter class of games capture many applications of practical relevance. They model situations in which a set of strategically acting players (or jobs) compete for a limited number of facilities (or machines). Every player chooses one of the resources available to it and assigns its weight (or load) to this resource. The delay of a resource depends on the total weight of the players using it. The goal of each player is to select a resource such that the delay that it experiences on this resource is minimized.

As load balancing games model situations where there is a set of autonomous players that utilize distributed processors upon which a system is built, the need for studying the price of anarchy of these games is clear. The social cost objective of an assignment of loads to processors is usually measured by the the completion time of the most loaded machine, i.e., (1.5). Load balancing games were studied extensively for a variety of different *machine environments*, including

- *identical machines* [Koutsoupias and Papadimitriou, 1999], where the delay functions on the facilities are identical.
- *uniformly related machines* [Czumaj and Vöcking, 2007, Gairing et al., 2006, Koutsoupias et al., 2003, Koutsoupias and Papadimitriou, 1999], where each facility j has an associated number $a_j \in \mathbb{R}_{\geq 0}$ and the delay of a facility is the sum of the weights of the players that have j in their strategy set.
- *restricted assignment* [Awerbuch et al., 2006, Gairing et al., 2006], where the case is allowed that $\Sigma_i \subset [m]$ for $i \in [n]$ (as opposed to *unrestricted assignment*, where for all $i \in [n]$ it holds that $\Sigma_i = [m]$.
- *unrelated machines* [Andelman et al., 2007], which are actually a generalization of the load balancing games that we just defined them. These load balancing games are (non-linear) bottleneck congestion games where the strategy set of each player consists only of singleton sets, such that there is a given number $a_{i,j}$ for all $i \in [n], j \in [m]$ such that $d_j(P) = \sum_{i \in P} a_{i,j}$ for all $j \in [m], P \subseteq [n]$.

Clearly, a natural extension of load balancing games are the *bottleneck congestion games*. They generalize load balancing games by allowing a player to choose among sets of resources, instead of just single resources. Also, in bottleneck congestion games, the delay functions may be more general. This generalization brings the model closer to practice, as in most large scale computing systems the workload of a player occupies different components of a system simultaneously. For example, instances of such games emerge if the components form paths in networks, or if they correspond to parallel processors. It is often natural to assume that each player wants to balance its load across the different components available to it and hence attempts to minimize the maximum delay of a facility that it uses.

A more restrictive definition of network bottleneck congestion games was considered first by Banner and Orda [2007]. The authors showed existence of pure equilibria and provided a $\Theta(m)$ bound on the price of anarchy for facilities with identical delay functions. Busch and Magdon-Ismail studied in Busch and Magdon-Ismail [2009] the (non-strong) price of anarchy of network bottleneck congestion games for identically weighted players. Harks et al. [2009] introduced the general bottleneck congestion games (i.e., the variant that we defined above) and showed that strong equilibria are guaranteed to exist in these games when the delay functions are monotone. In Harks et al. [2010], the problem of computation of pure and strong equilibria is addressed. The paper shows several hardness results and proposes polynomial time algorithms for special cases. In a work of Busch and Kannan [2012], the (non-strong) price of anarchy of bottleneck congestion games is analyzed under the assumption that the player has bounded *stretch*. Stretch is a measure of variation in the resource utilization in the strategy sets of the players.

Load balancing games were first studied by Koutsoupias and Papadimitriou [1999]. Among other results, the authors provided a lower bound on the mixed price of anarchy for the case of identical machines. Koutsoupias et al. [2003] and, independently Czumaj and Vöcking [2007], proved a matching upper bound. They also proved that the pure price of anarchy is $\Theta(\log m/\log\log m)$. The same bound on the price of anarchy was shown by Awerbuch et al. [2006] for the case of identical machines and restricted assignment. Gairing et al. [2006] obtained independently the same bounds and proved that the price of anarchy is at least $m-1$ and at most m for the case of restricted assignment and uniformly related machines.

Andelman et al. [2007] were the first to study strong and k-strong equilibria in the context of load balancing games. They proved that the strong price of anarchy lies between m and $2m-1$ for the case of unrelated machines, which was tightened to exactly m by Fiat et al. [2007]. In the latter work it was also shown that for uniformly related machines, the strong price of anarchy is $\Theta(\log m/(\log\log m)^2)$. For results in the context of more general scheduling games and associated scheduling policies (termed *coordination mechanisms*), the interested reader is referred to Immorlica et al. [2009] and the references therein.

Bottleneck congestion games owe their name to their similarity to *congestion games* (see Section 1.3.1.10). The pure price of anarchy was derived by Christodoulou and Koutsoupias [2005a], and independently by Awerbuch et al. [2005]. It is shown in

Christodoulou and Koutsoupias [2005a] that the pure price of anarchy is in $\Theta(\sqrt{n})$ for linear congestion games with respect to the social cost function being the maximum over the players' cost. Christodoulou and Koutsoupias [2005a] also shows that the price of anarchy is $5/2$ for linear congestion games with respect to the social cost function being the sum of the players' costs. Bounds on the pure price of anarchy for polynomial delay functions were also derived in Christodoulou and Koutsoupias [2005a]. Exact bounds for polynomial delay functions and also for weighted players were developed in Aland et al. [2006].

Related to the topic of coalition formation in congestion games is the work in Hayrapetyan et al. [2006], where the impact of "collusion" in network congestion games is studied, where players form coalitions to minimize their collective cost. These coalitions are assumed to be formed exogeneously, i.e., conceptually, each coalition is replaced by a "super-player" that acts on behalf of its members. The authors show that collusion in network congestion games can lead to Nash equilibria that are inferior to the ones of the collusion-free game (in terms of social cost). They also derive bounds on the the price of anarchy caused by collusion.

3.2 Contributions and Outline

We study the price of anarchy and strong price of anarchy of bottleneck congestion games. Our choice of social cost function C is the maximum among the player costs, i.e., (1.5) in Chapter 1. We restrict our study to *linear* bottleneck congestion games, which still constitute a rich class of games. For example, they include as a special case load balancing games with identical or uniformly related machines with or without restricted strategy sets. We provide upper and lower bounds on the (strong) price of anarchy for symmetric and asymmetric linear bottleneck congestion games.

1. In Section 3.3, we show that both the price of anarchy and the strong price of anarchy of linear bottleneck congestion games is $\Theta(m)$. More precisely, we show that the price of anarchy is between m and $2m - 1$ and the strong price of anarchy is between $m - 1$ and m.

2. We derive better bounds for games in which all players have identical weights, in Section 3.3.2. We prove that the strong price of anarchy equals 2 for symmetric linear network bottleneck congestion games and is at most $O(\sqrt{n})$ and $O(\sqrt{mC(s^*)})$ for asymmetric linear bottleneck congestion games, where $C(s^*)$ denotes the minimum social cost.

	general	identical players	identical facilities
symmetric	$O(m)$	2	2
asymmetric	$\Theta(m)$	$O(\sqrt{n}), O(\sqrt{mC(s^*)})$	$\Theta(\sqrt{(m)})$

Table 3.1: Summary of the bounds obtained for the strong price of anarchy and the k-strong price of anarchy of linear bottleneck congestion games. The left column displays the bounds on the strong price of anarchy for the general case. The middle column displays our bounds on the strong price of anarchy for the special case that the players have identical weights. The right column displays our bounds on the strong price of anarchy for the special case that the facilities have identical delay functions. The top row displays the bounds for the symmetric case, and the bottom row displays the bounds for the asymmetric case.

3. We consider in Section 3.4 the special case where all facilities have identical linear delay functions, and show that the strong price of anarchy is $\Theta(\sqrt{m})$.

Most of our results are summarized in Table 3.1. A result not displayed there is that the price of anarchy of linear bottleneck congestion games is at most $2m-1$ and there is an asymptotically matching lower bound showing that the strong price of anarchy is at least $m-1$. Moreover, we remark that the asymptotically tight lower bounds that we provide are all for the class of linear network bottleneck congestion games.

3.3 Arbitrary Facilities

In this section, we derive bounds on the price of anarchy and strong price of anarchy of linear bottleneck congestion games. We consider both the case of general player weights and identical player weights.

3.3.1 General Player Weights

We first consider the most general case of arbitrary linear delay functions with player weights that may be distinct. We show that the price of anarchy is at most $2m-1$ in this case. We obtain a better bound of m on the strong price of anarchy and present an almost tight lower bound.

Theorem 64. *The price of anarchy of linear bottleneck congestion games is at most $2m-1$ and at least m.*

Proof. Let s be a pure equilibrium with cost $C(s) = \alpha C(s^*)$ for some $\alpha \geq 1$ and let s^* be a strategy profile that minimizes C. We prove by induction that for every integer κ such that $1 \leq \kappa < (\alpha+3)/2$, there is a set $S_\kappa \subseteq [m]$, $|S_\kappa| = \kappa$, such that for every $j \in S_\kappa$, $d_j(P_j(s)) \geq (\alpha - \kappa + 1)C(s^*)$.

The claim holds for $\kappa = 1$ because there must exist a facility $j \in [m]$ with delay $d_j(P_j(s)) = \alpha C(s^*)$. Suppose that the induction hypothesis holds for $\kappa < (\alpha+1)/2$. We will prove that there exists a set $S_{\kappa+1} \subseteq [m]$, $|S_{\kappa+1}| = \kappa+1$ such that $d_j(P_j(s)) \geq (\alpha - \kappa)C(s^*)$ for every $j \in S_{\kappa+1}$.

Choose from S_κ a facility $\hat{j}$ with minimal a_j, i.e., $\hat{j} = \arg_j \min\{a_j \mid j \in S_\kappa\}$. By the induction hypothesis, we have

$$d_{\hat{j}}(P_{\hat{j}}(s)) \geq (\alpha - \kappa + 1)C(s^*) > \kappa C(s^*).$$

Note that

$$\sum_{i \in P_{\hat{j}}(s)} w_i = \frac{d_{\hat{j}}(P_{\hat{j}(s)}(s))}{a_{\hat{j}}} > \frac{\kappa C(s^*)}{a_{\hat{j}}}.$$

Consider the strategies that the players in $P_{\hat{j}}(s)$ choose under s^* and suppose for the sake of a contradiction that for every $i \in P_{\hat{j}}(s)$, $s_i^* \cap S_\kappa \neq \varnothing$. Then there is a facility $j \in S_\kappa$ with

$$\sum_{i \in P_j(s^*)} w_i \geq \frac{\sum_{i \in P_{\hat{j}}(s)} w_i}{\kappa} > \frac{C(s^*)}{a_{\hat{j}}}.$$

By the choice of $\hat{j}$, we have

$$d_j(P_j(s^*)) = a_j \sum_{i \in P_j(s^*)} w_i > C(s^*),$$

which is in contradiction with the definition of C. Thus, there is a player $i \in P_{\hat{j}}(s)$ that chooses a strategy s_j^* that is disjoint from S_κ. Note that for every $j \in s_i^*$ we have $a_j w_i \leq C(s^*)$. Since s is a pure equilibrium, player i cannot decrease its cost by deviating to s_i^* and thus there is some facility $j' \in s_i^*$ such that:

$$\begin{aligned} d_{j'}(P_{j'}(s)) &= (d_{j'}(s) + a_{j'} w_i) - a_{j'} w_i \\ &\geq c_i(s) - a_{j'} w_i \\ &\geq d_{\hat{j}}(P_j(s)) - C(s^*) \\ &\geq (\alpha - \kappa)C(s^*). \end{aligned}$$

The inductive step follows by setting $S_{\kappa+1} = S_\kappa \cup \{j'\}$. By choosing $\kappa = \lceil(\alpha + 1)/2\rceil < (\alpha+3)/2$, we obtain that there is a set $S_\kappa \subseteq [m]$ with $|S_\kappa| = \kappa$ and thus

$m \geq |S_\kappa| = \kappa \geq (\alpha + 1)/2$. We conclude that $C(s)/C(s^*) = \alpha \leq 2m - 1$, and therefore the price of anarchy is at most $2m - 1$.

The following instance shows that the price of anarchy is at least m, even for symmetric bottleneck congestion games with identical facilities and identical players. Consider a bottleneck congestion game with $m = n$. Every player $i \in [n]$ has unit weight $w_i = 1$ and the delay $d_j(P_j(s))$ of every $j \in [m]$ is given by $d_j(P_j(s)) = \sum_{i \in P_j(s)} w_i$ for all $s \in \Sigma$. Suppose that each player $i \in [n]$ has strategy set $\Sigma_i = 2^{[m]}$. If every player chooses a distinct facility we obtain an optimal strategy profile s^* with $C(s^*) = 1$. On the other hand, consider the strategy profile s in which every player chooses $[m]$ as its strategy. This is a pure equilibrium of cost $C(s) = m$. □

We derive a better upper bound on the strong price of anarchy for linear bottleneck congestion games. The following key lemma will be used several times in this chapter.

Lemma 65. *Let s be a strong equilibrium, let s^* be a strategy profile minimizing C, let $\lambda \in \mathbb{R}_{\geq 1}$ and let $P_\lambda \subseteq [n], P \neq \varnothing$ be such that for every $i \in P_\lambda$ we have $c_i(s) \geq \lambda C(s^*)$. Then, the following two statements hold.*

1. *There is a player $i \in P_\lambda$ and a facility $j \in s_i^*$ such that $d_j(\{i \in [n] \setminus P_\lambda \mid j \in s_i\}) \geq (\lambda - 1)C(s^*)$.*
2. *Suppose that P_λ is maximal, i.e., there is no $i \in [n] \setminus P_\lambda$ such that $c_i(s) \geq \lambda C(s^*)$. Then there is a set of players $Q_\lambda \subseteq [n] \setminus P_\lambda$ and a $j \in [m]$ such that $j \in s_i$ for all $i \in Q_\lambda$. Moreover, it holds that $a_j \sum_{i \in Q_\lambda} w_i = d_j(Q_\lambda) \geq (\lambda - 1)C(s^*)$, and $\sum_{i \in Q_\lambda} w_i \geq \lambda - 1$, and for every $i \in Q_\lambda$ we have $(\lambda - 1)C(s^*) \leq c_i(s) < \lambda C(s^*)$.*

Proof. We first prove the first claim. Note that for every player $i \in P_\lambda$ and every $j \in s_i^*$ we have

$$d_j(\{i \in P_\lambda \mid s_i^* = j\}) \leq d_j(P_j(s^*)) \leq C(s^*). \tag{3.1}$$

Suppose for the sake of a contradiction that for every player $i \in P_\lambda$ and for every $j \in S_i^*$ it holds that $d_j(\{i \in [n] \setminus P_\lambda \mid j \in s_i\}) < (\lambda - 1)C(s^*)$. Consider the strategy profile $s' = (s^*_{P_\lambda}, s_{-P_\lambda})$ in which the players in P_λ deviate to their optimal strategies in s^*. Using (3.1), we obtain for every $i' \in P_\lambda$ and for every $j \in s_i^*$:

$$\begin{aligned} d_j(s') &= d_j(\{i \in P_\lambda \mid s_i^* = j\}) + d_j(\{i \in [n] \setminus P_\lambda \mid j \in s_i\}) \\ &< C(s^*) + (\lambda - 1)C(s^*) \end{aligned}$$

$$= \quad \lambda C(s^*). \tag{3.2}$$

Thus, for every $i \in P_\lambda$, $c_i(s') = \max_{j \in s_i^*} d_j(s') < \lambda C(s^*)$, which is in contradiction with s being a strong equilibrium.

We next prove the second part of the lemma. Let $i \in P_\lambda$ be a player and $j \in s_i^*$ be a facility satisfying $d_j(\{i \in [n] \setminus P_\lambda \mid j \in s_i\}) \geq (\lambda - 1)C(s^*)$. Define Q_λ as the set of players that choose j under s but are not in P_λ, i.e., $Q_\lambda = P_j(s) \setminus P_\lambda \subseteq [n] \setminus P_\lambda$. We have

$$a_j \sum_{i \in Q_\lambda} w_i = d_j(Q_\lambda) \geq (\lambda - 1)C(s^*). \tag{3.3}$$

Since $j \in s_i^*$ and $w_i \geq 1$ for every $i \in [n]$, we have $a_j \leq C(s^*)$. Thus, $\sum_{i \in Q_\lambda} w_i \geq \lambda - 1$. Consider an arbitrary player $i \in Q_\lambda$. By the above we have, $c_i(s) \geq d_j(P_j(s)) \geq d_j(Q_\lambda) \geq (\lambda - 1)C(s^*)$. Moreover, by the maximality of P_λ and since $i \notin P_\lambda$, we have $c_i(s) < \lambda C(s^*)$. □

Remark 66. Observe that in the above proof we exploit the linearity of the delay functions only in (3.2). In fact, we can draw exactly the same conclusion if all delay functions are *set-subadditive*, i.e., for every $j \in [m]$, $d_j(P_1 \cup P_2) \leq d_j(P_1) + d_j(P_2)$ for every $P_1, P_2 \subseteq [n]$. As a consequence, all our upper bounds on the strong price of anarchy that exploit the first part of Lemma 65 hold for set-subadditive delay functions.

Theorem 67. *The strong price of anarchy of linear bottleneck congestion games is at most m.*

Proof. Let s^* be a strategy profile minimizing C and let s be a strong equilibrium with cost $C(s) = \alpha C(s^*)$ for some $\alpha \in \mathbb{R}_{>1}$. For $\lambda \in (1, \alpha]$, let P_λ be the maximal non-empty set of players $\{i \in [n] \mid c_i(s) \geq \lambda C(s^*)\}$. Applying Lemma 65, we obtain a player set Q_λ such that for every $i \in Q_\lambda$ we have $(\lambda - 1)C(s^*) \leq c_i(s) < \lambda C(s^*)$. Moreover, $\sum_{i \in Q_\lambda} w_i \geq \lambda - 1 > 0$ because $\lambda > 1$ and thus Q_λ is non-empty. Therefore, there exists a set $F = \{Q_\alpha, Q_{\alpha-1}, \ldots, Q_{\alpha-\kappa}\}$ of $\kappa + 1$ player sets that are non-empty and pairwise disjoint, where κ is the largest integer satisfying $\alpha - \kappa > 1$. Every set $Q_\lambda \in F$ contains at least one distinct facility $j \in [m]$ with $(\lambda - 1)C(s^*) \leq d_j(s) < \lambda C(s^*)$. Moreover, there is one facility $j \in [m]$ with $d_j(s) = \alpha C(s^*)$. We conclude that $m \geq |F| + 1 = \kappa + 2 \geq \alpha$ and thus the strong price of anarchy is $\alpha \leq m$. □

Theorem 68. *The strong price of anarchy is at least $m - 1$ in linear bottleneck congestion games and at least $(m + 1)/3$ in linear network bottleneck congestion games.*

Proof. We describe first the lower bound example for linear network bottleneck congestion games, and subsequently adapt it to the desired lower bound example for linear bottleneck congestion games. Define $0! = 1$ and let player $i \in [n]$ have weight $w_i = 1/(i-1)!$. The source and destination vertices of the players in the network are as follows. For each player i there is a distinct source vertex s_i, and all players share the same destination vertex, which we call t. Besides the source and destination vertices, there are $n-1$ auxiliary vertices v_i, $i \in [n] \setminus \{1\}$, in the network. The set of arcs in the network $E = E_1 \cup E_2 \cup \{(s_1, t), (s_n, t)\}$, where:

- $E_1 = \{(s_i, v_{i+1}) \mid i \in [n-1]\} \cup \{(s_i, v_i) \mid i \in [n] \setminus \{1\}\}$.
- $E_2 = \{(v_i, t) \mid i \in [n] \setminus \{1\}\}$.

Then $m = |E| = |E_1| + |E_2| + 2 = 2(n-1) + (n-1) + 2 = 3n - 1$. For each arc $e \in E_1$ define $a_e = 1$.[1] For $i \in [n-1] \setminus \{1\}$ let $a_{(v_i,t)} = (i-1)!$. Also, set $a_{(s_1,t)} = 1$ and $a_{(s_n,t)} = n!$. An example for the case $n = 4$ is depicted in Figure 3.1. Each player has two strategies, an *upper* path and a *lower* path. For almost every player i, the upper path is $\{(s_i, v_{i+1}), (v_{i+1}, t)\}$ and the lower path is $\{(s_i, v_i), (v_i, t)\}$. Exceptions are the upper path of player n, which is $\{(s_n, t)\}$, and the lower path of player 1, which is $\{(s_1, t)\}$.

Under configuration s where all players choose their upper path as their strategy, we have $c_i(s) = (1/(i-1)!)i! = i$, thus $C(s) = n = (m+1)/3$. We claim s is a strong equilibrium. Consider any coalition $P \subseteq [n]$ changing its strategies, and call s' the resulting profile. Let i be the least player in P and assume first $i \geq 2$. Then under s and s', $i-1$ plays its upper path, $\{(s_{i-1}, v_i), (v_i, t)\}$. The only strategies available to i are (s_i, v_i) and (v_i, t). Then:

$$
\begin{aligned}
c_i(s') &= \max\{d_{(s_i,v_i)}(P_{(s_i,v_i)}(s')), d_{(v_i,t)}(P_{(v_i,t)}(s'))\} \\
&= \max\left\{\frac{1}{(i-1)!}, (i-1)!\left(\frac{1}{(i-1)!} + \frac{1}{(i-2)!}\right)\right\} \\
&= \max\left\{\frac{1}{(i-1)!}, 1 + (i-1)\right\} = i = c_i(s).
\end{aligned}
$$

Player 1 will not participate in any deviating coalition either, because a cost of 1 is incurred to it under s and when playing its lower path. So s is a strong equilibrium. In the socially optimum configuration every player plays its lower path alone and has cost $c_i(s^*) = (1/(i-1)!)(i-1)! = 1$. Thus, the strong price of anarchy is at least

[1] We abuse notation slightly and identify the arcs of the graph with the facilities, instead of explicitly constructing a bijection b conform Definition 59.

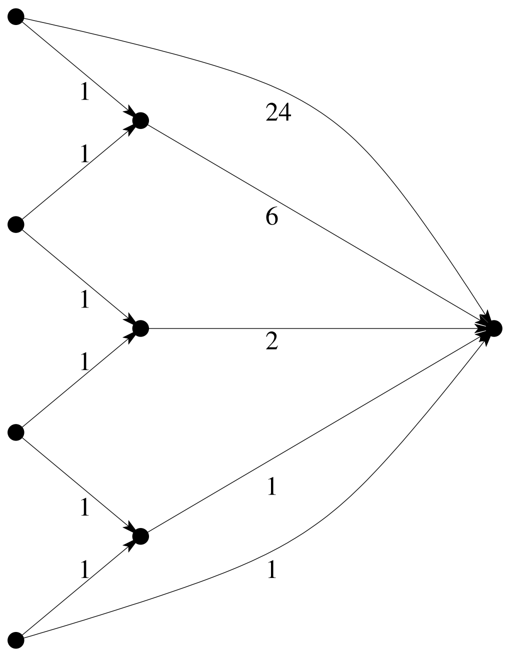

Figure 3.1: A worst-case example for the strong price of anarchy of linear network bottleneck congestion games.

$(m+1)/3$. For the general case we simply regard links in $\{(s_1,t)\} \cup E_2 \cup \{(s_n,t)\}$ as $m = n+1$ facilities $e_{i'}$, $i' \in [n+1]$ and restrict the strategy space of every player i to $\{e_i, e_{i+1}\}$. If every player plays e_{i+1} we obtain a strong equilibrium similar to s described above. The social optimum occurs when player i plays e_i. We omit a detailed analysis, since a similar construction has appeared in Gairing et al. [2006]. □

3.3.2 Identical Players

We next derive an upper bound on the strong price of anarchy for linear bottleneck congestion games if the weights of all players are identical. We assume without loss of generality that the weight w_i of each player $i \in [n]$ is 1.

Theorem 69. *Let Γ be a bottleneck congestion game and s^* be a strategy profile of Γ that minimizes C. The strong price of anarchy of Γ is at most*

$$\min\left\{\frac{1}{2}+\sqrt{2n-3/2}, \frac{1}{2}+\sqrt{mC(s^*)-3/2}\right\}$$

for linear bottleneck congestion games with identical player weights, and the strong price of anarchy is 2 for linear symmetric network bottleneck congestion games with identical player weights.

Proof. We start with proving the first claim. Let s be a strong equilibrium with cost $C(s) = \alpha\gamma^*$ for some $\alpha \in \mathbb{R}_{\geq 1}$. As in the proof of Theorem 67, we can apply Lemma 65 to identify a set $F = \{Q_\alpha, Q_{\alpha-1}, \ldots, Q_{\alpha-\kappa}\}$ of $\kappa+1$ player sets that are non-empty and pairwise disjoint, where κ is the largest integer satisfying $\alpha - \kappa > 1$. Each such set $Q_\lambda \in F$ contains at least $\lambda - 1$ players, i.e., $|Q_\lambda| \geq \lceil \lambda - 1 \rceil$. Moreover, there is at least one player that experiences a congestion of $\alpha C(s^*)$. Thus

$$n \geq 1 + \sum_{\lambda \in [\lceil \alpha - 1 \rceil]} \lambda \geq 1 + \frac{\alpha(\alpha-1)}{2}.$$

Solving for α we obtain $\alpha \leq \frac{1}{2} + \sqrt{2n - 3/2}$. Recall that we assume without loss of generality that $a_j \geq 1$ for every $j \in [m]$ and thus $C(s^*) \geq n/m$. We therefore also obtain $\alpha \leq \frac{1}{2} + \sqrt{mC(s^*) - 3/2}$.

We continue with the second claim of the theorem. In a strong equilibrium s, at least one player $i \in [n]$ must have cost $c_i(s) \leq C(s^*)$ since otherwise the grand coalition could deviate to the socially optimal strategy profile. Suppose there is a player $i' \in [n]$ having cost more than two times larger than the cost of i. Consider the strategy profile $s' = (s_i, s_{-i'})$ where player i' deviates to the strategy of player i. Then $c_{i'}(s') \leq \max_{j \in s_i}(d_j(P_j(s)) + a_j) \leq \max_{j \in s_i} 2d_j(P_j(s)) \leq 2c_i(s)$, which is a contradiction to s being a strong equilibrium.

Tightness of this bound is shown by Example 70 below, which is a symmetric network bottleneck congestion game with identical player weights. □

Example 70. Let $n = 3$ and $m = 6$. The strategy set of every player is $\{S_1 = \{1\}, S_2 = \{2,3\}, S_3 = \{4,5\}, S_4 = \{2,5,6\}\}$. The delay functions of the facilities are identical and are given by $d_j(P) = |P|$ for all $j \in [m], P \subseteq [n]$. The optimal social cost is attained for the strategy profile s^*, where player $i \in [3]$ plays S_i. Then, $C(s^*) = 1$. A strong equilibrium s is given by $s_1 = S_4$ and $s_2 = s_3 = S_1$. The cost of s is $C(s) = 2$. It is easy to verify that this example is a network bottleneck congestion game.

3.4 Identical Facilities

In this section, we study the strong price of anarchy for the case of linear bottleneck congestion games with facilities that have identical delay functions. We refer to this case as bottleneck congestion games with *identical facilities*. For this section, we assume without loss of generality that the delay function of every facility $j \in [m]$ is given by $d_j(P) = \sum_{i \in P} w_i$ for all $P \subseteq [n]$.

Theorem 71. *The strong price of anarchy of linear bottleneck congestion games with identical facilities is at most* $-\frac{1}{2}+\sqrt{2m+\frac{1}{4}}$. *The strong price of anarchy of symmetric linear bottleneck congestion games with identical facilities is* 2.

Proof. Let s^* denote a strategy profile attaining minimum social cost. For the symmetric case observe that in any strong equilibrium s, there is at least one player i' with $c_{i'}(s) \leq C(s^*)$. Indeed, if $c_i(s) > C(s^*)$ for all $i \in [n]$, then $[n]$ would deviate to s^*. Let $i \in [n]$ and $j \in [m]$ be such that $j \in s_i$ and $c_i(s) = d_j(s) = C(s)$. Consider player i deviating to $s_{i'}$. Then, because s is a pure equilibrium, $C(s) = c_i(s) \leq c_{i'}(s) + w_i \leq 2C(s^*)$. From Example 70 it follows that this bound is tight.

For the asymmetric case, let the cost of a strong equilibrium s be $C(s) = \alpha C(s^*)$, for some $\alpha \in \mathbb{R}_{>1}$. Similar to the proof of Theorem 67, let P_λ be the (maximal) (non-empty) set of players $\{i \in [n] \mid c_i(s) \geq \lambda C(s^*)\}$ for $\lambda \in (1, \alpha]$. By Lemma 65, we obtain a player set Q_λ such that for every $i \in Q_\lambda$ we have $(\lambda - 1)C(s^*) \leq c_i(S) < \lambda C(s^*)$, and there is a $j \in [m]$ such that $j \in s_i$ for all $i \in Q_\lambda$. Because the facilities are identical, and by non-emptiness of Q_λ it holds that $\sum_{i \in Q_\lambda} w_i = a_j \sum_{i \in Q_\lambda} w_i \geq (\lambda - 1)C(s^*) > 0$ by Lemma 65. That is, we can identify a set $F = \{Q_\alpha, Q_{\alpha-1}, \dots, Q_{\alpha-\kappa}\}$ of $\kappa + 1$ player sets that are non-empty and pairwise disjoint, where κ is the largest integer satisfying $\alpha - \kappa > 1$. Moreover, by construction we have $P_\alpha \cap Q_\lambda = \varnothing$ for every $Q_\lambda \in F$ and $\sum_{i \in P_\alpha} w_i \geq \alpha C(s^*)$ since facilities are identical. The total weight $\sum_{i \in [n]} w_i$ is then:

$$
\begin{aligned}
\sum_{i \in [n]} w_i &\geq \alpha C(s^*) + \sum_{\lambda=\alpha-\kappa}^{\alpha} \sum_{i \in Q_\lambda} w_i \\
&\geq \alpha C(s^*) + \sum_{\lambda=\alpha-\kappa}^{\alpha} (\lambda - 1)C(s^*) \\
&\geq \alpha C(s^*) + \sum_{\lambda \in [\alpha-1]} \lambda C(s^*).
\end{aligned}
$$

The latter equals $(1/2)\alpha(1+\alpha)C(s^*)$. Observe that $C(s^*) \geq \sum_{i\in[n]} w_i/m$ because the facilities are identical. We obtain $2m \geq \alpha(1+\alpha)$ or equivalently $\alpha \leq -\frac{1}{2} + \sqrt{2m+1/4}$. □

Theorem 72. *The strong price of anarchy of linear bottleneck congestion games with identical players and identical facilities is at least* $(-1/2)+\sqrt{2m+(1/4)}$. *The strong price of anarchy of linear network bottleneck congestion games with identical players and identical facilities is at least* $(-1/4)+\sqrt{2+2m}/2$.

Proof. We construct a sequence of bottleneck congestion game instances where m grows to infinity. That the price of anarchy of each of these instances will be at least $(-1/2)+\sqrt{2m+(1/4)}$. We turn these instances subsequently into network bottleneck congestion game instances in order to prove the second claim.

Fix $q \in \mathbb{N}_{>0}$ and consider a partition of the set of players $[n]$ into q subsets, $[n] = \bigcup_{\ell\in[q]} P_\ell$, where $|P_\ell| = \ell$ for all $\ell \in [q]$. Denote players in P_ℓ by $p_{\ell i}$, $i \in [\ell], \ell \in [q]$. Likewise, for each $\ell \in [q]$ there is a set of facilities $E_\ell = \{e_{\ell j} \mid j \in [\ell]\}$. Define moreover $E_{q+1} = E_1$. For every player $p_{\ell i} \in P_\ell, i \in [\ell], \ell \in [q]$, the strategy set of $p_{\ell i}$ is

$$\Sigma_{p_{\ell i}} = \{\{e\} \mid e \in E_\ell\} \cup \{E_{\ell+1}\}$$

The socially optimal configuration s^* is given by $s^*_{p_{\ell i}} = \{e_{\ell i}\}$ for all $i \in [\ell], \ell \in [q]$. Then $C(s^*) = 1$. Now consider the configuration s where $s_{p_{\ell i}} = E_{\ell+1}$ for all $i \in [\ell]$, $\ell \in [q]$. For every player $i \in P_\ell, \ell \in [q]$, we now have $c_i(s) = \ell$. $C(s)$ is therefore the delay of the unique facility $j = e_{11} \in E_1$ and is $d_j(P_j(s)) = |P_q| = q$.

We claim that s is a strong equilibrium. Assume for the sake of contradiction that this is not the case. Then there must exist a player set $P \subseteq [n]$ that can change its strategies such that it decreases the cost of every player in P. Let s' be the strategy profile resulting from the change in strategies of P. Note that for the player $p = p_{11} \in P_1$ we have $c_p(s) = 1$, hence no deviation can decrease its cost and $P_1 \cap P = \varnothing$.

Let $\ell = \min\{\ell' \mid P_{\ell'} \cap P \neq \varnothing\}$. Then $\ell \geq 2$. For all $\ell - 1$ players $p_{\ell-1,i} \in P_{\ell-1}$ it holds that $s_{p_{\ell-1,i}} = E_\ell$, because $P \cap P_{\ell-1} = \varnothing$. Hence, $c_\ell(s') = \ell - 1 + 1 = \ell = c_\ell(s)$. This is in contradiction with the definition of P, hence s is a strong equilibrium, and the strong price of anarchy is at least q. We have

$$m = \left| \bigcup_{\ell\in[q]} E_j \right| = \sum_{\ell\in[q]} \ell = \frac{q(q+1)}{2},$$

which yields $q \geq (-1/2) + \sqrt{2m+(1/4)}$.

We convert the example into a network bottleneck congestion game. The player set will remain the same, but the facility set will be different and will correspond to the edge set of a directed graph G that we will now describe. The strategy set of a player in $P_\ell, \ell \in [q]$ consist of the paths in G from a source node s_ℓ to a destination node t. I.e., all players share the same destination node The remainder of the vertex set of G follows implicitly from the description of the paths available to each player. It is recommended to inspect Figure 3.2 along with the description that follows, which depicts the network for $q = 4$.

For each $\ell \in [q], i \in [\ell-1]$ there is a path of length 3: $\{(s_\ell, u_{\ell i}), (u_{\ell i}, v_{\ell i}), (v_{\ell i}, t)\}$, and there is additionally a path of length 2: $\{(s_\ell, u_{\ell\ell}), (u_{\ell\ell}, t)\}$. Let A_ℓ be the union of these paths. There are auxiliary arcs $A'_\ell = \{(v_{\ell i}, u_{\ell,i+1}) \mid i \in [\ell - 1]\}, \ell \in [q]$. And finally, there is an arc $(s_{\ell-1}, u_{\ell 1})$, $\ell \in [q] \setminus \{1\}$. For the last group of players we add an arc (s_q, t).

We illustrate the analog of s on the constructed network. For $\ell \in [q-1]$, all players $p_{\ell i} \in P_\ell$ choose the strategy corresponding to the following path:

$$\begin{aligned} s_{\ell i} &= \{(s_\ell, u_{\ell+1,1}\} \\ &\cup \{(u_{\ell+1,r}, v_{\ell+1,r}), (v_{\ell+1,r}, u_{\ell+1,r+1}) \mid r \in [\ell - 1]\} \\ &\cup \{(u_{\ell+1,\ell}, v_{\ell+1,\ell}), (v_{\ell+1,\ell}, t)\}, \end{aligned}$$

and $s_{p_{qi}} = (s_q, t)$ for $i \in [q]$. The proof that s is strong is analogous to the proof given for the non-network example. In the optimal configuration, every player plays a disjoint path. The number of links m is

$$m = \sum_{\ell \in [q]} (|A_j| + |A'_j|) + q = \sum_{j \in [q]} (3\ell - 1 + (\ell - 1)) + q - 1 = 2q^2 + q - 1,$$

which yields $q \geq (-1/4) + \sqrt{2 + 2m}/2$. $\square$

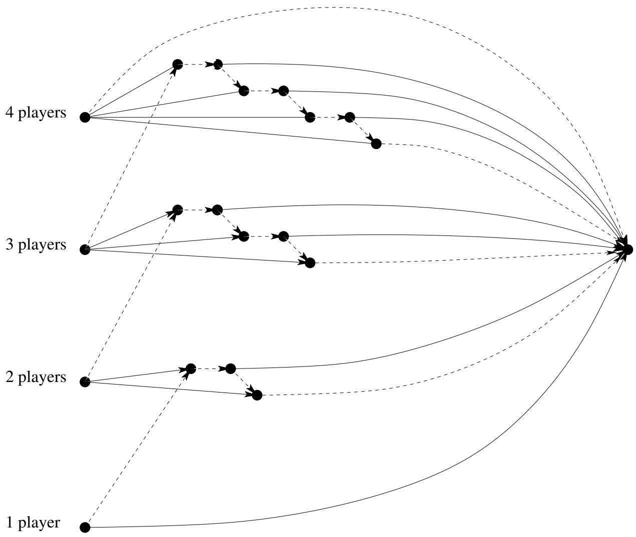

Figure 3.2: An example instance from the set of instances constructed in the proof of Theorem 72. The dashed edges indicate the strategy played by the players according to the strong equilibrium.

Chapter 4

The Robust Price of Anarchy of Altruistic Games*

Strategic games are often studied under the assumption that players strive to optimize a utility function that reflects an entirely selfish objective. In many realistic scenarios however, we may assume that players do not behave entirely selfishly but instead also take into account the well-being of others. In this chapter, we introduce a simple model for introducing altruistic behavior in games. We focus on establishing bounds on the price of anarchy for various well-known classes of games, extended with altruism. The bounds we obtain are in many cases expressed as a function of the degree of altruism introduced into the game. In order to establish these bounds, we adapt the smoothness notion of Roughgarden [2009] for use in altruistic games. Because we make use of (an adaptation of) smoothness, many of our bounds on the price of anarchy will turn out to hold for the broad equilibrium concept of coarse equilibria.

Among the preliminaries given in Chapter 1, the relevant material for this chapter is Section 1.3 up to Section 1.3.1.4, and Section 1.3.1.7 up to Section 1.3.1.10, with emphasis on Sections 1.3.1.9 and 1.3.1.10.

The classes of games we study (all full information games) are linear congestion games, symmetric singleton linear congestion games, fair cost sharing games, and *valid utility games*. For the first three classes, see Section 1.3.1.10. The fourth class is defined as follows.

*The contents of this chapter have been published as Chen et al. [2011a,b]. Part of the content of this chapter appears in the Ph.D. thesis of Po-An Chen. Both authors have contributed to this content in equal shares.

Definition 73 (Valid utility game [Vetta, 2002], valid social welfare function)**.** A *valid utility game* [Vetta, 2002] is a utility maximization game $\Gamma = (n, \Sigma, u)$ where there is a set $[m]$ of facilities and the strategy sets Σ_i are subsets of $2^{[m]}$. Moreover, there must exist a *valid social welfare function* for Γ, which is a social welfare function $U : \Sigma \to \mathbb{R}$ for which there exists a function $V : 2^{[m]} \to \mathbb{R}$ such that:

- $U(s) = V\left(\bigcup_{i\in[n]} s_i\right)$ for all $s \in \Sigma$;
- V is *submodular*, i.e., for all $S_1, S_2 \subseteq [m], j \in E, S_1 \subseteq S_2$, it holds that $V(S_1 \cup \{j\}) - V(S_1) \geq V(S_2 \cup \{j\}) - V(S_2)$.
- V is *non-decreasing*, i.e., for all $S_1, S_2 \subseteq [m], S_1 \subseteq S_2$, it holds that $V(S_1) \leq V_(S_2)$.
- $u_i(s) \geq U(s) - V\left(\bigcup_{i'\in[n]\setminus\{i\}} s_{i'}\right)$ for all $i \in [n], s \in \Sigma$. Intuitively, a player's utility is at least its contribution to U.
- $\sum_{i\in[n]} u_i(s) \leq U(s)$ for all $s \in \Sigma$.

Examples of games falling into the class of valid utility games include natural game-theoretic variants of the facility location, k-median and network routing problems [Vetta, 2002]. A bound of 2 was proved on the pure price of anarchy of valid utility games with respect to any valid social welfare function, by Vetta. Roughgarden [2009] showed how this bound can be achieved via smoothness, lifting the result up to the coarse price of anarchy. We extend this result to the altruistic extensions of these games.

Thus, the price of anarchy for altruistic valid utility games will be studied with respect to any valid social welfare function. The price of anarchy for altruistic linear congestion games, symmetric singleton linear congestion games, and fair cost sharing games, will be studied with respect to the standard sum-of-cost social cost, i.e., (1.3.1.8).

We model altruistic behavior by assuming that player i's perceived cost is a convex combination of $1 - \alpha_i$ times its direct cost and α_i times the social cost. Tuning the parameters α_i allows smooth interpolation between purely selfish and purely altruistic behavior. Our altruism model is based on one used (among others) by Chen et al. [2010], and similar to ones introduced by Caragiannis et al. [2010], Chen and Kempe [2008], Chen et al. [2010], Hoefer and Skopalik [2009b], Ledyard [1997].

Definition 74 (Altruistic extension of a cost-minimization game, $\hat{\alpha}$, $\check{\alpha}$, (non)-uniform altruism)**.** Let $\Gamma = (n, \Sigma, c)$ be a cost minimization game and let C be a social cost function for Γ. Let $\alpha \in [0,1]^n$. The *α-altruistic extension* of G with respect to C (or simply *α-altruistic game*) is defined as the strategic game $\Gamma^\alpha = (n, \Sigma, c^\alpha)$, where for every $i \in [n]$ and $s \in \Sigma$,

$$c_i^\alpha(s) \quad = \quad (1-\alpha_i)c_i(s) + \alpha_i C(s).$$

The function c_i is called the *selfish cost*, and c_i^α is called the *perceived cost function*, in the context of Γ. Thus, the perceived cost that player i experiences is a convex combination of its direct (selfish) cost c_i and the social cost C. We call player i in game Γ^α an *α_i-altruistic player*. When $\alpha_i = 0$, we say that player $i \in [n]$ is *entirely selfish*. A player $i \in [n]$ with $\alpha_i = 1$ is entirely altruistic. When $\alpha_i = \alpha_j$ for all $i, j \in [n]$, we speak of *uniform* altruism, and the game is a *uniformly* altruistic extension. Otherwise, we speak of *non-uniform* altruism, and the game is a *non-uniformly* altruistic extension.

Given an altruism vector $\alpha \in [0,1]^n$, we let $\hat{\alpha} = \max_{i \in N} \alpha_i$ and $\check{\alpha} = \min_{i \in N} \alpha_i$ denote the maximum and minimum altruism levels, respectively.

We define the altruistic extension of a utility maximization game analogously.

Definition 75 (Altruistic extension of a utility maximization game)**.** Let $\Gamma = (n, \Sigma, u)$ be a utility maximization game and let U be a social welfare function for Γ. Let $\alpha \in [0,1]^n$. The *α-altruistic extension* of G with respect to U (or simply *α-altruistic game*) is defined as the strategic utility maximization game $\Gamma^\alpha = (n, \Sigma, u^\alpha)$, where for every $i \in [n]$ and $s \in \Sigma$,

$$u_i^\alpha(s) \quad = \quad (1-\alpha_i)u_i(s) + \alpha_i U(s).$$

The function u_i is called the *selfish utility*, and u_i^α is called the *perceived utility function*, in the context of Γ.

We note that the altruistic part of a player's perceived cost does not recursively take other players' *perceived* cost into account. Such recursive definitions of altruistic utility have been studied, e.g., by Bergstrom [1999], and can be reduced to our definition under suitable technical conditions.

The price of anarchy and price of stability are studied with respect to the social cost function C of the *original* game, and not its altruistic extension. This reflects our desire to understand the overall quality of the equilibria of the strategic game, which is not affected by different *perceptions* of costs of the individual players. Note, however, that if all players have a uniform altruism level $\alpha_i = \alpha \in [0,1]$ and the social cost function C is equal to the sum of all players' individual costs, then for every strategy profile $s \in \Sigma$, $C^\alpha(s) = (1 - \alpha + \alpha n)C(s)$, where $C^\alpha(s) = \sum_{i \in N} C_i^\alpha(s)$ denotes the sum of all players' perceived costs. In particular, bounding the price of anarchy with respect to C is equivalent to bounding the price of anarchy with respect to total perceived cost C^α in this case.

The bounds on the price of anarchy that we establish in this chapter show that for congestion games and cost-sharing games, the worst-case price of anarchy increases with increasing altruism, while for valid utility games, it remains constant and is not affected by altruism. However, the increase in the price of anarchy is not a universal phenomenon: for symmetric singleton linear congestion games, we derive a bound on the pure price of anarchy that decreases as the level of altruism increases. Since the bound is also strictly lower than the mixed price of anarchy, it exhibits a natural example in which pure equilibria are more efficient than more permissive solution concepts.

4.1 Background

The accuracy of the price of anarchy as a predictor for how bad the natural stable outcomes of a game can be, has been criticized by a variety of arguments. One of these arguments is that the assumption that players seek only to maximize their own utility is at odds with altruistic behavior routinely observed in the real world. While modeling human incentives and behavior accurately is a formidable task, several papers have proposed natural models of altruism [Ledyard, 1997, Levine, 1998] and analyzed its impact on the equilibria of various games [Caragiannis et al., 2010, Chen et al., 2010, Chen and Kempe, 2008, Elias et al., 2010].

The work of Hayrapetyan et al. [2006], where the impact of collusion in network congestion games is studied (see Section 3.1 for a discussion) can be regarded as being orthogonal to the viewpoint that we adopt in this chapter: in their setting, players are assumed to be entirely altruistic but locally attached to their coalitions. In contrast, in our setting, players may have different levels of altruism but locality does not play a role.

Several recent studies investigate "irrational" player behavior in games. Examples include studies on malicious (or spiteful) behavior [Babaioff et al., 2007, Brandt et al., 2007, Chen and Kempe, 2008, Karakostas and Viglas, 2007] and unpredictable (or Byzantine) behavior [Blum et al., 2008, Moscibroda et al., 2006, Roth, 2008]. The work that is most related to our work in this context is the one by Blum et al. [2008]. The authors consider games that the players play repeatedly, where every player is assumed to minimize its own regret (see Section 1.3.1.4). They derive bounds on the inefficiency of the resulting outcomes for certain classes of games, including congestion games and valid utility games. This degree of inefficiency is termed the *total price of anarchy*. The bounds they prove exactly match the respective price of anarchy and even continue to hold if only part of the players minimize their regret, while the other players behave arbitrarily. The latter result is surprising in the context of valid utility games because it means that the price of total anarchy remains at 2, even if additional players are added to the game that behave arbitrarily. Our findings allow us to draw an even more dramatic conclusion: Our bounds extend to the total price of anarchy of the respective repeated games (see Section 4.3). As a consequence, our result for valid utility games implies that the price of total anarchy would remain at 2, even if the arbitrarily behaving players were to act altruistically. That is, while the result in Blum et al. [2008] suggests that arbitrary behavior does not harm the inefficiency of the final outcome, our result shows that altruistic behavior does not help.

If players' altruism levels are not uniform, then the existence of pure equilibria is not obvious. Hoefer and Skopalik established it for several subclasses of atomic congestion games [Hoefer and Skopalik, 2009b]. For the generalization of arbitrary player-specific cost functions, Milchtaich [1996] showed existence of pure equilibria for singleton congestion games, and Ackermann et al. [2006] did so for matroid congestion games, in which the strategy space of each player are the bases of a matroid on the set of facilities.

Models of Altruism Models of altruism either identical or very similar to the one used in this chapter have been studied in several papers. Perhaps the first published suggestion of a similar model is due to Ledyard [1997], but since then, variations of it have been studied more extensively, e.g., in Caragiannis et al. [2010], Chen et al. [2010], Chen and Kempe [2008], Elias et al. [2010]. The main difference is that in some of these models, linear combinations (rather than convex combinations) of cost functions are considered. For most of these variations, a straightforward scaling of the coefficients shows equivalence with the model we consider here.

Our altruism model can be naturally extended to include $\alpha_i < 0$, modeling spiteful behavior (see, e.g., Chen and Kempe [2008]). While this modeling extension is natural, several results of this chapter and other papers do not continue to hold directly for negative α_i.

Besides models based on linear combinations of individual players' costs (as well as social welfare), several other approaches have been studied as well. Generally, altruism or other "other-regarding" social behavior has received some attention in the behavioral economics literature (e.g., Gintis et al. [2005]). Alternative models of altruism and spite have been proposed by Levine [1998], Rabin [1993], and Geanakoplos et al. [1989]. These models are designed more with the goal of modeling the psychological processes underlying spite or altruism (and reciprocity): they involve players forming beliefs about other players. As a result, they are well-suited for experimental work, but perhaps not as directly suited for the type of analysis done in this chapter.

4.2 Contributions and Outline

In Section 4.3 we present our extension of smoothness, and prove some preliminary properties about it in order to motivate this extension.

Subsequently, we analyze the price of anarchy of altruistic extensions of four classes of games. In many cases, we use our smoothness extension to do this.

1. Using our framework, we derive in Section 4.4 a bound of $n/(1-\hat{\alpha})$ on the coarse price of anarchy of fair cost sharing games, where $\hat{\alpha}$ is the maximum altruism level of a player. This bound is tight for uniformly altruistic players.

2. For valid utility games, in Section 4.5 we show that the bound of 2 on the price of anarchy that was proved by Vetta [2002], remains 2 in the altruistic setting, this bound holds with respect to all valid social welfare functions and is independent of the altruism level. Moreover, we show that this bound is tight.

3. For linear congestion games, Caragiannis et al. [2010] derive a tight bound of $(5+4\alpha)/(2+\alpha)$ on the pure price of anarchy when all players have the *same* altruism level α.

   Remark 76. Caragiannis et al. [2010] model uniformly altruistic players by defining the perceived cost of player i as $(1-\xi)C_i(s) + \xi(C(s) - C_i(s))$, where $\xi \in [0, 1]$. It is not hard to see that in the range $\xi \in [0, \frac{1}{2}]$ this definition is equivalent to ours by setting $\alpha = \xi/(1-\xi)$ or $\xi = \alpha/(1+\alpha)$.[1] Therefore, various bounds we cite here are stated differently in Caragiannis et al. [2010].

   Our framework makes it an easy observation that their proof in fact bounds the coarse price of anarchy. In Section 4.6, we generalize their bound to the case when different players have different altruism levels, obtaining a bound in terms of the maximum and minimum altruism levels. This partially answers an open question from Caragiannis et al. [2010].

4. In Section 4.7, for the special case of symmetric singleton congestion games, we extend our study of non-uniform altruism and obtain an improved bound of $(4-2\alpha)/(3-\alpha)$ on the price of anarchy when an α-fraction of the players are entirely altruistic and the remaining players are entirely selfish.

Notice that many of these bounds on the coarse price of anarchy reveal a counter-intuitive trend: For valid utility games, the bound is independent of the level of altruism, and for congestion games and cost-sharing games, it actually *increases* in the

[1] The model of Caragiannis et al. [2010] with $\xi \in (\frac{1}{2}, 1]$ has players assign strictly more weight to others than to themselves, a possibility not present in our model since we consider altruism to be caring about others' costs at most as much as about one's own cost.

altruism level, unboundedly so for cost-sharing games. Intuitively, this phenomenon is explained by the fact that a change of strategy by player i may affect many players. An altruistic player will care more about these other players than a selfish player; hence, an altruistic player accepts more states as "stable". This suggests that the best stable solution can also be chosen from a larger set, and the price of stability should thus decrease. Our results on the price of stability lend support to this intuition: for congestion games, we derive an upper bound on the price of stability which decreases as $2/(1+\alpha)$ (Section 4.6); similarly, for fair cost sharing games, we establish an upper bound which decreases as $(1-\alpha)H_n + \alpha$ (Section 4.4).

The increase in the price of anarchy is not a universal phenomenon, demonstrated by our results for symmetric singleton congestion games. Caragiannis et al. [2010] showed a bound of $4/(3+\alpha)$ for pure equilibria with uniformly altruistic players, which decreases with the altruism level α. Our bound of $(4-2\alpha)/(3-\alpha)$ for mixtures of entirely altruistic and selfish players is also decreasing in the fraction of entirely altruistic players. We also extend an example of Lücking et al. [2008] to show that symmetric singleton congestion games may have a mixed price of anarchy arbitrarily close to 2 for arbitrary altruism levels (Section 4.7). In light of the above bounds, this establishes that pure Nash equilibria can result in strictly lower price of anarchy than weaker solution concepts.

Additionally, we provide in Section 4.8 some general insights in how the robust price of anarchy behaves as a function of the altruism levels of the players: this function is always quasi-convex, leading to the corollary that for classes of games for which the robust price of anarchy gives tight bounds, the worst case price of anarchy is attained when the players are either fully altruistic or fully selfish.

We end this chapter with Section 4.9, providing some concluding remarks and a discussion about future research.

4.3 Smoothness for Altruistic Extensions of Games

In extending the definition of smoothness to altruistic games, we have to exercise some care. Simply applying the classical smoothness technique to the new game does not work, as the social cost function we wish to bound is the sum of all *direct* costs instead of the sum of all *perceived costs*. Moreover, the social cost function is in general not sum-bounded for an altruistic extension of a game Γ, even if it is sum-bounded for Γ. Sum-boundedness is needed for applying the smoothness technique. For this reason, we propose a revised definition of smoothness that is suitable for use with altruistic extensions of games.

For notational convenience, for a cost minimization game $\Gamma = (n, \Sigma, c)$ we define

$c_{-i}(s) = C(s) - c_i(s) \leq \sum_{j \in [n] \setminus \{i\}} c_j(s)$. Note that when the social cost is the sum of all players' costs, the inequality is an equality.

Definition 77 ((λ, μ, α)-smoothness)**.** Let $\Gamma = (n, \Sigma, c)$ be a cost minimization game, let $\alpha \in [0,1]^n$, and let C be a social cost function for Γ. Γ is (λ, μ, α)-*smooth* with respect to C iff for any two strategy profiles $s, s^* \in \Sigma$,

$$\sum_{i \in [n]} c_i(s_i^*, s_{-i}) + \alpha_i(c_{-i}(s_i^*, s_{-i}) - c_{-i}(s)) \leq \lambda C(s^*) + \mu C(s). \tag{4.1}$$

If Γ is instead a utility maximization game (n, Σ, μ) and U is a social welfare function for Γ, then we define Γ^α to be (λ, μ, α)-smooth with respect to U iff for any two strategy profiles $s, s^* \in \Sigma$,

$$\sum_{i \in [n]} u_i(s_i^*, s_{-i}) + \alpha_i(u_{-i}(s_i^*, s_{-i}) - u_{-i}(s)) \geq \lambda U(s^*) - \mu U(s). \tag{4.2}$$

For $\alpha = \mathbf{0}$, this definition coincides with Roughgarden's notion of (λ, μ)-smoothness. To gain some intuition, consider two strategy profiles $s, s^* \in \Sigma$, and a player $i \in [n]$ that switches from its strategy s_i under s to s_i^*, while the strategies of the other players remain fixed at s_{-i}. The contribution of player i to the left-hand side of (4.1) then accounts for the individual cost that player i perceives after the switch plus α_i times the difference in social cost caused by this switch excluding player i. The sum of these contributions needs to be bounded by $\lambda C(s^*) + \mu C(s)$. We will see that this definition of (λ, μ, α)-smoothness allows us to derive the coarse price of anarchy of altruistic extensions of some large and important classes of games.

Preliminary Results

We first show that many of the results of Roughgarden [2009] following from regular (λ, μ)-smoothness carry over to our altruistic setting using the extended (λ, μ, α)-smoothness notion (Definition 77). Even though some care has to be taken in extending these results, most of the proofs of the propositions in this section follow along similar lines as the analogues of Roughgarden [2009].

Proposition 78. *Let $\Gamma = (n, \Sigma, c)$ be a cost minimization game, let $\alpha \in [0,1]^n$, let Γ^α be a α-altruistic extension of a cost minimization game Γ, and let C be a social cost function that is sum-bounded with respect to Γ. If Γ is (λ, μ, α)-smooth with respect to C for some $\lambda, \mu \in \mathbb{R}$ and $\mu < 1$, then the coarse price of anarchy of Γ^α with respect to C is at most $\lambda/(1-\mu)$.*

Likewise, suppose that Γ *is instead a utility maximization game* Γ*, and* U *is a social welfare function that is sum-bounded with respect to* Γ*. If* Γ *is* (λ, μ, α)*-smooth with respect to* U *for some* $\lambda, \mu \in \mathbb{R}$ *and* $\mu > -1$*, then the coarse price of anarchy of* Γ^α *with respect to* U *is at most* $(1 + \mu)/\lambda$*.*

Proof. We prove the claim for cost minimization games. The proof of the claim for utility maximization games is analogous.

Let σ be a coarse equilibrium of Γ^α, and $s^* \in \Sigma$ an arbitrary strategy profile. The coarse equilibrium condition implies that for every player $i \in [n]$:

$$\mathbf{E}_{s\sim\sigma}[(1-\alpha_i)c_i(s) + \alpha_i C(s)] \leq \mathbf{E}_{s\sim\sigma}[(1-\alpha_i)c_i(s_i^*, s_{-i}) + \alpha_i C(s_i^*, s_{-i})].$$

By linearity of expectation, for every player $i \in [n]$:

$$\mathbf{E}_{s\sim\sigma}[c_i(s)] \leq \mathbf{E}_{s\sim\sigma}[c_i(s_i^*, s_{-i}) + \alpha_i(C(s_i^*, s_{-i}) - c_i(s_i^*, s_{-i})) - \alpha_i(C(s) - c_i(s))].$$

By summing over all players and using sum-boundedness of C and linearity of expectation, we obtain

$$\mathbf{E}_{s\sim\sigma}[C(s)] \leq \mathbf{E}_{s\sim\sigma}\left[\sum_{i\in[n]} c_i(s_i^*, s_{-i}) + \alpha_i(c_{-i}(s_i^*, s_{-i}) - c_{-i}(s))\right].$$

Now we use (4.1) to conclude

$$\mathbf{E}_{s\sim\sigma}[C(s)] \leq \mathbf{E}_{s\sim\sigma}[\lambda C(s^*) + \mu C(s)] = \lambda C(s^*) + \mu \mathbf{E}_{s\sim\sigma}[C(s)].$$

Solving for $\mathbf{E}[C(s)]$ now proves the claim. $\square$

As we show later, for many important classes of games, the bounds obtained by (λ, μ, α)-smoothness arguments are tight, even for pure equilibria. Therefore, as in Roughgarden [2009], we define the *robust price of anarchy* as the best possible bound on the coarse price of anarchy obtainable by a (λ, μ, α)-smoothness argument.

Definition 79. Let $\Gamma = (n, \Sigma, c)$ be a cost minimization game and let $\alpha \in [0,1]^n$. The α*-robust price of anarchy* of Γ with respect to a social cost function C for Γ, is defined as

$$\text{RPoA}_\Gamma(\alpha) = \inf\left\{\frac{\lambda}{1-\mu} \;\middle|\; \Gamma \text{ is } (\lambda, \mu, \alpha)\text{-smooth with respect to } C,\, \lambda, \mu \in \mathbb{R}, \mu < 1\right\}.$$

If Γ is instead a utility maximization game, and U is a social welfare function for Γ, then the α-robust price of anarchy of Γ with respect to U is defined as

$$\text{RPoA}_\Gamma(\alpha) = \inf\left\{\frac{\lambda}{1-\mu} \;\middle|\; \Gamma \text{ is } (\lambda, \mu, \alpha)\text{-smooth with respect to } U,\, \lambda, \mu \in \mathbb{R}, \mu > -1\right\}.$$

For a class $\mathcal{G}$ of games (each game equipped with a social cost function or social welfare function), we define the α-robust price of anarchy of $\mathcal{G}$ as the supremum of the α-robust price of anarchy of the games in $\mathcal{G}$ (with respect to their social cost functions or social welfare functions).

The smoothness condition also proves useful in the context of no-regret sequences and the *price of total anarchy*, introduced by Blum et al. [2008] (see Section 1.3.1.4).

Proposition 80. *Let $\Gamma = (n, \Sigma, c)$ be a cost minimization game, let $\alpha \in [0,1]^n$, let $s^* \in \Sigma$ be a strategy profile minimizing a sum-bounded social cost function C for Γ, and let $(s^1, s^2, \dots)$ be a vanishing regret sequence of strategy profiles in Σ, with respect to the α-altruistic extension Γ^α, with respect to C. This sequence then satisfies that the factor by which the average social cost deviates from the optimum social cost, converges to at most the α-robust price of anarchy. I.e.,*

$$\lim_{T\to\infty} \frac{1}{T} \sum_{t\in[T]} C(s^t) \leq \mathit{RPoA}_\Gamma(\alpha) C(s^*).$$

If Γ is instead a profit maximization game (n, Σ, u), and U is a sum-bounded social welfare function for Γ, and $(s^1, s^2, \dots)$ is a vanishing regret sequence of strategy profiles in Σ with respect to at most the α-altruistic extension Γ^α with respect to U, then

$$\lim_{T\to\infty} \frac{1}{T} \sum_{t\in[T]} U(s^t) \geq \frac{U(s^*)}{\mathit{RPoA}_\Gamma(\alpha)}.$$

Proof. We prove the claim for cost minimization games. For profit maximization games, the proof is analogous. Let $\lambda, \mu \in \mathbb{R}, \mu < 1$, be such that Γ is (λ, μ, α)-smooth. By sum-boundedness of C and by the definition of $c_i^\alpha, i \in [n]$ it holds for all $t \in \mathbb{N}_{>0}$ that

$$\begin{aligned} C(s^t) &\leq \sum_{i\in[n]} (c_i^\alpha(s^t) - c_i^\alpha(s_i^*, s_{-i}^t)) \\ &\qquad + \sum_{i\in[n]} (c_i(s_i^*, s_{-i}^t) + \alpha_i(c_{-i}(s_i^*, s_{-i}^t) - c_{-i}(s^t))) \\ &\leq \sum_{i\in[n]} (c_i^\alpha(s^t) - c_i^\alpha(s_i^*, s_{-i}^t)) + \lambda C(s^*) + \mu C(s^t), \end{aligned}$$

where the second inequality follows from (λ, μ, α)-smoothness. This is equivalent to

$$C(s^t) \leq \frac{\lambda}{1-\mu} C(s^*) + \frac{1}{1-\mu} \sum_{i\in[n]} (c_i^\alpha(s^t) - c_i^\alpha(s_i^*, s_{-i}^t)). \tag{4.3}$$

The vanishing regret property of $(s^1, s^2, \ldots)$ implies that for all $i \in [n]$,

$$\begin{aligned}
&\lim_{T\to\infty} \frac{1}{T} \sum_{t\in[T]} (c_i^\alpha(s^t) - c_i^\alpha(s_i^*, s_{-i}^t)) \\
&\quad \le \lim_{T\to\infty} \frac{1}{T} \max\left\{ 0, \sum_{t\in[T]} (c_i^\alpha(s^t) - c_i^\alpha(s_i', s_{-i}^t)) \;\middle|\; s_i' \in \Sigma_i \right\} \\
&\quad = 0.
\end{aligned}$$

Combining this with (4.3) proves the claim:

$$\begin{aligned}
&\lim_{T\to\infty} \frac{1}{T} \sum_{t\in[T]} C(s^t) \\
&\le \lim_{T\to\infty} \frac{1}{T} \sum_{t\in[T]} \left(\frac{\lambda}{1-\mu} C(s^*) + \frac{1}{1-\mu} \sum_{i\in[n]} (c_i^\alpha(s^t) - c_i^\alpha(s_i^*, s_{-i}^t)) \right) \\
&\le \left(\frac{\lambda}{1-\mu} C(s^*) + \lim_{T\to\infty} \frac{1}{1-\mu} \sum_{i\in[n]} \sum_{t\in[T]} (c_i^\alpha(s^t) - c_i^\alpha(s_i^*, s_{-i}^t)) \right) \\
&= \frac{\lambda}{1-\mu} C(s^*).
\end{aligned}$$

□

4.4 Fair Cost-sharing Games

It is well-known that the pure price of anarchy of fair cost-sharing games is n [Nisan et al., 2007]. We show that it can get significantly worse in the presence of altruistic players: the following theorem gives a much worse upper bound, which we subsequently show to be tight.

Theorem 81. *Let $n \in \mathbb{N}_{>0}$ and $\alpha \in [0,1]^n$. The α-robust price of anarchy of an n-player cost-sharing game is at most $\frac{n}{1-\hat{\alpha}}$ (where $n/0 = \infty$), with respect to the sum-of-costs social cost function C (1.3.1.8). This bound is tight, even for pure equilibria with uniform altruism. I.e., the pure price of anarchy of the class of uniformly altruistic extensions of n-player fair cost sharing games is $\frac{n}{1-\hat{\alpha}}$.*

Proof. The claim is true for $\hat{\alpha} = 1$ because RPoA$(\alpha) \leq \infty$ holds trivially.

Let Γ be a fair cost sharing game, given by (n, m, Σ, d) (conform Definition 22). Let $a_j, j \in [m]$ denote the costs of the facilities (conform Definition 23). We show that Γ^α is $(n, \hat{\alpha}, \alpha)$-smooth for $\alpha \in [0, 1)^n$. For $s \in \Sigma$ we define

$$B(s) = \left\{ j \;\middle|\; j \in \bigcup_{i \in [n]} s_i \right\}$$

Let $s, s^* \in \Sigma$ be two strategy profiles. Fix an arbitrary player $i \in [n]$. We have

$$C(s_i^*, s_{-i}) - C(s) = \sum_{j \in B(s_i^*, s_{-i})} a_j - \sum_{j \in B(s)} a_j \leq \sum_{j \in s_i^* \setminus B(s)} a_j.$$

We use this inequality to obtain the following bound. We remind the reader that we write $P_j(s)$ to denote the set of players choosing facility $j \in [m]$ under strategy profile $s \in \Sigma$.

$$\begin{aligned}
&(1 - \alpha_i) c_i(s_i^*, s_{-i}) + \alpha_i (C(s_i^*, s_{-i}) - C(s)) \\
&\quad \leq (1 - \alpha_i) \sum_{j \in s_i^*} \frac{a_j}{P_j(s_i^*, s_{-i})} + \alpha_i \sum_{j \in s_i^* \setminus B(s)} \frac{a_j}{P_j(s_i^*, s_{-i})} \\
&\quad \leq \sum_{j \in s_i^*} \frac{a_j}{P_j(s_i^*, s_{-i})} \\
&\quad \leq \sum_{j \in s_i^*} \frac{n a_j}{P_j(s^*)}.
\end{aligned}$$

The first inequality holds because $P_j(s_i^*, s_{-i}) = 1$ for every $j \in s_i^* \setminus B(s)$, and the last inequality follows from $P_j(s_i^*, s_{-i}) \geq P_j(s^*)/n$ for every $j \in s_i^*$. The left-hand side of the smoothness condition (4.1) is equivalent to

$$\begin{aligned}
&\sum_{i \in [n]} \left((1 - \alpha_i) c_i(s_i^*, s_{-i}) + \alpha_i (C(s_i^*, s_{-i}) - C(s)) + \alpha_i c_i(s) \right) \\
&\quad \leq \sum_{i \in [n]} \left(\sum_{j \in s_i^*} \frac{n a_j}{P_j(s^*)} \right) + \hat{\alpha} C(s) \\
&\quad = n C(s^*) + \hat{\alpha} C(s).
\end{aligned}$$

We conclude that the robust price of anarchy is at most $\frac{n}{1 - \hat{\alpha}}$. Example 82 shows that this bound is tight, even for pure equilibria with uniform altruism. $\square$

Example 82. Let $\alpha \in [0,1]^n$ be a uniform altruism vector. Consider a cost-sharing game with n players and 2 facilities, where $a_1 = 1$ and $a_2 = n/(1-\alpha)$. The strategy set of each players consists of the two singleton sets. Let $s^* = (e_1, \ldots, e_1)$ and $s = (e_2, \ldots, e_2)$. Then $C(s^*) = 1$ and $C(s) = n/(1-\alpha)$. Note that s is a pure Nash equilibrium of the α-altruistic extension of this game, because for every player i,

$$(1-\alpha_i)c_i(s) + \alpha_i C(s) = 1 + \alpha_i \frac{n}{1-\alpha_i} = c_i^\alpha(\{1\}, s_{-i}).$$

The pure price of anarchy of this game is therefore at least $n/(1-\alpha)$.

We turn to the pure price of stability of uniformly α-altruistic extensions of cost-sharing games. Clearly, an upper bound on the pure price of stability extends to the mixed, correlated and coarse price of stability. As opposed to the price of anarchy, the price of stability does improve with increased altruism. The proof of the following proposition exploits a standard technique to bound the pure price of stability of exact potential games (see, e.g., Nisan et al. [2007]).

Proposition 83. *Let Γ be a fair cost sharing game, let $\alpha \in [0,1]^n$ be a uniform altruism vector, and let C be the sum-of-costs social cost function (1.3.1.8) for Γ. The pure price of stability of the uniformly α-altruistic extension Γ^α of Γ with respect to C is at most $(1-\check{\alpha})H_n + \check{\alpha}$, where H_n denotes the nth harmonic number $\sum_{k\in[n]} 1/k$.*

Proof. It is not hard to verify that Γ^α is a potential game with potential function $\Phi^\alpha(s) = (1-\check{\alpha})\Phi(s) + \check{\alpha}C(s)$, where $\Phi(s) = \sum_{j\in[m]} \sum_{k\in[|P_j(s)|]} a_j/k$. Observe that for $s \in \Sigma$,

$$\begin{aligned}
\Phi^\alpha(s) &= (1-\check{\alpha}) \sum_{j\in[m]} \sum_{k\in[|P_j(s)|]} \frac{a_j}{k} + \check{\alpha} \sum_{j\in B(s)} a_j \\
&\leq ((1-\check{\alpha})H_n + \check{\alpha}) \sum_{j\in B(s)} a_j \\
&= ((1-\check{\alpha})H_n + \check{\alpha})C(s).
\end{aligned}$$

We therefore have that $C(s) \leq \Phi^\alpha(s) \leq ((1-\check{\alpha})H_n + \check{\alpha})C(s)$.

Let $s \in \Sigma$ be a strategy profile that minimizes Φ^α, and let $s^* \in \Sigma$ be a strategy profile that minimizes C. Note that s is a pure equilibrium of Γ^α. We have

$$C(s) \leq \Phi^\alpha(s) \leq \Phi^\alpha(s^*) \leq ((1-\check{\alpha})H_n + \check{\alpha})C(s^*),$$

which proves the claim. $\square$

4.5 Valid Utility Games

Theorem 84. *Let Γ be a valid utility game and let $\alpha \in [0,1]^n$. The α-robust price of anarchy of Γ with respect to any valid social welfare function, is at most 2. This bound is tight, i.e., there exists a valid utility game Γ and valid social welfare function U such that the α-robust price of anarchy of Γ with respect to U is 2.*

Proof. We show that a valid utility game $\Gamma = (n, \Sigma, u)$ is $(1,1,\alpha)$-smooth with respect to a valid social welfare function U.

Fix two strategy profiles $s, s^* \in \Sigma$ and consider an arbitrary player $i \in [n]$. By assumption, we have $u_i(s) \geq U(s) - V\left(\bigcup_{i' \in [n]\setminus\{i\}} s_{i'}\right)$, or equivalently,

$$-U(s) + u_i(s) \geq -V\left(\bigcup_{i' \in [n]\setminus\{i\}} s_{i'}\right)$$

Now let $S_i = \bigcup_{i' \in [n]} s_{i'} \cup \bigcup_{i' \in [i]} s^*_{i'}$. Summing over all $i \in [n]$,

$$\begin{aligned}
&\sum_{i \in [n]} ((1-\alpha_i) u_i(s_i^*, s_{-i}) + \alpha_i (U(s_i^*, s_{-i}) - U(s) + u_i(s))) \\
&\quad \geq \sum_{i \in [n]} \left(U(s_i^*, s_{-i}) - V\left(\bigcup_{i' \in [n]\setminus\{i\}} s_{i'}\right)\right) \\
&\quad = \sum_{i \in [n]} \left(V\left(s_i^* \cup \bigcup_{i' \in [n]\setminus\{i\}} s_i\right) - V\left(\bigcup_{i' \in [n]\setminus\{i\}} s_{i'}\right)\right) \\
&\quad \geq \sum_{i \in [n]} (V(S_i) - V(S_{i-1})) \\
&\quad \geq U(s^*) - U(s).
\end{aligned}$$

Here, the first inequality follows from (4.5) and

$$u_i(s_i^*, s_{-i}) \geq U(s_i^*, s_{-i}) - V\left(\bigcup_{i' \in [n]\setminus\{i\}} s_{i'}\right)$$

for all $i \in [n]$. The second inequality holds because V is submodular, and the final inequality follows from V being non-decreasing. We conclude that Γ^α is $(1,1,\alpha)$-smooth, which proves an upper bound of 2 on the α-robust price of anarchy of Γ with respect to U. This bound is tight, as shown by Example 85. $\square$

Example 85. Consider a valid utility game with $n = 2$ and $m = 2$. The strategy sets are $\Sigma_1 = \{\{1\}, \{2\}\}$, $\Sigma_2 = \{\varnothing, \{1\}\}$. Define $V(S) = |S|$ for every subset $S \subseteq [2]$. Note that V is non-negative, non-decreasing and submodular. Of course, we define $U(s)$ as $V(s_1 \cup s_2)$ for all $s \in \Sigma$.

For a given strategy profile $s \in \Sigma$, the utility functions $u_1(s)$ and $u_2(s)$ are defined as follows: $u_1(s) = 1$ for all strategy profiles s. $u_2(s) = 1$ if $s = (\{2\}, \{1\})$ and $u_2(s) = 0$ otherwise. Observe that U is sum-bounded. Moreover, it is not hard to verify that for every player i and every strategy profile $s \in \Sigma$ we have $u_i(s) \geq U(s) - V(s_{-i})$. We conclude that Γ is a valid utility game and U a valid social welfare function.

Let $\alpha \in [0, 1]^2$, and consider the α-altruistic extension Γ^α of Γ. We claim that $s = (\{1\}, \varnothing)$ is a pure equilibrium of Γ^α:It holds that $u_1(s) = (1 - \alpha_1) + \alpha_1 = 1$. This utility remains 1 if player 1 switches to strategy $\{2\}$. $u_2(s) = \alpha_2$. If player 2 switches to strategy $\{1\}$, then its profit is α_2 as well. Thus, s is a pure equilibrium. Since $U(s) = 1$ and $U(\{2\}, \{1\}) = 2$, the pure price of anarchy of Γ with respect to U is 2.

4.6 Linear Congestion Games

Pure equilibria of altruistic extensions of linear congestion games always exist [Hoefer and Skopalik, 2009b]. This may not be the case for arbitrary (non-linear) congestion games.

It has been shown by [Christodoulou and Koutsoupias, 2005a] that the pure price of anarchy of linear congestion games is $5/2$. Recently, Caragiannis et al. [2010] extended this result to uniformly altruistic extensions of linear congestion games. Applying the transformation outlined in Remark 76, their result can be stated as follows:

Theorem 86 (Caragiannis et al. [2010]). *Let Γ be a linear congestion game, let $\alpha \in [0, 1]^n$ be a uniform altruism vector, and let C be the sum-of-costs social cost function (1.3.1.8) for Γ. The pure price of anarchy of the α-altruistic extension Γ^α of Γ with respect to C is at most $(5 + 4\bar{\alpha})/(2 + \bar{\alpha})$, where $\bar{\alpha}$ is the uniform altruism level $\alpha_i, i \in [n]$.*

The proof in Caragiannis et al. [2010] implicitly uses a smoothness argument conform the smoothness framework we defined for altruistic games. Thus, without any additional work, our framework allows the generalization of Theorem 86 to the α-robust price of anarchy, for uniform altruism vectors α. Caragiannis et al. [2010] also showed that the bound of Theorem 86 is asymptotically tight. A simpler example given below (Example 92) proves tightness of this bound (not only asymptotically). Thus, the robust price of anarchy for uniformly α-altruistic extensions of linear congestion games is *exactly* $(5+4\bar{\alpha})/(2+\bar{\alpha})$. We give a refinement of Theorem 86 to non-uniform altruism distributions, obtaining a bound in terms of the maximum and minimum altruism levels.

Theorem 87. *Let* Γ *be a linear congestion game, let* $\alpha \in [0,1]^n$*, and let* C *be the sum-of-costs social cost function (1.3.1.8) for* Γ*. The* α*-robust price of anarchy of a linear congestion game with respect to* C *is at most* $(5+2\hat{\alpha}+2\check{\alpha})/(2-\hat{\alpha}+2\check{\alpha})$.

As a first step, we show that without loss of generality we can focus on simpler instances of linear congestion games.

Definition 88. *A simple linear congestion game* is a linear congestion game $\Gamma = (n, m, \Sigma, d)$ such that $d_j(k) = k$ for all $k \in [n]$.

Lemma 89. *Let* $\Gamma = (n, m, \Sigma, d)$ *be a linear congestion game. Then there exists a simple linear congestion game* $\Gamma' = (n, m', \Sigma, d')$ *and a bijection* $f_i : \Sigma_i \to \Sigma'_i$ *for every player* $i \in [n]$ *such that* $c_i^{\Gamma}(s_1, \ldots, s_n) = c_i^{\Gamma'}(d_1(s_1), \ldots, d_n(s_n))$*, where for* $i \in [n]$*,* c_i^{Γ} *denotes the cost function of* i *in* Γ *and* $c_i^{\Gamma'}$ *denotes the cost function of* i *in* Γ'.

Proof. We will show how to transform Γ into a simple linear congestion game that possesses the required properties.

First, we may assume that for every delay function $d_j, j \in [m]$, the numbers a_j and b_j are integers. This can be ensured by multiplying all coefficients among all facilities by their least common multiple. In the resulting game, all coefficients are integers, the price of anarchy is the same, and so is the set of all equilibria.

Next, we can transform Γ into a game such that $b_j = 0$ for all $j \in m$. To show this, we replace any facility $j \in [m]$ by $n+1$ facilities $j_0, \ldots, j_n$ with delay functions $d_{j_0}(k) = a_j k$ and $d_{j_i}(k) = b_j k$ for $i \in [n]$. We then adapt the strategy space Σ_i of each player i as follows: we replace every strategy $s_i \in \Sigma_i$ in which j occurs by the strategy $s_i \setminus \{j\} \cup \{j_0, j_i\}$. Now there is an obvious bijection between the strategy profiles in the original game and those in the new game, preserving the values of individual cost functions.

Finally, for similar reasons, we can transform our linear congestion game such that $a_j = 1$ for all $j \in [m]$: We replace j with facilities $j_1, \ldots, j_{a_j}$, where $d_{j_i}(k) = k, k \in [n], j \in [a_j]$. We adapt the strategy space Σ_i of each player $i \in [n]$ by replacing each strategy $s_i \in \Sigma_i$ in which j occurs by $s_i \setminus \{j\} \cup \{j_1, \ldots, a_{a_j}\}$.

The resulting game is a simple linear congestion game. □

Therefore, it suffices to bound the α-robust price of anarchy of simple linear congestion games. The next step in the proof of Theorem 87 is the following technical lemma:

Lemma 90. *For all* $x, y \in \mathbb{N}_{\geq 0}$ *and* $\hat{\alpha}, \check{\alpha} \in [0,1]$ *with* $\hat{\alpha} \geq \check{\alpha}$*,*

$$((1+\hat{\alpha})x+1)y + \check{\alpha}(1-x)x \leq \frac{5+2\hat{\alpha}+2\check{\alpha}}{3}y^2 + \frac{1+\hat{\alpha}-2\check{\alpha}}{3}x^2.$$

To prove this lemma, we make use of the following result:

Lemma 91. *For all* $x, y \in \mathbb{N}_{\geq 0}$*,* $\beta, \gamma \in [0,1]$ *it holds that*

$$((1+\beta)x+1)y + \gamma\beta(1-x)x \leq (2+\beta-\delta)y^2 + \delta x^2$$

for all $\delta \in [(1/3)(1+\beta-2\gamma\beta), 1+\beta]$.

Proof. The inequality is equivalent to

$$((1+\beta)x+1)y + \gamma\beta(1-x)x - (2+\beta)y^2 \leq \delta(x^2-y^2).$$

Assume that $x = y$. The inequality is then trivially satisfied because $x \leq x^2$ for all $x \in \mathbb{N}_{\geq 0}$. Next suppose that $x > y$. Then

$$\delta \geq \frac{((1+\beta)x+1)y + \gamma\beta(1-x)x - (2+\beta)y^2}{x^2-y^2}.$$

We show that the maximum of the expression on the right-hand side is attained by $x = 2$ and $y = 1$. First, we fill in these values and conclude that for these values, $\delta \geq (1/3)(1+\beta-2\gamma\beta) \geq 0$. We now write x as $y+a, a \geq 1$, and rewrite the right-hand side as

$$f(y,a) = \frac{(1+\beta)y+\gamma\beta}{2y+a} + \frac{(1+\gamma\beta)(y-y^2)}{a(2y+a)} - \gamma\beta. \tag{4.4}$$

Because we know that there are choices of y and a for which $f(y,a)$ is positive (e.g., when $y = 1$ and $a = 1$), and because a only occurs in the denominators, we know that (4.4) reaches its maximum when $a = 1$. So we assume $a = 1$. When we then fill in $y = 0$, we see that $f(0,1) = 0$, so $f(1,1) \geq f(0,1)$. When $y > 1$ we write y as $w+2$, where $w \geq 0$, and we further rewrite $f(y,a)$ as

$$f(w+2,1) = \frac{2\beta-6\gamma\beta}{2w+5} - \frac{(2-\beta+5\gamma\beta)w + (1+\gamma\beta)w^2}{2w+5} \leq \frac{2\beta-6\gamma\beta}{2w+5}.$$

When $2\beta-6\gamma\beta$ is negative, this term is certainly less than $f(1,1)$. When $2\beta-6\gamma\beta$ is positive, we have

$$f(w+2,1) \leq \frac{2\beta-6\gamma\beta}{2w+5} \leq \frac{2\beta-6\gamma\beta}{5} \leq \frac{1}{3}(2\beta-6\gamma\beta) \leq \frac{1}{3}(1+\beta-2\gamma\beta) = f(1,1).$$

This shows that $\delta \geq f(1,1) = \frac{1}{3}(1+\beta-2\gamma\beta)$.

The final case is when $x < y$. Then,

$$\delta \leq \frac{(2+\beta)y^2 - ((1+\beta)x+1)y - \gamma\beta(1-x)x}{y^2 - x^2}.$$

We show that the minimum of the expression on the right-hand side is attained by $x = 0$ and $y = 1$. First, we fill in these values and conclude that for these values, $\delta \leq 1+\beta$. We now write y as $x + a$, $a \geq 1$, and rewrite the right-hand side as

$$g(x,a) = \frac{(1+\gamma\beta)x^2 - (1+a+(a+\gamma)\beta)x - a}{a(2x+a)} + 2 + \beta.$$

Suppose first that $x = 0$ and that $a \geq 2$. Then we can write a as $1+b$, $b > 0$, and therefore

$$g(0,1+b) = 2+\beta - \frac{1}{1+b} \geq \frac{3}{2} + \beta \geq 1+\beta = f(0,1).$$

When $x \geq 1$, we can write x as $1+b$, $b \geq 0$. We then have

$$g(1+b,a) = 2+\beta - \frac{2+\beta+(1-\beta)b}{2b+2+a} + \frac{(1+\gamma\beta)(b^2+b)}{a(2b+2+a)}.$$

The last of these terms is positive, hence

$$\begin{aligned} g(1+b,a) &\geq 2+\beta - \frac{2+\beta+(1-\beta)b}{2b+2+a} \geq 2+\beta - \frac{2+1+b}{2b+2+a} \\ &\geq 2+\beta-1 = 1+\beta = f(0,1). \end{aligned}$$

This shows that $\delta \leq f(0,1) = 1+\beta$. □

Now we can complete the proof of Lemma 90.

Proof of Lemma 90. Choose $\gamma \in [0,1]$ such that $\check{\alpha} = \gamma\hat{\alpha}$. Using Lemma 91 above, we obtain

$$((1+\hat{\alpha})x+1)y+\check{\alpha}(1-x)x = ((1+\hat{\alpha})x+1)y+\gamma\hat{\alpha}(1-x)x \leq (2+\hat{\alpha}-\delta)y^2+\delta x^2,$$

for $\delta \in [(1/3)(1+\hat{\alpha}-2\gamma\hat{\alpha}), 1+\hat{\alpha}]$. By choosing $\delta = (1/3)(1+\hat{\alpha}-2\gamma\hat{\alpha})$, we obtain

$$((1+\hat{\alpha})x+1)y + \check{\alpha}(1-x)x \leq \frac{5+2\hat{\alpha}+2\gamma\hat{\alpha}}{3}y^2 + \frac{1+\hat{\alpha}-2\gamma\hat{\alpha}}{3}x^2.$$

Substituting $\gamma\hat{\alpha} = \check{\alpha}$ yields the claim. □

We remark that the choice of δ in the proof above has been made in order to minimize the expression $\lambda/(1-\mu)$ (which is an increasing function of δ).

Lemma 90 is essentially the part that generalizes the proof of Caragiannis et al. [2010], and allows us to complete the proof of Theorem 87.

Proof of Theorem 87. Let $\Gamma = (n, m, \Sigma, d)$ be a linear congestion game. We show that Γ is $((1/3)(5 + 2\hat{\alpha} + 2\check{\alpha}), (1/3)(1 + \hat{\alpha} - 2\check{\alpha}), \alpha)$-smooth. By Lemma 89, we may assume without loss of generality that Γ is a simple linear congestion game.

Let $s, s^* \in \Sigma$. The left-hand side of the smoothness condition (4.1) is equivalent to

$$
\begin{aligned}
&\sum_{i\in[n]} ((1-\alpha_i)c_i(s_i^*, s_{-i}) + \alpha_i(C(s_i^*, s_{-i}) - C(s)) + \alpha_i c_i(s)) \\
&\quad = \sum_{i\in[n]} \left((1-\alpha_i) \left(\sum_{j\in s_i^*\setminus s_i} (|P_j(s)|+1) + \sum_{j\in s_i\cap s_i^*} |P_j(s)| \right) \right. \\
&\qquad \left. + \alpha_i \sum_{j\in[m]} (|P_j(s_i^*, s_{-i})|^2 - |P_j(s)|^2) + \alpha_i c_i(s) \right) \\
&\quad = \sum_{i\in[n]} \left((1-\alpha_i) \left(\sum_{j\in s_i^*\setminus s_i} (|P_j(s)|+1) + \sum_{j\in s_i\cap s_i^*} |P_j(s)| \right) \right. \\
&\qquad \left. + \alpha_i \sum_{j\in(s_i^*\setminus s_i)\cup(s_i\setminus s_i^*)} (|P_j(s_i^*, s_{-i})|^2 - |P_j(s)|^2) + \alpha_i c_i(s) \right) \\
&\quad = \sum_{i\in[n]} \left((1-\alpha_i) \left(\sum_{j\in s_i^*\setminus s_i} (|P_j(s)|+1) + \sum_{j\in s_i\cap s_i^*} |P_j(s)| \right) \right. \\
&\qquad + \alpha_i \left(\sum_{j\in s_i^*\setminus s_i} ((|P_j(s)|+1)^2 - |P_j(s)|^2) \right. \\
&\qquad \left. \left. + \sum_{j\in s_i\setminus s_i^*} ((|P_j(s)|-1)^2 - |P_j(s)|^2) \right) + \alpha_i c_i(s) \right) \\
&\quad = \sum_{i\in[n]} \left((1-\alpha_i) \left(\sum_{j\in s_i^*\setminus s_i} (|P_j(s)|+1) + \sum_{j\in s_i\cap s_i^*} |P_j(s)| \right) \right.
\end{aligned}
$$

$$+\alpha_i \left(\sum_{j \in s_i^* \setminus s_i} (2|P_j(s)|+1) + \sum_{j \in s_i \setminus s_i^*} (1-2|P_j(s)|) \right) + \alpha_i c_i(s) \Bigg)$$
$$\leq \sum_{i \in [n]} \left(\sum_{j \in s_i^*} ((1+\alpha_i)|P_j(s)|+1) + \alpha_i \sum_{j \in s_i} (1-|P_j(s)|) \right)$$
$$\leq \sum_{j \in [m]} \left(((1+\hat{\alpha})|P_j(s)|+1)|P_j(s^*)| + \check{\alpha}(1-|P_j(s)|)|P_j(s)| \right).$$

In the above derivation, the first inequality follows from the fact that

$$(1-\alpha_i)|P_j(s)| \leq (1+\alpha_i)|P_j(s)| + 1 + \alpha_i(1-2|P_j(s)|)$$

for every $j \in s_i \cap s_i^*, i \in [n]$. Therefore, it is possible to replace $(1-\alpha_i)|P_j(s)|$ (in the third summation operator of the left hand side of the first inequality) by $(1+\alpha_i)|P_j(s)| + 1 + \alpha_i(1-2|P_j(s)|)$, write $c_i(s)$ as $\sum_{j \in s_i} |P_j(s)|$, and finally rewrite the resulting expression into the form of the right hand side of the first inequality. The second inequality holds because for every $i \in [n]$ and $j \in s_i$, $1 - |P_j(s)| \leq 0$ and by the definition of $\hat{\alpha}$ and $\check{\alpha}$. The bound on the robust price of anarchy now follows from Lemma 90, letting $x = |P_j(s)|$ and $y = |P_j(s^*)|$. $\square$

The following is a simple example that shows that the bound of $\frac{5+2\hat{\alpha}+2\check{\alpha}}{2-\hat{\alpha}+2\check{\alpha}}$ on the robust price of anarchy for *uniformly* α-altruistic linear congestion games is tight, even for pure equilibria. It improves the lower bound example of Caragiannis et al. [2010], because it is simpler and it shows tightness of the bound not only asymptotically.

Example 92. Consider a game $\Gamma = (3, 6, \Sigma, d)$. The facility set $[6]$ is partitioned into two sets E_1 and E_2, of three facilities each. Denote the three facilities in S_1 by h_0, h_1, and h_2. Denote the three facilities in S_2 by g_0, g_1, and $g_2\}$. Let $\alpha \in [0,1]^3$ be any uniform altruism vector, and let $\bar{\alpha}$ be the uniform altruism level of the players, i.e., $\bar{\alpha} = \alpha_i$ for $i \in [3]$. The delay functions are given by $d_j(k) = (1+\bar{\alpha})k$ for $j \in E_1, k \in [n]$, and $d_j(k) = k$ for $j \in E_2, k \in [n]$. We have $\Sigma_i = \{S_i^1 = \{h_{i-1}, g_{i-1}\}, S_i^2 = \{h_{(i-2) \text{ (mod 3)}}, h_{i \text{ (mod 3)}}, g_{i \text{ (mod 3)}}\}\}$. The strategy profile (S_1^1, S_2^1, S_3^1) is a social optimum of cost $3(1+\bar{\alpha}) + 3 = 3(2+\bar{\alpha})$.

We argue that the strategy profile (S_1^2, S_2^2, S_3^2) is a pure equilibrium for the α-altruistic extension Γ^α of Γ. In Γ^α, each player's perceived cost is $(1-\bar{\alpha})(4(1+\bar{\alpha})+1) + 3\bar{\alpha}(5+4\bar{\alpha})$. If a player switches to its first strategy, the new social cost would become $3(5+4\bar{\alpha}) + 1 - \bar{\alpha}$, so that the player's new perceived individual cost is $(1-\bar{\alpha})(3(1+\bar{\alpha})+2) + 3\bar{\alpha}((5+4\bar{\alpha})+1-\bar{\alpha})$, which is not an improvement over its old social cost. So s is a pure equilibrium, of cost $12(1+\bar{\alpha}) + 3 = 3(5+4\bar{\alpha})$. We conclude that the pure price of anarchy of Γ^α with respect to the sum-of-costs social cost function (1.3.1.8)is at least $(5+4\bar{\alpha})/(2+\bar{\alpha})$ for $\alpha \in [0,1]$.

We turn to the pure price of stability of altruistic extensions of linear congestion games. Again, note that an upper bound on the pure price of stability extends to the mixed, correlated and coarse price of stability.

Proposition 93. *Let $\Gamma = (n, m, \Sigma, d)$ be a linear congestion game, let $\alpha \in [0,1]^n$ be a uniform altruism vector, and let C be the sum-of-costs social cost function (1.4) for Γ. The pure price of stability of the α-altruistic extension Γ^α of Γ with respect to C is at most $\frac{2}{1+\breve{\alpha}}$.*

Proof. By Lemma 89, we may assume without loss of generality that Γ is a simple linear congestion game. It is not hard to verify that Γ^α is an exact potential game with potential function $\Phi^\alpha(s) = (1-\breve{\alpha})\Phi(s)+\breve{\alpha}C(s)$, where $\Phi(s) = \sum_{j\in[m]}\sum_{i\in[|P_j(s)|]} i$. Observe that

$$\begin{aligned}\Phi^\alpha(s) &= (1-\breve{\alpha})\sum_{j\in[m]}\sum_{i\in[|P_j(s)|]} i + \breve{\alpha}C(s)\\ &= \frac{1-\breve{\alpha}}{2}\sum_{j\in[m]}(|P_j(s)|^2+|P_j(s)|) + \breve{\alpha}\sum_{j\in[m]}|P_j(s)|^2\\ &= \frac{1+\breve{\alpha}}{2}C(s) + \frac{1-\breve{\alpha}}{2}\sum_{j\in[m]}P_j(s).\end{aligned}$$

Therefore, $\frac{1+\breve{\alpha}}{2}C(s) \le \Phi^\alpha(s) \le C(s)$.

Let $s \in \Sigma$ be a strategy profile that minimizes Φ^α, and let $s^* \in \Sigma$ be a strategy profile that minimizes C. Note that s is a pure equilibrium of Γ^α. We have

$$C(s^*) \ge \Phi^\alpha(s^*) \ge \Phi^\alpha(s) \ge \frac{1+\breve{\alpha}}{2}C(s),$$

which proves the claim. □

4.7 Symmetric Singleton Linear Congestion Games

For symmetric singleton linear congestion games, we distinguish between the cases of uniform altruism levels and non-uniform altruism.

4.7.1 Uniform Altruism

Caragiannis et al. [2010] prove the following theorem (stated using the transformation described in Remark 76). It shows that the pure price of anarchy does not always increase with the altruism level. The relationship between α and the price of anarchy is thus rather subtle.

Theorem 94 (Caragiannis et al. [2010]). *Let $\Gamma = (n, m, \Sigma, d)$ be a symmetric singleton linear congestion game, let $\alpha \in [0,1]^n$ be a uniform altruism vector, and let C be the sum-of-costs social cost function (1.3.1.8). The pure price of anarchy of the uniformly α-altruistic extension Γ^α of Γ with respect to C is $4/(3+\bar{\alpha})$, where $\bar{\alpha}$ is the uniform altruism level $\alpha_i, i \in [n]$.*

We show that even the mixed price of anarchy (and thus also the robust price of anarchy) will be at least 2 regardless of the altruism levels of the players, by generalizing the proof of Theorem 5.4 of Lücking et al. [2008]. This implies that for uniform altruism, the benefits of higher altruism in singleton congestion games are only reaped in pure equilibria, and the gap between the pure and mixed price of anarchy increases in $\check{\alpha}$. Also it shows that singleton congestion games constitute a class of games for which a smoothness proof cannot deliver tight bounds on the pure price of anarchy.

Proposition 95. *Let Γ be a symmetric singleton linear congestion game, let $\alpha \in [0,1]^n$, and let C be the sum-of-costs social cost function (1.3.1.8). The mixed price of anarchy of the α-altruistic extension Γ^α of Γ with respect to C is at least* 2.

Proof. Let $m \in \mathbb{N}_{\geq 2}$ and consider the symmetric singleton linear congestion game with player set $[m]$ and facility set $[m]$, with $d_j(k) = k$ for $j, k \in [m]$. Denote by σ the mixed strategy profile where each player chooses each facility with probability $1/m$. It is straightforward to verify that $\mathbf{E}_{s\sim\sigma}[C(s)] = 2 - 1/m$. Moreover, the strategy profile s^* where each player $i \in [m]$ plays facility i is a minimizer for C, and $C(s^*) = m$. It therefore suffices to show that σ is a Nash equilibrium of Γ^α.

By symmetry, it suffices to show that the expected cost of player 1 increases if it deviates to the strategy where it chooses facility 1 with probability 1. Let $s_1^* = 1$. We have

$$\begin{aligned}\mathbf{E}_{s\sim\sigma}[c_1^\alpha(s_1^*, s_{-1})] &= \mathbf{E}_{s\sim\sigma}[(1-\alpha_1)c_1(s_1^*, s_{-1}) + \alpha_1 C(s_1^*, s_{-1})] \\ &= (1-\alpha_1)\mathbf{E}_{s\sim\sigma}[c_1(s_1^*, s_{-1})] + \alpha_1 \mathbf{E}_{s\sim\sigma}[C(s_1^*, s_{-1})].\end{aligned}$$

In Theorem 5.4 of Lücking et al. [2008] is shown that s is a Nash equilibrium when $\alpha = \mathbf{0}$, so it remains to show that

$$\mathbf{E}_{s\sim\sigma}[C(s_1^*, s_{-1})] \geq \mathbf{E}_{s\sim\sigma}[C(s)] = 2m - 1.$$

For $i, j \in [m]$, let $X_{i,j} : \Sigma \to \{0, 1\}$ be the indicator function that maps a strategy profile $s' \in \Sigma$ to 1 iff player i chooses facility j under s'. Then it is clear that $c_i(s') = \sum_{j\in[m]} X_{i,j}(s')d_j(s')$ and $d_j(s') = \sum_{i\in[m]} X_{i,j}(s')$ for all $i, j \in [m], s' \in \Sigma$. So $c_i(s') = \sum_{i,i',j\in[m]} X_{i,j}(s')X_{i',j}(s')$ for all $i \in [m], s \in \Sigma$. Using this last identity, along with symmetry, stochastic independence, and linearity of expectation, we derive the following:

$$
\begin{aligned}
\mathbf{E}_{s\sim\sigma}[C(s_1^*, s_{-1})] &= \sum_{i\in[m]} \mathbf{E}_{s\sim\sigma}[c_i(s_1^*, s_{-1})] \\
&= \mathbf{E}_{s\sim\sigma}[c_1(s_1^*, s_{-1})] + (m-1)\mathbf{E}_{s\sim\sigma}[c_2(s_1^*, s_{-1})] \\
&= \mathbf{E}_{s\sim\sigma}[d_1(s_1^*, s_{-1})] + (m-1) \sum_{i',j\in[m]} \mathbf{E}_{s\sim\sigma}[X_{2,j}(s_1^*, s_{-1})X_{i',j}(s_1^*, s_{-1})] \\
&= \sum_{i\in[m]} \mathbf{E}_{s\sim\sigma}[X_{i,1}(s_1^*, s_{-1})] \\
&\quad + (m-1)\left(\sum_{i'\in[m]} \mathbf{E}_{s\sim\sigma}[X_{2,1}(s_1^*, s_{-1})X_{i',1}(s_1^*, s_{-1})] \right. \\
&\quad \left. +(m-1) \sum_{i'\in[m]} \mathbf{E}_{s\sim\sigma}[X_{2,2}(s_1^*, s_{-1})X_{i',2}(s_1^*, s_{-1})] \right) \\
&= \left(1 + (m-1)\frac{1}{m}\right) + (m-1)\left(\frac{1}{m} + \frac{1}{m} + (m-2)\frac{1}{m^2}\right. \\
&\quad \left. +(m-1)\left(0 + \frac{1}{m} + (m-2)\frac{1}{m^2}\right)\right) \\
&= 2m - 1.
\end{aligned}
$$

□

We prove next that when every player in a symmetric singleton linear congestion game is completely altruistic, the pure price of anarchy is 1 for a broader class of delay functions, namely *semi-convex* delay functions. This is a corollary of the following lemma, which is stated for $(0, 1)$-altruism vectors. The reason for this is that the lemma is reused in the next section.

Definition 96. For a congestion game $\Gamma = (n, m, \Sigma, d)$, a delay function d_j of a facility $j \in [m]$ is called *semi-convex* iff $kd_j(k) - (k-1)d_j(k-1) \geq \ell d_j(\ell) - (\ell - 1)d_j(\ell - 1)$ for $k, \ell \in [n], k \geq \ell$.

Lemma 97. *Let $\Gamma = (n, m, \Sigma, d)$ be a symmetric singleton congestion game, let $\alpha \in \{0,1\}^n$, let C be the sum-of-costs social cost function (1.3.1.8), let $s \in \Sigma$ be a pure equilibrium, let*

$$S_1 = \{j \in [m] \mid \exists i \in [n] : \alpha_i = 1, s_i = \{j\}\},$$

and assume that all delay functions d are semi-convex. Then there is strategy profile $s^* \in \Sigma$ that minimizes C, such that $|P_j(s)| \leq |P_j(s^*)|$ for every facility $j \in S_1$.

Proof. Let $s^* \in \Sigma$ be a strategy profile minimizing C, and assume that $|P_j(s^*)| < |P_j(s)|$ for some $j \in S_1$. Then there is a facility $j' \in [m]$ with $|P_{j'}(s^*)| > |P_{j'}(s)|$. Consider a player $i \in P_1$ with $s_i = \{j\}$ and $\alpha_i = 1$. (Note that i exists by the definition of S_1.) Because s is a pure equilibrium, player i has no incentive to change its strategy from j to j', i.e., $C(\{j'\}, s_{-i}) \geq C(s)$, or, equivalently,

$$\begin{aligned}(|P_{j'}(s)| + 1)d_{j'}(|P_{j'}(s) + 1) - |P_{j'}(s)|d_{j'}(|P_{j'}(s)|)\\ \geq |P_j(s)|d_j(|P_j(s)|) - (|P_j(s)| - 1)d_j(|P_j(s)| - 1).\end{aligned} \tag{4.5}$$

Since $|P_j(s^*)| < |P_j(s)|$ and $|P_{j'}(s)| < |P_{j'}(s^*)|$, the semi-convexity of the delay functions implies

$$\begin{aligned}(|P_j(s^*)| + 1)d_j(|P_j(s^*)| + 1) - |P_j(s^*)|d_j(|P_j(s^*)|)\\ \leq |P_j(s)|d_j(|P_j(s)|) - (|P_j(s)| - 1)d_j(|P_j(s)| - 1),\end{aligned} \tag{4.6}$$

and

$$\begin{aligned}(|P_{j'}(s)| + 1)d_{j'}(|P_{j'}(s)| + 1) - |P_{j'}(s)|d_{j'}(|P_{j'}(s)|)\\ \leq |P_{j'}(s^*)|d_{j'}(|P_{j'}(s^*)|) - (|P_{j'}(s^*)| - 1)d_{j'}(|P_{j'}(s^*)| - 1).\end{aligned} \tag{4.7}$$

By combining (4.5), (4.6) and (4.7) and re-arranging terms, we obtain

$$\begin{aligned}(|P_j(s^*)| + 1)d_j(|P_j(s^*)| + 1) + (|P_{j'}(s^*)| - 1)d_{j'}(|P_{j'}(s^*)| - 1)\\ \leq |P_j(s^*)|d_e(|P_j(s^*)|) + |P_{j'}(s^*)|d_{j'}(|P_{j'}(s^*).\end{aligned}$$

The above inequality implies that by moving a player i with $s^*_i = \{j'\}$ from j' to j, we obtain a new strategy profile $s' = (\{j\}, s^*_{-i})$ of cost $C(s') \leq C(s^*)$. (Note that i must exist because $|P_{j'}(s^*)| > |P_{j'}(s)| \geq 0$.) Moreover, the number of players on j under the new strategy profile s' is increased by one. We can therefore repeat the above argument (with s' in place of s^*) until we obtain an optimal strategy profile that satisfies the claim. □

Corollary 98. *With respect to the sum-of-costs social cost function (1.3.1.8), the pure price of anarchy of $\mathbf{1}$-altruistic extensions of symmetric singleton congestion games with semi-convex delay functions, is 1.*

4.7.2 Non-Uniform Altruism

We analyze the case where the altruism vector consists of numbers in $\{0, 1\}$, i.e., each player is either completely altruistic or completely selfish.[2] The next theorem shows that in this case too, the pure price of anarchy *improves* with the overall altruism level.

Theorem 99. *Let* $\Gamma = (n, m, \Sigma, d)$ *be a symmetric singleton linear congestion game, let* $\beta \in [0, 1]$ *and let* $\alpha \in \{0, 1\}^n$ *such that* $\sum_{i \in [n]} \alpha_i / n = \beta$*, and let* C *be the sum-of-costs social cost function (1.3.1.8). Then, the pure price of anarchy of the* α*-altruistic extension* Γ^α *of* Γ *with respect to* C *is at most* $(4 - 2\beta)/(3 - \beta)$*.*

Proof. Let $s \in \Sigma$ be a pure equilibrium of Γ^α and let $s^* \in \Sigma$ be a strategy profile that minimizes C. Based on the strategy profile s, we partition $[m]$ into sets S_0 and S_1. The set S_1 is defined as in Lemma 97, i.e.,

$$S_1 = \{j \in [m] \mid \exists i \in [n] : \alpha_i = 1, s_i = \{j\}\},$$

is the set of facilities that have at least one player choosing it under s. We let $S_0 = [m] \setminus S_1$ be the set of facilities that are used exclusively by selfish players or not used at all. Let P_1 refer to the set of players $i \in [n]$ such that $s_i \cap S_1 \neq \varnothing$, and let P_0 refer to $[n] \backslash P_1$, i.e., the set of players that choose a facility in P_0, under s. The set P_1 may contain both altruistic and selfish players, while P_0 consists of selfish players only.

Let $\gamma \in [0, 1]$ be the number that satisfies that $\gamma C(s) = \sum_{j \in S_0} |P_j(s)| d_j(|P_j(s)|)$ and $(1 - \gamma) C(s) = \sum_{j \in S_1} |P_j(s)| d_j(|P_j(s)|)$. The high-level approach of this proof is as follows: We bound $\gamma C(s)$ and $(1 - \gamma) C(s)$ separately, to show that

$$\frac{3}{4} \gamma C(s) + (1 - \gamma) C(s) \leq C(s^*). \tag{4.8}$$

The pure price of anarchy is therefore at most $((3/4)\gamma + (1 - \gamma))^{-1} = 4(4 - \gamma)$. The bound then follows by deriving an upper bound on γ, which is done in Lemma 101.

By Lemma 97, we may assume without loss of generality that s^* satisfies the properties expressed in Lemma 97.

Lemma 100. *Define* $y^* = (y_1, \ldots, y_m)$ *as* $y_j^* = |P_j(s^*)| - |P_j(s)| \geq 0$ *for* $j \in S_1$*, and* $y_j^* = |P_j(s^*)|$ *for* $j \in S_0$*. Then,*

$$\sum_{j \in S_0} |P_j(s)| d_j(|P_j(s)|) \leq \frac{4}{3} \sum_{j \in [m]} y_j^* d_j(|P_j(s^*)|).$$

[2] This model relates naturally to *Stackelberg scheduling games* (see, e.g., Chen and Kempe [2008]).

Proof. Consider the game $\Gamma' = (n, m, \Sigma', d')$ obtained from Γ^α by removing for each player $i \in P_1$ from Σ_i all strategies except s_i. I.e., in Γ' players in P_1 are "fixed" on the facilities in S_1 according to s. It is clear that s is a pure equilibrium of the game Γ'. Let $\bar{s}^* \in \Sigma'$ be a strategy profile that minimizes C on the domain Σ'. Because Γ' is in essence a symmetric singleton linear congestion game with only selfish players (i.e., P_0), we can apply Theorem 94 to conclude that

$$C(s) \leq \frac{4}{3} C(\bar{s}^*).$$

Let $s' \in \Sigma'$ be any strategy profile such that $|\{i \in P_0 \mid s_i = \{j\}\}| = y_j^*$. Note that s' exists, because $\sum_{j \in [m]} y_j^* = |P_0|$. Because $\bar{s}^*$ minimizes C, it follows from the above inequality that

$$C(s) \leq \frac{4}{3} C(s').$$

We subtract $\sum_{j \in S_1} |P_j(s)| d_j(|P_j(s)|)$ from both sides, to obtain

$$
\begin{aligned}
\sum_{j \in S_0} |P_j(s)| d_j(|P_j(s)|) &\leq \frac{4}{3} \left(\sum_{j \in [m]} |P_j(s')| d_j(|P_j(s')|) - \sum_{j \in S_1} |P_j(s)| d_j(|P_j(s)|) \right) \\
&= \frac{4}{3} \left(\sum_{j \in S_0} |P_j(s^*)| d_j(|P_j(s^*)|) \right. \\
&\qquad \left. + \sum_{j \in S_1} ((y_j^* + P_j(s)) d_j(y_j^* + |P_j(s)|) - |P_j(s)| d_j(|P_j(s)|)) \right) \\
&= \frac{4}{3} \left(\sum_{j \in S_0} y_j^* d_j(|P_j(s^*)|) \right. \\
&\qquad \left. + \sum_{j \in S_1} ((y_j^* + P_j(s)) d_j(|P_j(s^*)|) - |P_j(s)| d_j(|P_j(s)|)) \right) \\
&\geq \frac{4}{3} \left(\sum_{j \in S_0} y_j^* d_j(|P_j(s^*)|) \right. \\
&\qquad \left. + \sum_{j \in S_1} (y_j^* d_j(|P_j(s^*)|) + (|P_j(s)| - |P_j(s)|) d_j(|P_j(s)|)) \right)
\end{aligned}
$$

$$= \frac{4}{3} \sum_{j \in [m]} y_j^* d_j(|P_j(s^*)|).$$

□

Lemma 101. *It holds that $\gamma \leq 2(1-\beta)/(2-\beta)$.*

Proof. The claim follows directly from Theorem 94 if $P_1 = \varnothing$. Assume therefore that $P_1 \neq \varnothing$, and let $i \in P_1$ with $s_i = \{\bar{j}\}$. Let $\bar{C}(s) = \sum_{i' \in P_0} c_{i'}(s)/|P_0|$ be the average cost experienced by players in $|P_0|$. We first show $c_i(s) \geq \bar{C}(s)/2$. If $|P_0| = \varnothing$, then $c_i(s) \geq \bar{C}(s)/2$ trivially holds. Suppose therefore that $P_0 \neq \varnothing$, and let $i' \in P_0$ with $s_i' = \{j\}$. Recall that i is selfish. Because s is a pure equilibrium, we have

$$c_{i'}(s) = a_j|P_j(s)| + b_j \leq a_{\bar{j}}(|P_{\bar{j}}(s) + 1) + b_{\bar{j}} \leq 2(a_{\bar{j}}|P_{\bar{j}}(s)| + b_{\bar{j}}) = 2c_i(s).$$

By summing over all selfish players in P_0, we obtain $c_i(s) \geq \bar{C}(s)/2$ and thus $\sum_{i \in P_1} c_i(s) \geq |P_1|\bar{C}(S)/2$. We have

$$\begin{aligned} \gamma &= \frac{\sum_{i \in P_0} c_i(s)}{\sum_{i \in P_0} c_i(s) + \sum_{i \in P_1} c_j(s)} \\ &\leq \frac{|P_0|\bar{C}(s)}{|P_0|\bar{C}(s) + \frac{1}{2}|P_1|\bar{C}(s)} \\ &= \frac{2|P_0|}{n + |P_0|} \\ &\leq \frac{2(1-\beta)}{2-\beta}, \end{aligned}$$

where the last inequality follows because $|P_0| \leq (1-\beta)n$. □

Using the above lemmas, we can show that the relation in (4.8) holds:

$$\begin{aligned} \frac{3}{4}\gamma C(s) + (1-\gamma)C(s) &= \frac{3}{4}\sum_{j \in S_0} |P_j(s)|d_j(|P_j(s)|) + \sum_{j \in S_1} |P_j(s)|d_j(|P_j(s)|) \\ &\leq \sum_{j \in [m]} y_j^* d_j(|P_j(s^*)|) + \sum_{j \in S_1} |P_j(s)|d_j(|P_j(s)|) \\ &= \sum_{j \in [m]} |P_j(s^*)|d_j(P_j(s^*)|) \\ &\qquad + \sum_{j \in S_1} (|P_j(s)|d_j(|P_j(s)|) - |P_j(s)|d_j(|P_j(s^*)|)) \end{aligned}$$

$$\leq \sum_{j \in [m]} |P_j(s^*)| d_j(|P_j(s^*)|) = C(s^*),$$

where the first inequality follows from Lemma 100 and the last inequality follows from the property that $|P_j(s^*)| \geq |P_j(s)|$ for all $j \in S_1$. We conclude that the pure price of anarchy is at most

$$\left(\frac{3}{4}\gamma + (1-\gamma)\right)^{-1} = \frac{4}{4-\gamma} \leq \frac{4-2\beta}{3-\beta},$$

Where the last inequality follows from Lemma 101. □

4.8 General Properties of Smoothness

For the game classes that we analyzed (with the exception of symmetric singleton congestion games), we used (λ, μ, α)-smoothness as our main tool to derive bounds on the price of anarchy. In this section, we provide some general results about (λ, μ, α)-smoothness.

Proposition 102. *Let $n \in \mathbb{N}_{\geq 0}$. Suppose that $\mathcal{G}$ is a class of cost minimization games with player set $[n]$, each game equipped with a sum-bounded social cost function. The set $S_{\mathcal{G}} = \{(\lambda, \mu, \alpha) \in \mathbb{R}^2 \times [0,1]^n \mid \forall \Gamma \in \mathcal{G} : \Gamma$ is (λ, μ, α)-smooth$\}$ is convex.*

The same holds when instead, $\mathcal{G}$ is a class of utility maximization games, each equipped with a sum-bounded social welfare function.

Proof. We proof the claim for cost minimization games. For the case of utility maximization games, the proof is analogous. Pick an arbitrary game $\Gamma = (n, \Sigma, d) \in \mathcal{G}$. It suffices to show that $S_\Gamma = \{(\lambda, \mu, \alpha) \in \mathbb{R}^2 \times [0,1]^n \mid G$ is (λ, μ, α)-smooth$\}$ is convex, because the intersection of any collection of convex sets is convex.

Let $(\lambda_1, \mu_1, \alpha^1), (\lambda_2, \mu_2, \alpha^2) \in S_\Gamma$ be two elements in S_Γ, and pick an arbitrary number $\gamma \in [0,1]$. For all pairs $s, s^* \in \Sigma$ of strategy profiles,

$$\begin{aligned}
&\gamma \sum_{i \in [n]} (c_i(s_i^*, s_{-i}) + \alpha_i^1(c_{-i}(s_i^*, s_{-i}) - c_{-i}(s))) \\
&\qquad + (1-\gamma) \sum_{i \in [n]} (c_i(s_i^*, s_{-i}) + \alpha_i^2(c_{-i}(s_i^*, s_{-i}) - c_{-i}(s))) \\
&\leq \gamma(\lambda_1 C(s^*) + \mu_1 C(s)) + (1-\gamma)(\lambda_2 C(s^*) + \mu_2 C(s)).
\end{aligned}$$

By rewriting both sides of the above inequality, we obtain

$$\sum_{i\in[n]}(c_i(s_i^*, s_{-i}) + (\gamma\alpha_i^1 + (1-\gamma)\alpha_i^2)(c_{-i}(s_i^*, s_{-i}) - c_{-i}(s)))$$
$$\leq (\gamma\lambda_1 + (1-\gamma)\lambda_2)C(s^*) + (\gamma\mu_1 + (1-\gamma)\mu_2)C(s).$$

We conclude that Γ is $(\gamma(\lambda_1, \mu_1, \alpha^1) + (1-\gamma)(\lambda_2, \mu_2, \alpha^2))$-smooth. Therefore, S_Γ is convex. ☐

A natural question to ask now is whether the robust price of anarchy is also a convex function of α. This turns out not to be the case. For instance, the robust price of anarchy of the class of uniformly α-altruistic congestion games with respect to the sum-of-costs social cost function is $\frac{5+4\breve{\alpha}}{2+\breve{\alpha}}$ (see Section 4.6), which is a non-convex function. However, we can prove a somewhat weaker statement. For a subset $S \subseteq \mathbb{R}^n$, we call a function $f : S \to \mathbb{R}$ *quasi-convex* iff $f(\gamma x + (1-\gamma)y) \leq \max\{f(x), f(y)\}$ for all $\gamma \in [0,1]$.

Theorem 103. *Let $n \in \mathbb{N}_{\geq 1}$ and let $\mathcal{G}$ be a class of cost minimization games with player set $[n]$, each game equipped with a sum-bounded social cost function. Then RPoA$_\mathcal{G}$ is a quasi-convex function on its domain, $[0,1]^n$.*

The same holds when instead, $\mathcal{G}$ is a class of utility maximization games, each equipped with a sum-bounded social welfare function.

Proof. We proof the claim for cost minimization games. For the case of utility maximization games, the proof is analogous. Let $(\Gamma, C) \in \mathcal{G}$, where Γ is a game and C is a social cost function for Γ. We show that for any $\alpha^1, \alpha^2 \in \mathbb{R}^n$ and $\gamma \in [0,1]$,

$$\text{RPoA}_\Gamma(\gamma\alpha^1 + (1-\gamma)\alpha^2) \leq \max\{\text{RPoA}_\Gamma(\alpha^1), \text{RPoA}_\Gamma(\alpha^2)\}.$$

Let $(\epsilon_1, \epsilon_2, \ldots)$ be a decreasing sequence of positive real numbers that tends to 0. Moreover, let

$$((\lambda_{1,1}, \mu_{1,1}, \alpha^1), (\lambda_{1,2}, \mu_{1,2}, \alpha^1), \ldots)$$

and

$$((\lambda_{2,1}, \mu_{2,1}, \alpha^2), (\lambda_{2,2}, \mu_{2,2}, \alpha^2), \ldots)$$

be sequences of elements in S_Γ (where S_Γ is as defined in the proof of Proposition 102) such that

$$\text{RPoA}_\Gamma(\alpha^1) + \epsilon_j = \frac{\lambda_{1,k}}{1-\mu_{1,j}}$$

and

$$\text{RPoA}_\Gamma(\alpha^2) + \epsilon_j = \tfrac{\lambda_{2,k}}{1-\mu_{2,j}}$$

for all $k \in \mathbb{N}_{\geq 1}$. By Proposition 102, we know that for all $k \in \mathbb{N}_{\geq 1}$,

$$\begin{aligned}\sum_{i\in[n]} &(c_i(s_i^*, s_{-i}) + (\gamma\alpha_i^1 + (1-\gamma)\alpha_i^2)(c_{-i}(s_i^*, s_{-i}) - c_{-i}(s))) \\ &\leq \gamma(\lambda_{1,k}C(s^*) + \mu_{1,k}C(s)) + (1-\gamma)(\lambda_{2,k}C(s^*) + \mu_{2,k}C(s)) \\ &\leq \max\{\lambda_{1,k}C(s^*) + \mu_{1,k}C(s), \lambda_{2,k}C(s^*) + \mu_{2,k}C(s)\}.\end{aligned}$$

Hence,

$$\begin{aligned}\text{RPoA}_\Gamma(\gamma\alpha^1 + (1-\gamma)\alpha^2) &\leq \max\left\{\frac{\lambda_{1,k}}{1-\mu_{1,k}}, \frac{\lambda_{2,k}}{1-\mu_{2,k}}\right\} \\ &\leq \max\{\text{RPoA}(\alpha^1), \text{RPoA}(\alpha^2)\} + \epsilon_k,\end{aligned}$$

for all $k \in \mathbb{N}_{\geq 0}$. By taking the limit as k goes to infinity, we conclude $\text{RPoA}_\Gamma(\gamma\alpha^1 + (1-\gamma)\alpha^2) \leq \max\{\text{RPoA}(\alpha^1), \text{RPoA}(\alpha^2)\}$, which proves the claim. □

The quasi-convexity of $\text{RPoA}_\mathcal{G}$ implies:

Corollary 104. *Let $n \in \mathbb{N}_{\geq 1}$ and let $\mathcal{G}$ be a class of cost minimization games with player set $[n]$, each game equipped with a sum-bounded social cost function. The points $\alpha \in [0,1]^n$ that minimize $\text{RPoA}_\mathcal{G}$ on the domain $[0,1]^n$ form a convex set. The set of points $\alpha \in [0,1]^n$ that maximize $\text{RPoA}_\mathcal{G}$ on the domain $[0,1]^n$ includes at least one point in $\{0,1\}^n$.*

The same holds when instead, $\mathcal{G}$ is a class of utility maximization games, each equipped with a sum-bounded social welfare function.

4.9 Conclusions

One might not expect that there are games in which the price of anarchy is greater than 1 when $\alpha = \mathbf{1}$. This phenomenon is a lot less surprising when approached from a local search point-of-view, as this is only equivalent to saying that there exist local optima in the objective function C with respect to the neighborhood set obtained by taking all strategies obtained by single-player deviations from a given strategy profile s. Nevertheless, it still seems to us rather surprising that the price of anarchy can get worse when the altruism levels α get closer to $\mathbf{1}$. This phenomenon has been observed before, by Caragiannis et al. [2010]. The fact that the price of anarchy does not *necessarily* get worse in all cases is exemplified by our analysis of the pure price of anarchy in symmetric singleton congestion games.

The most immediate research directions include analyzing singleton congestion games with more general delay functions than linear ones. While the price of anarchy of such functions increases (e.g., the price of anarchy for polynomials increases exponentially in the degree [Awerbuch et al., 2005, Christodoulou and Koutsoupias, 2005a]), this also creates room for potentially larger reductions in the price of anarchy as a consequence of altruism. Similarly, the characterization of the robust price of anarchy of altruistic congestion games with more general delay functions (e.g., polynomials) is left for future work.

For games where the smoothness argument cannot give tight bounds, would a refined smoothness argument like local smoothness [Roughgarden and Schoppmann, 2011] work? For symmetric singleton congestion games, this seems unlikely, as the price of anarchy bounds are already different between pure and mixed Nash equilibria. It is also worth trying to apply the smoothness argument or its refinements to analyze the price of anarchy for other dynamics in other classes of altruistic games, for example, (altruistic) network vaccination games [Chen et al., 2010], which are known to not always possess pure Nash equilibria, or to find examples to see why smoothness-based arguments do not work.

We have seen that the impact of altruism depends on the underlying game. It would be nice to identify general properties that enable to predict whether a given game suffers from altruism or not. What is it that makes valid utility game invariant to altruism? Furthermore, what kind of "transformations" (not just altruistic extensions) might be applied to a strategic game such that the smoothness approach can still be adapted to give (tight) bounds? More generally, while the existence of pure equilibria has been shown for singleton and matroid congestion games with player-specific latency functions [Ackermann et al., 2006, Milchtaich, 1996], the price of anarchy (for pure Nash equilibria or more general equilibrium concepts) has not yet been addressed. Studying the price of anarchy in such a general setting (in which our setting with altruism can be embedded) by either smoothness-based techniques or other methods, is undoubtedly intriguing.

Chapter 5

Inefficiency of Games with Social Context*

We study in this chapter a variant of the altruism framework introduced in the previous chapter, where we can conveniently model arbitrary altruistic *social structures*, causing players to direct their altruistic behavior in a refined player-specific way (depending, for example, on friendships that exist among the players). Instead of a single altruism coefficient α_i for each player i of a game, we now associate an altruism coefficient $\alpha_{i,i'}$ to each ordered pair of players (i, i') in the game, reflecting to what extent player i cares about the direct cost or direct utility of player j. For this altruism model, we will focus on establishing upper bounds on the coarse price of anarchy of altruistic extensions of four classes of full information games:

- linear congestion games,
- singleton linear congestion games in which all facilities have identical delay functions,
- minsum machine scheduling games,
- generalized second price auctions.

For a definition of the first two classes of games, see Section 1.3.1.10. Minsum machine scheduling games and generalized second price auctions are defined in Section 5.3.1

*A part of the contents of this chapter has been published as Anagnostopoulos et al. [2013].

Like in the previous chapter, an extension of the smoothness notion of Roughgarden [2009] will again be the main tool for obtaining most of our bounds.

The material from Chapter 1 that is necessary for understanding this chapter consists of Section 1.3 up to Section 1.3.1.5, and Section 1.3.1.7 up to Section 1.3.1.10, with emphasis on Sections 1.3.1.9 and 1.3.1.10.

We first define formally the altruism model we will work with.

Definition 105 (Altruistic social context)**.** Given a game Γ with player set $[n]$, an *altruistic social context* for Γ is an $n \times n$ matrix $\alpha \in (\mathbb{R}^n_{\geq 0})^n$.

Definition 106 (Altruistic extension of a game for altruistic social context)**.** Let $\Gamma = (n, \Sigma, c)$ be a game and let α be an altruistic social context for Γ. The α*-altruistic extension* of Γ is defined as the strategic game $\Gamma^\alpha = (n, \Sigma, c^\alpha)$, where for all $i \in [n]$ and $s \in \Sigma$,

$$c_i^\alpha(s) = \sum_{i' \in [n]} \alpha_{i,i'} c_{i'}(s). \tag{5.1}$$

In case Γ is a cost minimization game, c_i is called the *selfish cost function*, and c_i^α is called the *perceived cost function*. In case Γ is a utility maximization game (n, Σ, u), u_i is called the *selfish utility function*, and c_i^α is called the *perceived utility function*. Thus, the perceived cost or perceived utility that player i experiences is a non-negative linear combination of its selfish cost or selfish utility, and the selfish costs or selfish utilities of the other players.

We note some essential differences between the present altruism model and the altruism model of the previous chapter (Definitions 74 and 4): In the previous chapter, convex combinations are considered, while here we take linear combinations. Moreover, we do not need to define an altruistic extension with respect to a given social cost or welfare function, as we only consider linear combinations of the selfish costs of the players.

It is possible to extend the above definition into one where negative numbers are allowed in α. One could regard such a negative entry $\alpha_{i,i'}, i, i' \in [n]$ as player i having a *spiteful* attitude towards player i'. Our main focus however is on altruistic behavior, and we leave spiteful behavior unexplored in this chapter.[1]

This altruism model has a natural interpretation in terms of *social networks*: One could suppose that the players in $[n]$ are identified with the nodes of a complete directed graph $G = (N, A)$. The weight of an edge $(i, i') \in A$ is equal to $\alpha_{i,i'}$, specifying the extent to which player i cares about the cost or utility of player i'.

[1] However, we will investigate spiteful player behavior in Chapter 6, for the class of *procurement auctions*.

Just as in the previous chapter we will study the price of anarchy of some classes of altruistic extensions of games. This will in all cases be done with respect to the sum-of-costs social cost function and sum-of-utilities social welfare function (i.e., (1.4) and (1.3.1.8) in Chapter 1).

Remark 107. In this chapter, the social cost function and social welfare function for the altruistic extensions of the games Γ we consider will always be the sum-of-costs social cost function and sum-of-utilities social welfare function of Γ. Therefore in this chapter, contrary to the previous chapter, we sometimes omit mentioning explicitly this social cost function and social welfare function in the discussions and the statements of our results.

Note nonetheless that we only make this assumption for the particular classes of games we consider. The more general results that we present in Section 5.3.2 hold indeed for arbitrary social cost functions and social welfare functions.

We distinguish between *unrestricted* and *restricted* altruistic social contexts.

Definition 108 (Restricted altruistic social context)**.** A *restricted* altruistic social context α for a game $\Gamma = (n, \Sigma, c)$ is an altruistic social context for Γ where $\alpha_{i,i} > 0$ and $\alpha_{i,i'} \leq \alpha_{i,i}$ for all $i, i' \in [n]$. An *unrestricted* altruistic social context α for a game Γ is simply an altruistic social context for Γ.

For the class of unrestricted social contexts α, for most interesting classes of games one can prove trivially unboundedness of the price of anarchy of the α-altruistic extension of any game, just by setting $\alpha_{i,i'} = 0$ for all players i, i', because this causes every strategy profile to be an equilibrium. Also, with the additional property that $\alpha_{i,i'} > 0$, for all classes of games we consider there exist easy examples that show unboundedness of the price of anarchy.

For this reason we often make the assumption that the altruistic social context is a restricted altruistic social context. In a restricted altruistic social context, every player cares for itself at least as much as for any other player, which still captures a broad set of scenarios that can be considered realistic. In the case of a restricted altruistic social context α, note that for the sake of our analysis of the price of anarchy, we can normalize α without loss of generality such that $\alpha_{i,i} = 1$ for every player i.[2]

Our findings show that the increase in the price of anarchy is modest for congestion games and *minsum scheduling games*, whereas it is drastic for generalized second price auctions.

[2]To see this, note that, by dividing all $\alpha_{i,i'}$ by $\alpha_{i,i} > 0$, the set of equilibria and the social cost and social welfare of any strategy profile remain the same.

5.1 Background

In (algorithmic) game theory literature, considerable attention has been given to more general settings in which the players do not necessarily behave entirely selfishly, but may alternatively exhibit *spiteful* or *altruistic* behavior; see, for instance, Brandt et al. [2007], Buehler et al. [2011], Caragiannis et al. [2010], Chen and Kempe [2008], Chen et al. [2011a,b], Elias et al. [2010], Hoefer and Skopalik [2009b,a, 2012], Bilò et al. [2011a]. Studying such alternative behavior in games is motivated by the observation that altruism and spite are phenomena that frequently occur in real life (see, for example, Fehr and Schmidt [2006]). Consequently, it is desirable to incorporate such alternative behavior in game-theoretical analyses.

Previous work on the price of anarchy for spiteful and altruistic games has focused on simple models of spite and altruism, where a spite or altruism level α_i is associated to each player i denoting the extent to which its perceived cost is influenced by a nonspecific other player. In a series of papers [Buehler et al., 2011, Caragiannis et al., 2010, Chen et al., 2011a,b] as well as the previous chapter, it has been observed that altruistic behavior can actually be harmful in the sense that the price of anarchy may *increase* as players become more altruistic. This observation serves as a starting point for the investigations conducted in this chapter. The main question that we address here is: How severe can this effect be if one considers more refined models of altruism that capture complex social relationships between the players?

There are several other works that propose models of altruism and spite [Brandt et al., 2007, Buehler et al., 2011, Caragiannis et al., 2010, Chen and Kempe, 2008, Elias et al., 2010, Hoefer and Skopalik, 2009b,a, 2012]. All these models are special cases of the one studied here. The model studied in the previous chapter is a special case of the present model as well, in case the choice of social welfare function and social cost function is the sum-of-utilities social welfare function and the sum-of-costs social cost function. The inefficiency of equilibria in the presence of altruistic or spiteful behavior was studied for various games in Buehler et al. [2011], Caragiannis et al. [2010], Chen and Kempe [2008], Elias et al. [2010].

Related to the work discussed here is the work by Bilò et al. [2011a], where altruism is studied for the class of congestion games, for the case where the perceived cost of a player is the minimum, maximum, or sum of the immediate cost of its neighbors. Bilò et al. [2011a] establishes, among other results, an upper bound of $17/3$ on the pure price of anarchy of linear congestion games for a special case of the setting that we study here.

Related but different from our setting, is the concept of *graphical congestion games* [Bilò et al., 2011b, Fotakis et al., 2009]. In such games, the cost and the strategy set of a player depends only on a subset of the players.

5.2 Contributions and Outline

Using the smoothness extension that we propose in Section 5.3, we prove upper bounds on the price of anarchy for altruistic extensions of the classes of games mentioned above.

The perspective adopted in this chapter is slighly different from that of the previous chapter: We adopt here a worst-case point of view by trying to answer the question of how bad the price of anarchy can be over a large set of altruistic extensions, while in the previous chapter we were concerned with giving bounds on the price of anarchy as a function of the altruism levels.

We show in all cases that for unrestricted altruistic social contexts α, the pure price of anarchy of the α-altruistic extension of the game is unbounded, even if $\alpha > \mathbf{0}$. We therefore derive our upper bounds under the assumption that the altruistic social context is restricted, conform Definition 108. Under this assumption, we derive the following upper bounds on the coarse price of anarchy:

- A bound of 7 for altruistic linear congestion games (Section 5.4).
- A bound of $\varphi^3 \approx 4.2361$ for the special case of singleton linear congestion games with identical delay functions for all facilities, where $\varphi = (1 + \sqrt{5})/2$ denotes the golden ratio (Section 5.5). We use a novel proof approach to establish this bound, because following the proof template that we use for the other upper bounds turns out to be too weak. Instead, we use a more refined amortized argument, based on distributing "budget" in a sophisticated way among the facilities.
- A bound of $4 + 2\sqrt{3} \approx 7.4641$ and $12 + 8\sqrt{2} \approx 23.3137$ for altruistic min-sum machine scheduling games with related and unrelated machines, respectively (Section 5.6).
- A bound of $2(n + 1)$ for altruistic generalized second price auctions, where n is the number of players (Section 5.7). We additionally provide a family of examples of GSP auctions for which a smoothness proof cannot give an upper bound on the price of anarchy that is better than $n/2$. This stands in contrast with the small constant price of anarchy that is known for the purely selfish setting, which is achieved through smoothness ([Caragiannis et al., 2011]).

5.3 Preliminaries

We define in this section the classes of games studied in this chapter that are not already defined in Chapter 1, and we give our smoothness extension for games with altruistic social contexts.

5.3.1 Minsum Scheduling Games and Generalized Second Price Auctions

In a minsum scheduling game there is a set of jobs to be completed on a set of machines, and each job has a certain given processing time on each machine. Each job takes the role of a player that has to pick a machine to run itself on, and it wants to minimize its completion time. A machine executes its set of jobs in order of increasing processing time, breaking ties according to a given deterministic tie-breaking rule.

Definition 109 (Minsum machine scheduling game)**.** A *minsum machine scheduling game* is a cost minimization game $\Gamma = (n, \Sigma, c)$ for which there exists a number $m \in \mathbb{N}_{>0}$ of *machines*, *processing times* $p_{i,j} \in \mathbb{R}_{\geq 0}$ for all $i \in [n], j \in [m]$, and a strict total order $\preceq_j$ on $[n]$ for all $j \in [m]$ with the properties that:

- $i \prec_j i'$ or $i' \prec_j i$ for all $i, i' \in [n], j \in [m]$.
- If $i \prec_j i'$, then $p_{i,j} \leq p_{i',j}$, for all $i, i' \in [n], j \in [m]$.

It holds that $\Sigma_i = [m]$ for all $i \in [n]$, and $c_i(s) = \sum_{i' \in [n]: s_{i'} = j, i' \preceq_j i} p_{i',j}$. When, in a strategy profile $s \in \Sigma$, it holds for two players $i, i' \in [n]$ that $s_i = s_{i'}$ and $i \prec_{s_i} i'$, we say that i *is scheduled before* i' *under* s and i' *is scheduled after* i *under* s. A minsum scheduling game can therefore be represented as a quadruple $(n, m, p, \preceq)$ where $p \in (\mathbb{R}^n_{\geq 0})^n$ and $\preceq = (\preceq_1, \ldots, \preceq_n)$. In the context of Γ, the set of players $[n]$ is alternatively referred to as the set of *jobs*, and $\preceq_j$ is called the *tie-breaking* rule of machine j.

A minsum scheduling game with *related machines* is a minsum scheduling game $(n, m, p, \preceq)$ such that there is a given speed $a_j \in \mathbb{R}_{>0}$ for each machine $j \in [m]$, and a given length $w_i \in \mathbb{R}_{\geq 0}$ for each job $i \in [n]$, such that $p_{i,j} = w_i / a_j$.

A minsum scheduling game with *unrelated machines* is simply defined as a minsum scheduling game.

Generalized second price auctions are games motivated by online advertising. The players in this game are advertisers that want to place their advertisement on a web page. However, there are only a limited number of slots available on the web page that the advertisers can place their advertisements in, and some slots are more visible than other slots, and are thus more desirable to the advertisers. A prominent example of a generalized second price auction is the auction mechanism used in Google AdWords.

Definition 110 (Generalized second price auction)**.** A *generalized second price auction* is a utility maximization game with infinite strategy sets $\Gamma = (n, \Sigma, u)$, where for each player there is a *valuation* $v_i \in \mathbb{R}_{\geq 0}$ such that $\Sigma_i = [0, v_i]$. The utility functions u are defined as follows.

There are $m \in \mathbb{N}_{\geq 1}$ *slots*, and $[m]$ is called the set of slots. For each slot $j \in [m]$, there is an associated *click-through rate* $\beta_j \in [0, 1]$, where $\beta_j \geq \beta_{j'}$ for $j' \in [m], j' > j$. Moreover, there is a total order $\preceq$ on $[n]$ called the *tie-breaking rule*. For a strategy profile s we denote by $r(s)$ the *ranking* of the players under s: $r(s)$ is the m-dimensional vector such that for all $i \in [n], j \in [m]$ it holds that $r_j(s) = i$ iff $|\{i' \in [n] \setminus \{i\} \mid s_{i'} > s_i \text{ or } s_{i'} = s_i, i' \prec i\}| = j - 1$. A player $i \in [n]$ is said to *win slot* $j \in [m]$ *under* $s \in \Sigma$ iff $r_j(s) = i$. The *valuation of player* $i \in [n]$ *for a slot* $j \in [m]$ is given by $\beta_j v_i$. The *price paid by player* i who wins slot j under strategy profile $s \in \Sigma$ is given by $\beta_j s_{r_{j+1}(s)}$ (where $\beta_j = 0$ if $j > m$). For player $i \in [n]$ and strategy profile $s \in \Sigma$, the utility $u_i(s)$ is defined as the valuation that i has for the slot that it has won under s, minus the price i has to pay under s.

A generalized second price auction can thus be represented as a tuple $(n, m, v, \beta, \preceq)$, where $v = (v_1, \ldots, v_n)$ and $\beta = (\beta_1, \ldots, \beta_m)$.

5.3.2 Smoothness and a Proof Template

We define our smoothness notion for games with altruistic social context as follows.

Definition 111 ((λ, μ, α)-smoothness for games with altruistic social context)**.** Let $\Gamma = (n, \Sigma, c)$ be a cost minimization game, let C be a social cost function for Γ, and let $\alpha \in (\mathbb{R}^n)^n$ be an altruistic social context. Further, let $s^* \in \Sigma$ be a strategy profile that minimizes C. Γ is (λ, μ, α)-*smooth* iff there exists a strategy profile $\bar{s} \in \Sigma$ such that for every strategy profile $s \in \Sigma$ it holds that

$$\sum_{i \in [n]} \sum_{i' \in [n]} \alpha_{i,i'}(c_{i'}(\bar{s}_i, s_{-i}) - c_{i'}(s)) \leq \lambda C(s^*) + (\mu - 1)C(s). \tag{5.2}$$

When instead, Γ is a utility maximization game (n, Σ, u), U is a social welfare function for Γ, and $s^* \in \Sigma$ is a strategy profile maximizing U, then Γ is (λ, μ, α)-*smooth* iff there exists a strategy profile $\bar{s} \in \Sigma$ such that for every strategy profile $s \in \Sigma$ it holds that

$$\sum_{i \in [n]} \sum_{i' \in [n]} \alpha_{i,i'}(u_{i'}(\bar{s}_i, s_{-i}) - u_{i'}(s)) \geq \lambda U(s^*) - (\mu + 1)U(s). \tag{5.3}$$

Just like the smoothness extension presented in Chapter 4, we prove some analogues of results from Roughgarden [2009]. The following theorem shows that (λ, μ, α)-smoothness implies a bound on the coarse price of anarchy of α-altruistic games.

Theorem 112. *Let Γ be a cost minimization game, let C be a social cost function for Γ, and let $\alpha \in (\mathbb{R}^n_{\geq 0})^n$ be an altruistic social context for Γ. If Γ is (λ, μ, α)-smooth with respect to C for some values $\lambda, \mu \in \mathbb{R}$ with $\mu < 1$, then the coarse price of anarchy of Γ^α with respect to C is at most $\lambda/(1-\mu)$.*

If instead, Γ is a utility maximization game, U is a social welfare function for Γ, and Γ is (λ, μ, α)-smooth with respect to U for some values $\lambda, \mu \in \mathbb{R}$ with $\mu > -1$, then the coarse price of anarchy of Γ^α with respect to C is at most $(1+\mu)/\lambda$.

Proof. We prove this for the case of cost minimization games. For utility maximization games, the proof is analogous.

Let σ be a coarse equilibrium of Γ^α, let $\bar{s}$ be a strategy profile for which (5.2) holds, and let $s^* \in \Sigma$ be an strategy profile. The coarse equilibrium condition implies that for every player $i \in [n]$:

$$\sum_{i' \in [n]} \alpha_{i,i'} \mathbf{E}_{s\sim\sigma}[c_{i'}(\bar{s}_i, s_{-i})] - \sum_{i' \in [n]} \alpha_{i,i'} \mathbf{E}_{s\sim\sigma}[c_{i'}(s)] \geq 0.$$

By summing over all players and using linearity of expectation, we obtain

$$\mathbf{E}_{s\sim\sigma}[C(s)] \leq \mathbf{E}_{s\sim\sigma}[C(s)] + \mathbf{E}_{s\sim\sigma}\left[\sum_{i\in[n]}\sum_{i'\in[n]} \alpha_{i,i'}(c_{i'}(\bar{s}_i, s_{-i}) - c_{i'}(s))\right].$$

Now we use the smoothness property (5.2) and obtain

$$\mathbf{E}_{s\sim\sigma}[C(s)] \leq \mathbf{E}_{s\sim\sigma}[C(s)] + \mathbf{E}_{s\sim\sigma}[\lambda C(s^*) + (\mu-1)C(s)] = \lambda C(s^*) + \mu \mathbf{E}_{s\sim\sigma}[C(s)].$$

Since $\mu < 1$, this implies that the coarse price of anarchy is at most $\lambda/(1-\mu)$. □

Observe that in the purely selfish setting (i.e., when $\alpha_{i,i} = 1$ and $\alpha_{i,i'} = 0$ for every $i, i' \in [n]$, $i \neq j$), our smoothness notion is more permissive than the one in Roughgarden [2009] where (5.2) is required to hold for any arbitrary strategy profile s^* and with $\bar{s} = s^*$. Also in Roughgarden [2009], the analogue of Theorem 112 is shown under the additional assumption that the social welfare function or social cost function is *sum-bounded*. Here, we managed to get rid of this assumption.

The smoothness condition also proves useful in the context of no-regret sequences and the *price of total anarchy*, introduced by Blum et al. [2008] (see Section 1.3.1.4). The following proposition is the analogue of Proposition 80 for games with altruistic social contexts.

Proposition 113. *Let $\Gamma = (n, \Sigma, c)$ be a cost minimization game, let $\alpha \in (\mathbb{R}^n)^n$ be an altruistic social context for Γ, and suppose Γ is (λ, μ, α)-smooth, for some values $\lambda, \mu \in \mathbb{R}$. Let $s^* \in \Sigma$ be a strategy profile minimizing a social cost function C for Γ, and let $(s^1, s^2, \dots)$ be a vanishing regret sequence of strategy profiles in Σ, with respect to the α-altruistic extension Γ^α, with respect to C. This sequence then satisfies that factor by which the average social cost deviates from the optimum social cost, converges to at most $\lambda/(1-\mu)$. I.e.,*

$$\lim_{T\to\infty} \frac{1}{T} \sum_{t\in[T]} C(s^t) \le \frac{\lambda}{1-\mu} C(s^*).$$

If Γ is instead a profit maximization game (n, Σ, u), and U is a social welfare function for Γ, and $(s^1, s^2, \dots)$ is a vanishing regret sequence of strategy profiles in Σ with respect to the α-altruistic extension Γ^α, then

$$\lim_{T\to\infty} \frac{1}{T} \sum_{t\in[T]} U(s^t) \ge \frac{1+\mu}{\lambda} U(s^*).$$

Proof. We prove the claim for cost minimization games. For profit maximization games, the proof is analogous.

The (λ, μ, α)-smoothness condition (5.2) is equivalent to

$$C(s^t) \le \frac{\lambda}{1-\mu} C(s^*) + \frac{1}{1-\mu} \sum_{i\in[n]} (c_i^\alpha(s^t) - c_i^\alpha(s_i^*, s_{-i}^t)), \tag{5.4}$$

by the definitions of $c_i^\alpha, i \in [n]$. The vanishing regret property of $(s^1, s^2, \ldots)$ implies that for all $i \in [n]$,

$$\begin{aligned}
&\lim_{T\to\infty} \frac{1}{T} \sum_{t\in[T]} (c_i^\alpha(s^t) - c_i^\alpha(s_i^*, s_{-i}^t)) \\
&\quad \le \lim_{T\to\infty} \frac{1}{T} \max\left\{ 0, \sum_{t\in[T]} (c_i^\alpha(s^t) - c_i^\alpha(s_i', s_{-i}^t)) \;\middle|\; s_i' \in \Sigma_i \right\} \\
&\quad = 0.
\end{aligned}$$

Combining this with (5.4) proves the claim:

$$\lim_{T\to\infty} \frac{1}{T} \sum_{t\in[T]} C(s^t)$$

$$
\begin{aligned}
&\leq \lim_{T\to\infty} \frac{1}{T} \sum_{t\in[T]} \left(\frac{\lambda}{1-\mu} C(s^*) + \frac{1}{1-\mu} \sum_{i\in[n]} (c_i^\alpha(s^t) - c_i^\alpha(s_i^*, s_{-i}^t)) \right) \\
&\leq \left(\frac{\lambda}{1-\mu} C(s^*) + \lim_{T\to\infty} \frac{1}{1-\mu} \sum_{i\in[n]} \sum_{t\in[T]} (c_i^\alpha(s^t) - c_i^\alpha(s_i^*, s_{-i}^t)) \right) \\
&= \frac{\lambda}{1-\mu} C(s^*).
\end{aligned}
$$

□

A proof template for upper bounding the coarse price of anarchy through smoothness. A template we use for two of our proofs below, is the following: Suppose that a cost minimization game $\Gamma = (n, \Sigma, c)$ is known to be (λ_1, μ_1)-smooth with respect to a sum-bounded social cost function C for Γ, in the classical sense, for some values $\lambda_1, \mu_1 \in \mathbb{R}$ (i.e., it is (λ, μ, α)-smooth with respect to C for $\alpha_{i,i} = 1$ and $\alpha_{i,i'} = 0$ for all $i, i' \in [n], i \neq i'$). This means that for all strategy profiles $s \in \Sigma$, there is a strategy profile $\bar{s} \in \Sigma$ such that

$$\sum_{i\in[n]} c_i(\bar{s}_i, s_{-i}) \leq \lambda_1 C(s^*) + \mu_1 C(s), \tag{5.5}$$

where $s^* \in \Sigma$ minimizes C. In order to establish $(\lambda_1 + \lambda_2, \mu_1 + \mu_2, \alpha)$-smoothness of Γ with respect to C, for some values $\lambda_2, \mu_2 \in \mathbb{R}$ and some altruistic social context $\alpha \in (\mathbb{R}^n_{\geq 0})^n$, by our definition of smoothness it suffices to prove that for $\bar{s}$ it holds that

$$\sum_{i\in[n]} \sum_{i'\in[n]\setminus\{i\}} \alpha_{i,i'}(c_{i'}(\bar{s}_i, s_{-i}) - c_{i'}(s)) \leq \lambda_2 C(s^*) + \mu_2 C(s). \tag{5.6}$$

Suppose instead that Γ is a utility maximization game (n, Σ, u) and U is a sum-bounded social welfare function for Γ. Now, if Γ is (λ_1, μ_1)-smooth with respect to U in the classical sense, for some values $\lambda_1, \mu_1 \in \mathbb{R}$, then this means that for all strategy profiles $s \in \Sigma$, there is a strategy profile $\bar{s} \in \Sigma$ such that

$$\sum_{i\in[n]} u_i(\bar{s}_i, s_{-i}) \geq \lambda_1 U(s^*) - \mu_1 U(s), \tag{5.7}$$

where $s^* \in \Sigma$ maximizes U. In order to establish $(\lambda_1 + \lambda_2, \mu_1 + \mu_2, \alpha)$-smoothness of Γ with respect to U, for some values $\lambda_2, \mu_2 \in \mathbb{R}$ and some altruistic social context

$\alpha \in (\mathbb{R}^n_{\geq 0})^n$, by our definition of smoothness it suffices to prove that for $\bar{s}$ it holds that

$$\sum_{i \in [n]} \sum_{i' \in [n] \setminus \{i\}} \alpha_{i,i'}(u_{i'}(\bar{s}_i, s_{-i}) - u_{i'}(s)) \geq \lambda_2 U(s^*) - \mu_2 U(s). \tag{5.8}$$

5.4 Linear Congestion Games

The following example shows that for unrestricted altruistic social contexts α, the price of anarchy of the class of α-altruistic extensions of linear congestion games is unbounded, even when $\alpha \neq \mathbf{0}$.

Example 114. Consider a linear congestion game (n, m, Σ, d) with $n = 2$, and $m = 4$. Let the delay functions be defined by $d_1(k) = d_3(k) = k$ and $d_2(k) = d_4(k) = Mk$ for all $k \in [n]$, with $M \in \mathbb{R}_{\geq 0}$. The strategy sets are $\Sigma_1 = \{\{1,2\},\{3\}\}$, $\Sigma_2 = \{\{3,4\},\{1\}\}$. Suppose furthermore that α is given by $\alpha_{1,1} = \alpha_{2,2} = 0$ and $\alpha_{1,2} = \alpha_{2,1} = 1$.

Observe that in the α-altruistic extension of this game, the strategy profile $(\{1,2\}, \{3,4\})$ is a pure equilibrium with social cost $2K + 2$. The optimal social cost is 2, and is attained by the strategy profile $(\{3\},\{1\})$. The price of anarchy in this game is therefore $K + 1$, and K can be taken arbitrarily large, so the price of anarchy of the class of α-altruistic extensions of linear congestion games is unbounded, where α is an unrestricted altruistic social context.

We therefore prove in this section an upper bound on the coarse price of anarchy in case the altruistic social context is restricted.

Theorem 115. *Let* $\Gamma = (n, \Sigma, c)$ *be linear congestion game, and let* $\alpha \in (\mathbb{R}^n_{\geq 0})^n$ *be a restricted altruistic social context.* Γ *is* $(7/3, 2/3, \alpha)$*-smooth.*

We need the following lemma for the proof of Theorem 115.

Lemma 116. *For every two integers* $x, y \in \mathbb{N}$

$$(x+1)y + xy \leq \frac{7}{3}y^2 + \frac{2}{3}x^2 \tag{5.9}$$

Proof. We write x as $(y + z)$, where $z \in \mathbb{Z}$. Substituting this into (5.9) results in

$$((y+z)+1)y + (y+z)y \leq \frac{7}{3}y^2 + \frac{2}{3}(y+z)^2, \tag{5.10}$$

which can be rewritten as

$$\frac{2}{3}yz + y - y^2 \leq \frac{2}{3}z^2. \tag{5.11}$$

When $z \leq 0$, (5.11) is easily seen to hold, since the first term of the left hand side is non-positive, the second and third terms together are non-positive (taking into account that y is non-negative), and the right hand side is non-negative.

When $z > 0$ and $y \leq z$, (5.11) holds because

$$\frac{2}{3}yz + y - y^2 \leq \frac{2}{3}yz \leq \frac{2}{3}z^2.$$

Lastly, when $z > 0$ and $y > z$, (5.11) holds because

$$\frac{2}{3}yz + y - y^2 < \frac{2}{3}y^2 + y - y^2 = y - \frac{1}{3}y^2.$$

When $y \geq 3$, the latter expression is non-positive and thus not exceeds $\frac{2}{3}z^2$. For the values of (y, z) not yet covered in this case analysis (i.e., $(y, z) \in \{(1, 1), (2, 1)\}$, it can be checked by hand that (5.11) holds. □

Proof of Theorem 115. Let $\Gamma = (n, \Sigma, c)$ be a linear congestion game, let $s, s^* \in \Sigma$ be such that s^* minimizes C. By Lemma 89, we may assume without loss of generality that $d_j(k) = k$ for all $j \in [m]$. We show that (5.2) holds for $(\lambda_2, \mu_2) = (7/3, 2/3)$. We remind the reader that we write $P_j(s)$ to denote the set of players choosing facility $j \in [m]$ under strategy profile $s \in \Sigma$.

Observe that $c_i(s_i^*, s_{-i}) \leq \sum_{j \in s_i^*}(|P_j(s)| + 1)$. Taking the sum over all players, we obtain

$$\begin{aligned}
\sum_{i \in [n]} (c_i(s_i^*, s_{-i}) - c_i(s)) \quad &\leq \quad \sum_{i \in [n]} \sum_{j \in s_i^*} (|P_j(s)| + 1) - C(s) && (5.12)\\
&= \quad \sum_{j \in [m]} |P_j(s^*)|(|P_j(s)| + 1) - C(s). && (5.13)
\end{aligned}$$

Let $i \in [n]$. Note that $|P_j(s^*, s_{-i})| = |P_j(s)| + 1$ for $j \in s_i^* \setminus s_i$, $|P_j(s^*, s_{-i})| = |P_j(s^*, s_{-i}| - 1$ for $j \in s_i \setminus s_i^*$ and $|P_j(s^*, s_{-i})| = |P_j(s)|$ if $j \in s_i^* \cap s_i$. Using these relations, we obtain

$$\begin{aligned}
&\sum_{i' \in [n] \setminus \{i\}} \alpha_{i,i'}(c_{i'}(s_i^*, s_{-i}) - c_{i'}(s))\\
&\quad = \sum_{i' \in [n] \setminus \{i\}} \left(\sum_{j \in s_{i'} \cap (s_i^* \setminus s_i)} \alpha_{i,i'} - \sum_{j \in s_{i'} \cap (s_i \setminus s_i^*)} \alpha_{i,i'} \right)\\
&\quad = \sum_{j \in s_i^* \setminus s_i} \sum_{i' \in [n] \setminus \{i\} : j \in s_{i'}} \alpha_{i,i'} - \sum_{j \in s_i \setminus s_i^*} \sum_{i' \in [n] \setminus \{i\} : j \in s_{i'}} \alpha_{i,i'}.
\end{aligned}$$

Summing over all players and exploiting that in the restricted case $0 \leq \alpha_{i,i'} \leq 1$ for all $i, i' \in [n]$, $i \neq i'$, we can bound

$$\begin{aligned}
&\sum_{i\in[n]} \left(\sum_{j\in s_i^*\setminus s_i} \sum_{i'\in[n]\setminus\{i\}:j\in s_{i'}} \alpha_{i,i'} - \sum_{j\in s_i\setminus s_i^*} \sum_{i'\in[n]\setminus\{i\}:j\in s_{i'}} \alpha_{i,i'} \right) \\
&\quad \leq \sum_{i\in[n]} \sum_{j\in s_i^*} \sum_{i'\in[n]:j\in s_{i'}} 1 \\
&\quad = \sum_{j\in[m]} |P_j(s)||P_j(s^*)|. \qquad (5.14)
\end{aligned}$$

Combining (5.13) and (5.14) and using Lemma 116, we conclude that

$$\sum_{i,i'\in[n]} \alpha_{i,i'}(c_{i'}(s_i^*, s_{-i}) - c_{i'}(s)) \leq \frac{7}{3}C(s^*) + (\frac{2}{3} - 1)C(s),$$

as desired. □

Applying Theorem 112, we obtain:

Corollary 117. *The coarse price of anarchy the class of α-altruistic extensions of singleton linear congestion games, for restricted altruistic social contexts α, is at most 7.*

5.5 Singleton Linear Congestion Games with Identical Delay Functions

We derive a better smoothness result for singleton linear congestion games in which all facilities have identical delay functions.

Theorem 118. *Let $\Gamma = (n, \Sigma, c)$ be a singleton congestion game in which all facilities have identical delay functions and let $\alpha \in (\mathbb{R}^n_{\geq 0})^n$ be a restricted altruistic social context for Γ. Γ is $(\varphi^2, 1/\varphi^2, \alpha)$-smooth, where $\varphi = (1 + \sqrt{5})/2$ is the golden ratio.*

In most smoothness-based proofs for congestion games, one first massages the smoothness condition to derive an equivalent condition summing over all facilities (instead of players), after which one establishes smoothness by reasoning for each facility separately. If we follow this approach here naively, we again obtain an upper bound of

7. In order to improve this bound, we use an *amortized* argument here to derive our improved bound.

For the proof of Theorem 118, we need the following lemma.

Lemma 119. *Let* $\varphi = \frac{1+\sqrt{5}}{2}$ *be the golden ratio. For every two integers* $x, y \in \mathbb{N}_{\geq 0}$*,* $2xy - \varphi x + \varphi^2 y \leq \frac{1}{\varphi^2}x^2 + \varphi^2 y^2$.

Proof. It is easy to check that the claim holds if $x = 0$ or $y = 0$. Let $x \geq 1$ and $y \geq 1$. Recall that $1 + \varphi = \varphi^2$. We have

$$\begin{aligned}
\varphi^2 y^2 - 2xy + \frac{1}{\varphi^2}x^2 + \varphi x - \varphi^2 y &= \left(\varphi y - \frac{1}{\varphi}x\right)^2 + \varphi x - (1+\varphi)y \\
&\geq 2\varphi y - \frac{2}{\varphi}x - 1 + \varphi x - (1+\varphi)y \\
&= (\varphi - 1)y + \left(\varphi - \frac{2}{\varphi}\right)x - 1 \\
&= \frac{1}{\varphi}y + \left(1 - \frac{1}{\varphi}\right)x - 1 \\
&\geq 0,
\end{aligned}$$

where the first inequality holds because $z^2 \geq 2z - 1$ for every $z \in \mathbb{R}$ and the last inequality holds because $x \geq 1$ and $y \geq 1$. □

Proof of Theorem 118. Let $\lambda = \varphi^2$ and let $\mu = 1/\varphi^2$. Let $s^* \in \Sigma$ be a strategy profile minimizing C. We assume that $\alpha_{i,i} = 1$ without loss of generality. To satisfy the smoothness condition (5.2) it therefore suffices to show that

$$\sum_{i \in [n]} \left(c_i(s_i^*, s_{-i}) + \sum_{i' \in [n] \setminus \{i\}} \alpha_{i,i'}(c_{i'}(s_i^*, s_{-i}) - c_i(s)) \right) \leq \lambda C(s^*) + \mu C(s).$$

We rewrite the left hand side of the above inequality as follows.

$$\begin{aligned}
&\sum_{i \in [n]} \left(c_i(s_i^*, s_{-i}) + \sum_{i' \in [n] \setminus \{i\}} \alpha_{i,i'}(c_{i'}(s_i^*, s_{-i}) - c_{i'}(s)) \right) \\
&\quad = \sum_{i \in [n]} \left(1 + \sum_{i' \in [n] \setminus \{i\}} |s_i^* \cap s_{i'}| + \sum_{i \in [n] \setminus \{i\}} \alpha_{i,i'}|s_i^* \cap s_{i'}| - \alpha_{i,i'}|s_i \cap s_{i'}| \right)
\end{aligned} \tag{5.15}$$

$$= n + \sum_{i \in [n]} \sum_{i' \in [n] \setminus \{i\}} ((1 + \alpha_{i,i'})|s_i^* \cap s_{i'}| - \alpha_{i,i'}|s_i \cap s_{i'}|) \tag{5.16}$$

Note that

$$C(s) = \sum_{i \in [n]} \sum_{i' \in [n]} |s_i \cap s_j| + \sum_{i \in [n]} |s_i| = \sum_{i \in [n]} \sum_{i' \in [n]} |s_i \cap s_j| + n,$$

and $C(s^*)$ can be expressed similarly.

Thus, the smoothness condition is equivalent to

$$\begin{aligned}
&\sum_{i \in [n]} \sum_{i' \in [n] \setminus \{i\}} (\lambda |s_i^* \cap s_{i'}^*| + (\mu + \alpha_{i,i'})|s_i \cap s_{i'}| - (1 + \alpha_{i,i'})|s_i^* \cap s_{i'}|) \\
&\qquad + (\lambda + \mu - 1)n \\
&\qquad \geq 0.
\end{aligned} \tag{5.17}$$

To prove this inequality we first derive, for $j \in [m]$,

$$\begin{aligned}
\lambda &|\{(i,i') \mid i,i' \in [n], i \neq i', s_i^* = s_{i'}^* = j\}| \\
&\quad + (\mu + \alpha_{i,i'})|\{(i,i') \mid i,i' \in [n], i \neq i', s_i = s_{i'} = j\}| \\
&\quad - (1 + \alpha_{i,i'})|\{(i,i') \mid i,i' \in [n], i \neq i', s_i = s_{i'}^* = j\}| \\
&\quad + (\lambda + \mu - 1)|\{i \mid s_i = \{j\}\}| \\
&\geq \lambda|\{(i,i') \mid i,i' \in [n], i \neq i', s_i^* = s_{i'}^* = j\}| \\
&\quad + \mu|\{(i,i') \mid i,i' \in [n], i \neq i', s_i = s_{i'} = j\}| \\
&\quad - 2|\{(i,i') \mid i,i' \in [n], i \neq i', s_i = s_{i'}^* = j\}| \\
&\quad + (\lambda + \mu - 1)|\{i \mid s_i = \{j\}\}| \\
&\geq 2\lambda \frac{|P_j(s^*)|(|P_j(s^*)| - 1)}{2} + 2\mu \frac{|P_j(s)|(|P_j(s)| - 1)}{2} - 2|P_j(s)||P_j(s^*)| \\
&\quad + (\lambda + \mu - 1)|P_j(s)| \\
&= \lambda|P_j(s^*)|^2 + \mu|P_j(s)|^2 - 2|P_j(s)|P_j(s^*)| + (\lambda - 1)|P_j(s)| \\
&\geq 0.
\end{aligned}$$

Where the last inequality follows from Lemma 119 and the observation that $\varphi^2 - 1 = \varphi$.

Summing over $j \in [m]$, we obtain

$$\sum_{j \in [m]} (\lambda|\{(i,i') \mid i,i' \in [n], i \neq i', s_i^* = s_{i'}^* = j\}|$$

$$
\begin{aligned}
&+ (\mu + \alpha_{i,i'})|\{(i,i') \mid i,i' \in [n], i \neq i', s_i = s_{i'} = j\}| \\
&- (1 + \alpha_{i,i'})|\{(i,i') \mid i,i' \in [n], i \neq i', s_i = s^*_{i'} = j\}| \\
&+(\lambda + \mu - 1)|\{i \mid s_i = \{j\}\}|) \\
\geq\ & 0.
\end{aligned}
$$

The left hand side of the above inequality is equal to the left hand side of (5.17). This proves that (5.17) holds. □

Applying Theorem 112, we obtain:

Corollary 120. *The coarse price of anarchy of the class of α-altruistic extensions of singleton linear congestion games in which all facilities have identical delay functions, for restricted altruistic social contexts α, is at most $\varphi^3 \approx 4.2361$.*

5.6 Minsum Machine Scheduling

Cole et al. [2011] show that minsum scheduling games with unrelated machines (without altruism) are $(2, 1/2)$-smooth, resulting in a coarse price of anarchy of at most 4.[3] Hoeksma and Uetz [2012] prove that minsum machine scheduling games with related machines (without altruism) are $(2, 0)$-smooth, leading to the conclusion that the coarse price of anarchy is at most 2.

The following example shows that for unrestricted social contexts $\alpha > \mathbf{0}$, the pure price of anarchy of α-altruistic extensions of minsum machine scheduling games with related machines is unbounded.

Example 121. Fix a number $M \in \mathbb{R}_{>0}$ arbitrarily, and consider the minsum machine scheduling game with related machines $\Gamma = (n, m, a, w, \preceq)$, with $n = m = 2$. The speeds are given by $a_1 = M, a_2 = 1$. The job lengths are given by $w_1 = w_2 = 1$. Suppose that the altruistic social context is as follows: $\alpha_{1,1} = \alpha_{2,2} = \alpha_{2,1} = 0$, $\alpha_{1,2} = 1$. Then the strategy profile $(1, 2)$ is a pure equilibrium. The social cost of this equilibrium is $M + 1$. When $M \geq 2$, it is a social optimum to schedule both jobs on machine 2, and this schedule achieves a social cost of 3. Therefore, for $M \geq 2$, the pure price of anarchy of the α-altruistic extension of Γ is $M + 1/3$. Because M can be picked arbitrarily large, this shows that the pure price of anarchy is unbounded.

[3] More precisely, this is shown to hold for the more general case when the social cost is an arbitrary nonnegative linear combination of the player's cost. From a scheduling game instance described by Correa and Queyranne [2012], it follows that this bound is tight, i.e., that the coarse price of anarchy is actually exactly 4.

We prove constant upper bounds on the coarse price of anarchy for α-altruistic extensions of minsum scheduling games, when α is a restricted altruistic social context.

Theorem 122. *Let $\Gamma = (n, m, p, \preceq)$ be a minsum machine scheduling game, let $\alpha \in (\mathbb{R}^n_{\geq 0})^n$ be a restricted altruistic social context for Γ, and let $x \in \mathbb{R}_{>0}$. Γ is $(2 + x, 1/2 + 1/x, \alpha)$-smooth if Γ is a minsum scheduling game with unrelated machines. Γ is $(2 + x, 1/x, \alpha)$-smooth if Γ is a minsum scheduling game with related machines.*

Proof. Let Σ be the set of strategy profiles of Γ. For $i \in [n], j \in [m], s \in \Sigma$, we define the value $N(i, j, s) = |\{i' \mid i \prec_j i', s_{i'} = j\}|$. Note that the social cost of a strategy profile $s \in \Sigma$ can then be written as

$$C(s) = \sum_{j \in [m]} \sum_{i : s_i = j} (N(i, j, s) + 1) p_{i,j}.$$

We use the proof template described in Section 5.3.2. In Hoeksma and Uetz [2012] it is proved that the base game for the case of related machines is $(2, 0)$-smooth, and from the proof of Theorem 3.2 in Cole et al. [2011], it follows that the base game for the case of unrelated machines is $(2, 1/2)$-smooth. Thus, let $s^* \in \Sigma$ be an arbitrary optimal schedule and let $s \in \Sigma$ be an arbitrary schedule. It suffices to show that

$$\sum_{i \in [n]} \sum_{i' \in [n] \setminus \{i\}} \alpha_{i,i'} (c_{i'}(s_i^*, s_{-i}) - c_{i'}(s)) \leq x C(s^*) + \frac{C(s)}{x}$$

for all $x > 0$.

Let $P_1 = \{(i, i') \mid s_i^* = s_{i'}, s_i^* \neq s_i, i \prec_{s_i^*} i'\}$, and let $P_2 = \{(i, i') \mid s_i = s_{i'}, s_i^* \neq s_{i'}, i \prec_{s_i} i'\}$. Note that for a pair of players (i, i') that is not in $P_1 \cup P_2$, we have $c_{i'}(s_i^*, s_{-i}) - c_{i'}(s) = 0$, and for $(i, i') \in P_2$ it holds that $\alpha_{i,i'}(c_{i'}(s_i^*, s_{-i}) - c_{i'}(s)) \leq 0$. Therefore:

$$\begin{aligned}
\sum_{i \in [n]} \sum_{i' \in [n] \setminus \{i\}} \alpha_{i,i'} (c_{i'}(s_i^*, s_{-i}) - c_{i'}(s)) &= \sum_{(i,i') \in P_1 \cup P_2} \alpha_{i,i'} (c_{i'}(s_i^*, s_{-i}) - c_{i'}(s)) \\
&\leq \sum_{(i,i') \in P_1} \alpha_{i,i'} (c_{i'}(s_i^*, s_{-i}) - c_{i'}(s)) \leq \sum_{(i,i') \in P_1} (c_{i'}(s_i^*, s_{-i}) - c_{i'}(s)) \\
&= \sum_{(i,i') \in P_1} p_{i,s_i^*}.
\end{aligned}$$

We now rewrite this last expression into a summation over the machines. We obtain:

$$\sum_{i \in [n]} \sum_{i' \in [n] \setminus \{i\}} \alpha_{i,i'} (c_{i'}(s_i^*, s_{-i}) - c_{i'}(s)) \leq \sum_{j \in [m]} \sum_{(i,i') \in P_1 : s_i^* = j} p_{i,j}$$

$$
\begin{aligned}
&= \sum_{j \in [m]} \sum_{i \in [n] : s_i^* = j} \sum_{i' \in [n] : (i,i') \in P_1} p_{i,j} \\
&= \sum_{j \in [m]} \sum_{\substack{i \in [n] : s_i^* = j, \\ s_i \neq j}} \sum_{i' : S_{i'} = j, i \prec_{s_i} i'} p_{i,j} \\
&\leq \sum_{j \in [m]} \sum_{\substack{i \in [n] : s_i^* = j, \\ s_i \neq j}} N(i,j,s) p_{i,j} \\
&\leq \sum_{j \in [m]} \sum_{\substack{i \in [n] : s_i^* = j, \\ s_i \neq j}} (xN(i,j,s^*) + x - 1 + N(i,j,s) - xN(i,j,s^*) - x + 1) p_{i,j} \\
&\leq \sum_{j \in [m]} \sum_{i \in [n] : s_i^* = j} (x(N(i,j,s^*) + 1) - 1) p_{i,j} \\
&\quad + \sum_{j \in [m]} \sum_{\substack{i \in [n] : s_i^* = j, \\ s_i \neq j}} (N(i,j,s) - xN(i,j,s^*) - x + 1) p_{i,j} \\
&\leq \sum_{j \in [m]} \sum_{i \in [n] : s_i^* = j} (x(N(i,j,s^*) + 1) - 1) p_{i,j} \\
&\quad + \sum_{j \in [m]} \sum_{\substack{i \in [n] : s_i^* = j, \\ s_i \neq j, \\ N(i,j,s) > xN(i,j,s^*) + x - 1}} \lceil N(i,j,s) - xN(i,j,s^*) - x + 1 \rceil p_{i,j} \\
&\leq \sum_{j \in [m]} \sum_{i \in [n] : s_i^* = j} (x(N(i,j,s^*) + 1) - 1) p_{i,j} \\
&\quad + \sum_{j \in [m]} \sum_{\substack{i \in [n] : s_i^* = j, \\ s_i \neq j, \\ N(i,j,s) > xN(i,j,s^*) + x - 1}} \lceil N(i,j,s) - xN(i,j,s^*) - x + 1 \rceil p_{i,j}.
\end{aligned}
$$

Consider a player $i \in [n]$ and machine $j \in [m]$ such that it holds that $s_i^* = j$, $s_i \neq j$, and $N(i,j,s) > xN(i,j,s^*) + x - 1$. Let $S(i,j)$ be the set consisting of the $\lceil N(i,j,s) - xN(i,j,s^*) - x \rceil$ players $i' \in [n]$ with smallest processing time on j such that $i' \succ_j i$ and $s_{i'} = j$. Note that $S(i,j)$ is well defined in the sense that this number of players with these properties exist, because $N(i,j,s) > xN(i,j,s^*) + x - 1$ implies $\lceil N(i,j,s) - xN(i,j,s^*) - x \rceil \geq 0$, and because there exist $N(i,j,s) \geq |S(i,j)|$ players $i' \in [n]$ such that $i' \succ_j i$ and $s_{i'} = j$. Note that for every player $i' \in S(i,i')$, it holds that $N(i',j,s) \geq N(i,j,s) - |S(i,j)| > xN(i,j,s^*) + x - 1$. We use this to upper bound the above as follows:

$$\begin{aligned}
&\sum_{i\in[n]}\sum_{i'\in[n]\setminus\{i\}} \alpha_{i,i'}(c_{i'}(s_i^*, s_{-i}) - c_{i'}(s)) \\
&\quad\le \sum_{j\in[m]}\sum_{i\in[n]:s_i^*=j} (x(N(i,j,s^*)+1)-1)p_{i,j} \\
&\qquad + \sum_{j\in[m]} \sum_{\substack{i\in[n]:s_i^*=j,\\ s_i\neq j,\\ N(i,j,s)>xN(i,j,s^*)+x-1}} \left(p_{i,j} + \sum_{i'\in S(i,j)} p_{i,j}\right) \\
&\quad\le xC(s^*) + \sum_{j\in[m]}\sum_{i'\in[n]:s_{i'}=j} \sum_{\substack{i\in[n]:s_i^*=j\\ s_i\neq j\\ i\prec_j i',\\ N(i',j,s)>xN(i,j,s^*)+x-1}} p_{i,j} \\
&\quad\le xC(s^*) + \sum_{j\in[m]}\sum_{i'\in[n]:s_{i'}=j} \sum_{\substack{i\in[n]:s_i^*=j\\ s_i\neq j\\ i\prec_j i',\\ N(i',j,s)>xN(i,j,s^*)+x-1}} p_{i',j}.
\end{aligned}$$

The next step in the derivation is made by observing that for each job $i' \in [n]$ and each machine $j \in [m]$ such that $s_{i'} = j$, there are at most $\lceil (N(i',j,s) - x + 1)/x\rceil$ players $i \in [n]$ such that $i \prec_j i'$, $s_i^* = j$, $s_i \neq j$ and $N(i',j,s) > xN(i,j,s^*) + x - 1$. To see this, assume for contradiction that there are *more* than $\lceil (N(i',j,s) - x + 1)/x\rceil$ jobs $i \prec_j i'$ such that $s_i^* = j$, $s_i \neq j$ and $N(i',j,s) > xN(i,j,s^*) + x - 1$. Let $i \in [n]$ be the $(\lceil (N(i',j,s) - x + 1)/x\rceil + 1)$-th largest player, with respect to processing time on j, for which these three properties hold. Then, there are at least $(\lceil (N(i',j,s)-x+1)/x\rceil+1)$ players scheduled on machine j that have these properties and that are scheduled after i under s^*. Therefore, we have that $xN(i,j,s^*) + x - 1 \ge x\lceil (N(i',j,s) - x + 1)/x\rceil + x - 1 \ge N(i',j,s)$, which is a contradiction. Through this observation, we obtain:

$$\begin{aligned}
&\sum_{i\in[n]}\sum_{i'\in[n]\setminus\{i\}} \alpha_{i,i'}(c_{i'}(s_i^*, s_{-i}) - c_{i'}(s)) \\
&\quad\le xC(s^*) + \sum_{j\in[m]}\sum_{i'\in[n]:s_{i'}=j} \left\lceil \frac{N(i',j,s)-x+1}{x} \right\rceil p_{i',k} \\
&\quad\le xC(s^*) + \sum_{j\in[m]}\sum_{i'\in[n]:s_{i'}=j} \frac{1}{x}(N(i',j,s)+1)p_{i',j}
\end{aligned}$$

$$= xC(s^*) + \frac{C(s)}{x},$$

where we use for the second inequality the basic fact that $\lceil a \rceil \leq a+1$ for all $a \in \mathbb{R}$. □

Corollary 123. *The coarse price of anarchy of the class of α-altruistic extensions of minsum machine scheduling games with unrelated machines, for restricted altruistic social contexts α, is at most $12 + 8\sqrt{2} \approx 23.3137$. The coarse price of anarchy of the class of α-altruistic extensions of minsum machine scheduling games with related machines, for restricted altruistic social contexts α, is at most $4 + 2\sqrt{3} \approx 7.4641$.*

Proof. For the first claim, apply Theorem 122 with $x = 2 + 2\sqrt{2}$, and then apply Theorem 112. For the second claim, apply Theorem 122 with $x = 1 + \sqrt{3}$, and then apply Theorem 112. □

5.7 Generalized Second Price Auctions

As said in Remark 107, we study the price of anarchy in this chapter with respect to the sum-of-utilities social welfare function (1.4). However, in case of generalized second price auctions, we consider the auctioneer to be a player in the auction with a single strategy, of which its utility is the total price charged to the other players. This is a standard convention when studying the price of anarchy of auction games (see also the discussion in Section 2.3 of Chapter 2). Therefore, for a generalized second price auction $\Gamma = (n, m, v, \beta, \preceq)$ and strategy profile s of Γ, it holds (asuming $m = n$ w.l.o.g.) that

$$U(s) = \sum_{j \in [m]} \beta_j v_{r_j(s)}.$$

We remark that in our definition of generalized second price auctions (Definition 110) we defined the strategy spaces such that a player cannot play a strategy greater than its valuation. This is equivalent to the no-overbidding assumption that we adopted in Chapter 2 for the case of uniform price auctions, and it is a standard assumption in the setting of generalized second price auctions. We refer to the discussion at the end of Section 2.3 for a justification of the no-overbidding assumption.

We prove an upper bound of $O(n)$ on the coarse price of anarchy of α-altruistic extensions of generalized second price auctions if α is a restricted altruistic social context.

Theorem 124. *Let $\Gamma = (n, m, v, \beta, \preceq)$ be a generalized second price auction, let $\alpha \in (\mathbb{R}^n)^n$ be a restricted altruistic social context for Γ. Then Γ is $(1/2, n, \alpha)$-smooth.*

Proof. Let Σ be the set of strategy profiles of Γ, and let $s^*, s \in \Sigma$. We may assume w.l.o.g. that $m = n$ by removing slots or by adding dummy slots with click-through rate 0. By renaming the players, we may moreover assume that for all $j \in [m]$, $r_j(s) = j$.

We apply the proof template of Section 5.3.2. Generalized second price auctions are known to be $(1/2, 1)$-smooth [Roughgarden, 2012], so it holds that

$$\sum_{i \in [n]} u_i(s_i^*, s_{-i}) \geq \frac{1}{2} U(s^*) - U(b).$$

We bound the altruistic part of the smoothness condition as follows.

$$\begin{aligned}
\sum_{i \in [n]} \sum_{i' \in [n] \setminus \{i\}} \alpha_{i,i'}(u_{i'}(s_i^*, s_{-i}) - u_{i'}(s)) &\geq \sum_{i \in [n]} \sum_{i' \in [n] \setminus \{i\}} \alpha_{i,i'}(-u_{i'}(b)) \\
&\geq \sum_{i \in [n]} \sum_{i' \in [n] \setminus \{i\}} \alpha_{i,i'}(-\beta_{i'} v_{i'}) \\
&\geq \sum_{i \in [n]} \sum_{i' \in [n] \setminus \{i\}} -\beta_{i'} v_{i'} \\
&\geq -(n-1)U(b).
\end{aligned}$$

Combining these inequalities proves $(\frac{1}{2}, n, \alpha)$-smoothness. $\square$

Theorem 112 now gives us:

Corollary 125. *The coarse price of anarchy of the class of α-altruistic extensions of generalized second price auctions, for restricted altruistic social contexts α, is at most $2n + 2$.*

The above smoothness analysis turns out to be asymptotically tight, as the next example shows.

Example 126. Let $m \in \mathbb{N}_{>0}$ be odd, let $\epsilon \in \mathbb{R}_{>0}$ be a sufficiently small number, and let Γ be a generalized second price auction $(n, m, v, \beta, \preceq)$ with $n = m$ bidders and m slots. Let $\beta_j = 1$ for all $j \in [(m-1)/2]$, let $\beta_j = \epsilon$ for all $j \in [m-1] \setminus [(m-1)/2]$. And let $\beta_m = \epsilon^2$. Further, define the player valuations as $v_i = \epsilon$ for $i \in [(m-3)/2]$, $v_i = 1$ for $i \in [m-1] \setminus [(m-3)/2]$ and $v_m = 0$. In particular, this implies that $\Sigma_m = \{0\}$.

Let Σ be the set of strategy profiles of Γ and let $s^* \in \Sigma$ be a strategy profile where every player chooses its valuation as its strategy. In this strategy profile, we have:

- $r_i(s^*) \in [m] \setminus \{1, \dots, (m-1/2), m\}$ for $i \in [(m-1)/2]$,
- $r_i(s^*) \in [m-1]$ for $i \in [m-1] \setminus [(m-1)/2]$,
- $r_m(s^*) = m$.

The total social welfare of s^* is therefore $U(b^*) = (m-1)/2 + \epsilon + (m-3)\epsilon^2/2$.

s is a strategy profile such that $s_1 > \cdots > s_m = 0$ and $s_1 < s_i^*$ for every $i \in [m-1]$. Note that this implies that $s_i < \epsilon$ for all $i \in [n]$ because a player does not play a strategy exceeding its valuation. Thus, by choosing ϵ sufficiently small, we ensure that the strategies of s become arbitrarily small but induce the order as indicated above. We have

$$U(s) = \sum_{i \in [n]} \beta_i v_i = (n-1)\epsilon + 1.$$

If player $i \in [n] \setminus [(m-1)/2]$ changes its strategy from s_i to s_i^*, then $r_1(s_i^*, s_{-i}) = i$, and therefore $r_j(s_i^*, s_{-i}) = r_{j-1}(s)$ for all $j \in [i-1]$. Observe that player $(m-1)/2$ (that has valuation $v_{(m-1)/2} = 1$) wins slot $(m-1)/2$ under s, with $\beta_{(m-1)/2} = 1$, but wins slot $(m+1)/2$ under (s_i^*, s_{-i}), with $\beta_{(m+1)/2} = \epsilon$.

We consider the left-hand side of the smoothness condition 5.3.

$$\sum_{i \in [n]} \sum_{i' \in [n]} \alpha_{i,i'} (u_{i'}(s_i^*, s_{-i}) - u_{i'}(s)).$$

We bound each term in the above summation separately. Let $\Delta_{i,i'}$ be the term corresponding to players $i, i' \in [n]$, i.e.,

$$\Delta_{i,i'} = \alpha_{i,i'} (u_{i'}(s_i^*, s_{-i}) - u_{i'}(s)). \tag{5.18}$$

Note that if i' or i is n, then $\Delta_{i,i'} = 0$. For $i', i \in [n-1]$, we distinguish four cases depending on the position of $i' \in [n]$ with respect to $i \in [n]$:

Case $i' > i$: The deviation of i does not affect player i' and thus $\Delta_{i,i'} = 0$.

Case $i' = i$: Player i wins slot 1 under (s_i^*, s_{-i}), and pays s_1. Thus

$$\Delta_{i,i} = \alpha_{i,i}(\beta_1(v_i - s_1) - \beta_i(v_i - s_{i+1})).$$

Case $i' = i - 1$: We have $r_{i'}(s) = r_{i'+1}(s_i^*, s_{-i}^*) = i'$, and i' pays $s_{i'+2}$ under (s_i^*, s_{-i}), so

$$\Delta_{i,i'} = \alpha_{i,i'}(\beta_{i'+1}(v_{i'} - s_{i'+2}) - \beta_j(v_{i'} - s_{i'+1})).$$

Case $i' < i-1$: We have $r_{i'}(s) = r_{i'+1}(s_i^*, s_{-i}^*) = i'$, and i' pays $s_{i'+1}$ (as under s). Thus

$$\begin{aligned}\Delta_{i,i'} &= \alpha_{i,i'}(\beta_{i'+1}(v_{i'} - s_{i'+1}) - \beta_{i'}(v_{i'} - s_{i'+1}) \\ &= \alpha_{i,i'}(\beta_{i'+1} - \beta_i')(v_{i'} - s_{i'+1}).\end{aligned}$$

Choosing ϵ sufficiently small we can make sure that the total contribution of the payments in each of the above four cases is negligible.

We consider the restricted social context, so without loss of generality we assume that $\alpha_{i,i} = 1$ for $i \in [n]$. Recall that $\beta_1 = 1$. Ignoring the effect of the payments (which we just argued is negligible if we make ϵ small enough), the total contribution to the left-hand side of the smoothness definition is

$$\begin{aligned}&\sum_{i\in[n-1]} \left(\alpha_{i,i}(\beta_1 v_i - \beta_i v_i) + \sum_{i'\in[i-1]} \alpha_{i,i'}(\beta_{i'+1} - \beta_{i'})v_{i'}\right) \\ &\quad = \sum_{i\in[n-1]} (1-\beta_i)v_i + \sum_{i\in[n-1]} \sum_{i'\in[i-1]} \alpha_{i,i'}(\beta_{i'+1} - \beta_{i'})v_{i'}.\end{aligned}$$

Note that $\beta_{i'+1} - \beta_{i'} = 0$ for all $i' \neq [n-2] \setminus \{(m-1)/2\}$ and $\beta_{(m+1)/2} - \beta_{(m-1)/2} = \epsilon - 1$. The above expression thus simplifies to

$$\begin{aligned}&\sum_{i\in[n-1]\setminus[(m-1)/2]} (1-\epsilon)v_i \\ &\quad + \sum_{i\in[n-1]\setminus[(m-1)/2]} \alpha_{i,(m-1)/2}(\beta_{(m+1)/2} - \beta_{(m-1)/2})v_{(m-1)/2} \\ &\quad = \frac{m-1}{2}(1-\epsilon) - (1-\epsilon) \sum_{i\in[n-1]\setminus[(m-1)/2]} \alpha_{i,(m-1)/2} \\ &\quad = (1-\epsilon)\left(\frac{m-1}{2} - \sum_{i\in[n-1]\setminus[(m-1)/2]} \alpha_{i,(m-1)/2}\right).\end{aligned}$$

By setting $\alpha_{i,(m-1)/2} = 1$ for every $i \in [n-1]$, the above contribution is equal to zero.

To show (λ, μ, α)-smoothness, we need to lower bound the latter expression by $\lambda U(s^*) - (\mu+1)U(s)$. That is, λ and μ need to satisfy

$$(1-\epsilon)\left(\frac{m-1}{2} - \sum_{i\in[n-1]\setminus[(m-1)/2]} \alpha_{i,(m-1)/2}\right)$$

$$
\begin{aligned}
&= 0 \\
&\geq \lambda \left(\frac{m-1}{2} + \epsilon + \frac{m-3}{2} \epsilon^2) \right) - (\mu + 1)((n-1)\epsilon + 1),
\end{aligned}
$$

which implies (letting $\epsilon \to 0$)

$$\mu + 1 \geq \lambda \frac{m-1}{2}.$$

This provides an asymptotic lower bound of

$$\frac{1+\mu}{\lambda} \geq \frac{m-1}{2} = \frac{n-1}{2}$$

on the best possible coarse price of anarchy upper bound achievable by our smoothness framework.

Chapter 6

Efficiency of Equilibria in Procurement Auctions with Spiteful Players*

In this chapter, we shift our focus from altruistic behavior to *spiteful* behavior. The necessary preliminary knowledge from Chapter 1 consists of Section 1.3 up to 1.3.1.1, and Sections 1.3.1.5, 1.3.1.8, and 1.3.1.11.

We study a model of spite, applied to *procurement auctions* (also known as *reverse auctions*). These form a well-established class of games. An *auctioneer* would like to purchase an item or service from exactly one of the players (also called *sellers*), each of which can supply or perform it equally well. Different players may have different *costs* to provide the item or service. They will submit *bids* to the auctioneer, which may misrepresent their true valuations. The set of possible bids that a player may submit, forms its strategy set. Based on the bids, the auctioneer selects a *winner* and a *payment* to the winner, which must be at least the winner's bid. The two most common mechanisms are the *first price procurement auctions* and the *second price procurement auctions*. In both, the winner is the player with the lowest bid. In a first price auction, the player is paid exactly its bid, while in a second price auction, the player is paid the second-lowest bid, i.e., the bid of the cheapest competitor. The utility of a player is defined as the payment it receives from the auctioneer, minus its valuation for providing the item or service. Our study is carried out for a full information model, for both *first*

*This chapter is based on unpublished work with Po-An Chen and David Kempe *price and second price* procurement auctions.

As a result of spite, players that lose the auction might derive negative utility when another player wins. Such spite will alter the choice of strategy of a player.

Our interest is in the loss in social welfare as a consequence of spite, and we therefore study a variation of the price of anarchy of procurement auctions with spiteful players, called the *price of spite* (defined in Section 6.3). We generalize the class of procurement auctions slightly: for each player we introduce a penalty multiplier, which is a positive real number. A player's effective bid in the auction is then its submitted bid multiplied by its penalty multiplier. The player with the lowest effective bid is the winner of the auction. However, in a first price generalized procurement auction, we still pay the winner a price that its equal to its *submitted* bid, rather than its *effective* bid. Likewise, for a second price auction, we pay the winner a price equal to the highest bid that it could have submitted in order to still win the auction. Thus, first price and second price generalized procurement auctions (without spite) are formally defined as follows.

Definition 127 (First price and second price generalized procurement auctions)**.** A *first price generalized procurement auction* is a game $\Gamma = (n, \Sigma, u)$ where $\Sigma_i = \mathbb{R}_{\geq 0}$ for all $i \in [n]$. The strategy profiles in $\Sigma_i, i \in [n]$ are alternatively called *bid vectors*. The utility functions u are defined as follows. There is a *tie-breaking rule* $\preceq$ which is a strict total order on n, a penalty multiplier $\mu_i \in \mathbb{R}_{\geq 0}$ and a *cost* or *valuation* $v_i \in \mathbb{R}_{\geq 0}$ for each player $i \in [n]$. For a bid vector $s \in \Sigma$, the *winner* of Γ under s, $w(s)$, is defined as the player $i \in [n]$ such that for all $i' \in [n] \setminus \{i\}$ it holds that $\mu_i v_i < \mu_{i'} v_{i'}$ or $\mu_i v_i = \mu_{i'} v_{i'}, i \prec i'$. For $s \in \Sigma, i \in [n]$,

$$u_i(s) = \begin{cases} s_i - v_i & \text{if } i = w(s) \\ 0 & \text{otherwise .} \end{cases}$$

A *second price generalized procurement auction* is a game (n, Σ, c) where $\Sigma_i = \mathbb{R}_{\geq 0}$ for all $i \in [n]$. The strategy profiles in $\Sigma_i, i \in [n]$ are alternatively called *bid vectors*. The utility functions u are defined as follows. There is a *tie-breaking rule* $\preceq$ which is a strict total order on n, a penalty multiplier $\mu_i \in \mathbb{R}_{\geq 0}$ and a *cost* or *valuation* $v_i \in \mathbb{R}_{\geq 0}$ for each player $i \in [n]$. For $s \in \Sigma$, the *winner* of Γ under s, $w(s)$, is defined in the same way as for first price procurement auctions. The *runner-up* of Γ under s, $r(s)$, is defined as the player $i \in [n] \setminus \{w(s)\}$ such that for all $i' \in [n] \setminus \{i, w(s)\}$ it holds that $\mu_i v_i < \mu_{i'} v_{i'}$ or $\mu_i v_i = \mu_{i'} v_{i'}, i \prec i'$. For $s \in \Sigma, i \in [n]$,

$$u_i(s) = \begin{cases} \frac{\mu_{r(s)} s_{r(s)}}{\mu_i} - v_i & \text{if } i = w(s) \\ 0 & \text{otherwise .} \end{cases}$$

A first price or second price generalized procurement auction Γ can therefore be represented as a quadruple $(n, v, \mu, \preceq)$, where $v = (v_1, \ldots, v_n)$ are the valuations and $\mu = (\mu_1, \ldots, \mu_n)$ are the penalty multipliers.

We study for such auctions the question of how to set the penalty multipliers in order to minimize the price of spite of the auction, given the spite levels of the players. This can be regarded as a mechanism design problem in a restricted domain. However, unlike regular mechanism design problems, we study this problem in a full information setting, and we are not concerned with the notion of *truthfulness*, which is usual in traditional mechanism design.[1]

The solution concepts we are interested in, in this chapter, are *ϵ-equilibria*, and *limit equilibria*. The latter solution concept we will define later on in this chapter. ϵ-equilibria are strategy profiles for which no player can improve its utility by more than ϵ by switching to a different strategy. They are formally defined as follows.

Definition 128 (ϵ-equilibrium)**.** Let $\Gamma = (n, \Sigma, u)$ be a full information utility maximization game and let $\epsilon \in \mathbb{R}_{\geq 0}$. An *$\epsilon$-equilibrium* is defined as a strategy profile $s \in \Sigma$ such that for all $s' \in \Sigma$ it holds that for all $i \in [n]$,

$$u_i(s'_i, s_{-i}) - u_i(s) \leq \epsilon.$$

The set of ϵ equilibria of Γ is denoted by ϵ-PE_Γ. The subscript will sometimes dropped in case no confusion can arise.

For two spiteful players, we provide necessary and sufficient conditions for a strategy profile being an ϵ-equilibrium when $\epsilon \in \mathbb{R}_{>0}$ is small enough, and we quantify the extent to which the spite of the players impact the quality of the equilibria, by deriving the price of spite for ϵ-equilibria when $\epsilon \to 0$. We find expressions for the penalty multipliers that minimize the divergence rate of the price of spite.

For procurement auctions with more than two spiteful players, we characterize all limit equilibria, give a polynomial time algorithm that computes them all, and derive the price of spite.

6.1 Background

Auctions are an area with a huge body of work, spanning decades of research [Krishna, 2002]. Yet, only a much smaller amount of work considers auctions embedded in a larger economic or social context. Typically, the analysis of auctions assumes that any

[1] However, we will study such a problem in Chapter 7. See Section 7.3 for a definition of truthfulness.

participant that does not win in an auction, and does not pay or get paid, will have zero utility. In reality, an outcome wherein a competitor wins could have significant impact on a player. The competitor's profits may lead to a significant future advantage, or an individual may simply resent a competitor and therefore assign negative value to the outcome of that competitor's winning. Such perceptions — whether they be motivated economically or socially — create what is known as *negative externalities*.

When studying auctions, one often makes the assumption that the player does not bid higher than its valuation (when buying an item) or does not bid lower than its valuation (when selling an item) (see e.g. the generalized second price auction studied in Chapter 5 and the uniform price auction studied in Chapter 2). In the presence of negative externalities, such an assumption is no longer justified: If losing in the auction leads to a negative utility for a player, then a smaller negative utility would be preferable if such an outcome can be obtained by overbidding or underbidding.

A point at which we depart from the standard auction model is the introduction of spite: For each player i we introduce a *spite level* α_i that captures how strongly affected i is by another player winning the auction. The model of spite that we use has been used in several prior works [Brandt et al., 2007, Brandt and Weiß, 2002, Maasland and Onderstal, 2007, Morgan et al., 2003, Vetsikas and Jennings, 2007]. Besides the social interpretation of the spite levels of the players, spite can be motivated from an economic perspective as well, in which case spite level of a player may be interpreted as the degree to which losing the auction causes a future economic disadvantage to the player.

We study procurement auctions with competitive externalities in a full information setting, wherein all players know each others' true valuations (see, e.g., Brandt and Weiß [2002]). While incomplete information settings have been more widely studied, they suffer — among others — from the drawback that a player would not even be able to evaluate its own utility, since it does not know another player's valuation.

Auctions with spite among players have been studied before [Brandt et al., 2007, Brandt and Weiß, 2002, Maasland and Onderstal, 2007, Morgan et al., 2003, Vetsikas and Jennings, 2007]. Among these works, the paper by Brandt and Weiß [2002] studies *min-max equilibria* for *complete information* two player auctions, both of which have spite level $\frac{1}{2}$. This is an exception to the remainder of these works, where it is assumed that there is incomplete information, and that the valuations are drawn randomly from some distribution. All these papers are concerned with the case of *uniform spite*, where the spite levels of all players are equal. Uniform spite is interesting as an analysis of the effects of general distrust or future competition between players, but does not take into account the effects of asymmetry among spiteful players on individual behavior.[2]

[2]Special cases of non-uniform spite with Bayesian priors were considered in Chapter 7 of Chen [2011].

Morgan et al. [2003] and Brandt et al. [2007] focus on Bayes-Nash equilibria of first price and second price auctions with uniform spite. The results in these two papers differ mostly in the precise model of the utility of the winner. Vetsikas and Jennings [2007] extend this work to multi-item auctions.

The notion of *spite* and *altruism* as defined here broadly falls into the class of *allocation externalities* in auctions: the utility of a player depends not exclusively on its own allocation, but also on the allocations of other players. There is a large amount of literature on various types of allocation externalities (see, e.g., Jehiel et al. [1999], Brocas [2003]). In specific scenarios, externalities in sponsored search auctions have also been considered [Kempe and Mahdian, 2008, Fotakis et al., 2011].

In [Fiat et al., 2012], "externality-resistant mechanisms" are designed when the spite levels are all below a very small constant, to achieve an approximately equal utility for each spiteful player. The present chapter aims at exploring the effects on the social welfare under a full spiteful range (from selfishness to extreme spitefulness).

Several papers have analyzed the impact of spiteful and altruistic behavior in other game-theoretic settings. In the context of non-atomic congestion games, Chen and Kempe [2008] show that when all players are at least β-altruistic, then the price of anarchy is bounded by $1/\beta$. In the context of network inoculation, Moscibroda et al. [2006] study the effect of Byzantine malicious players. In Chen et al. [2010], altruism to the whole society is considered, where pure equilibria may not exist anymore, so a bound on a measure similar to the price of anarchy for a specific type of dynamics is given.

6.2 Contributions and Outline

Our first result is a characterization of the ϵ-equilibria. More specifically, for two players, we explicitly characterize all ϵ-pure equilibria for any $\epsilon > 0$ in Section 6.4. In the same section, we derive exact expressions on the price of spite for the *limit equilibria* of a given procurement auction with two spiteful players. Limit equilibria are defined formally below. Roughly, they are the limit points of the set of ϵ-equilibria, as $\epsilon \to 0$.

In the auctions we study, each player's bid is multiplied by its penalty multiplier, and the winner is the player with the lowest such modified bid. The multipliers are supposed to mitigate the negative effects of competition, and will be chosen by the auctioneer with knowledge of the spite levels. The mechanism design problem we attempt to solve is to choose the penalty multipliers so as to minimize the worst-case price of spite over all possible valuation vectors. As the price of spite can grow unboundedly as the valuations grow unboundedly at different rates, our goal — stated more precisely — is to find penalty multipliers that minimize the rate of growth of the price of spite as the valuations diverge to infinity. For procurement auctions with two players, we derive expressions for the optimal penalty multipliers (also presented in Section 6.4).

For the general case, i.e., more than two players, we give in Section 6.5 a complete characterization of the set of limit equilibria of a given procurement auction, and we show that limit equilibria always exist. We moreover show that the set of limit equilibria of an auction can be compactly represented as a polynomial number of quadruples of numbers. This result implies that one can compute a limit equilibrium (and therefore an ϵ-equilibrium, for any $\epsilon \in \mathbb{R}_{>0}$) in polynomial time. Lastly, we provide precise expressions for the price of spite of an auction, which are computable in polynomial time.

Finally, in Section 6.6 we discuss some potentially interesting future research directions.

6.3 Preliminaries

We model spiteful behavior by associating a spite level $\alpha_i \in (-1, 0]$ with each player i, and assuming that player i's perceived utility is a combination of $1 + \alpha_i$ times its direct utility and α_i times the utility of the other players. Tuning the parameters α_i allows smooth interpolation between purely spiteful and purely selfish behavior. Note that this model of spite resembles the altruism model of Chapter 4, but is not entirely the spiteful analogue of it. This model can however be seen as a special case of the model in Chapter 5 when negative altruism is allowed.

For generalized procurement auctions, we therefore define their *spiteful extension* as follows.

Definition 129 (Spiteful extension of first price and second price generalized procurement auctions)**.** Let $\Gamma = (n, v, \mu, \preceq)$ be a generalized procurement auction, and let $\alpha \in (-1, 0]^n$. The α-spiteful extension of Γ is the game $\Gamma^\alpha = (n, \Sigma, u^\alpha)$ where $\Sigma_i = \mathbb{R}_{\geq 0}$ for all $i \in [n]$, and

$$u_i^\alpha(s) = \begin{cases} (1 + \alpha_i)(s_i - v_i) & \text{if } i = w(s) \\ \alpha_i(s_{w(s)} - v_{w(s)}) & \text{otherwise .} \end{cases} \tag{6.1}$$

for all $s \in \Sigma, i \in [n]$, in case Γ is a first price generalized procurement auction, and

$$u_i^\alpha(s) = \begin{cases} (1 + \alpha_i)\left(\frac{\mu_{r(s)} s_{r(s)}}{\mu_i} - v_i\right) & \text{if } i = w(s), \\ \alpha_i\left(\frac{\mu_{r(s)} s_{r(s)}}{\mu_{w(s)}} - v_{w(s)}\right) & \text{otherwise .} \end{cases}$$

for all $s \in \Sigma, i \in [n]$, in case Γ is a second price generalized procurement auction.

Observe that when $\alpha = \mathbf{0}$, then $\Gamma^\alpha = \Gamma$ for all generalized procurement auctions Γ. In contrast, when $a_i \to -1$ for all $i \in [n]$, players are not interested any more in their own selfish utilities, but only in preventing positive selfish utility of others. We note that the spite model we use here is a special case of the one used by Brandt et al. [2007]. A second strongly related model of spite is the one used by Morgan et al. [2003].

Throughout the chapter, we use the word *auction* as a shorthand for *generalized procurement auction*. Our goal is to understand the loss of social welfare caused by spiteful behavior of the player. To define this loss formally, we first need to define the outcomes of auctions. While *pure equilibria* would be natural candidates for outcomes, they generally do not exist for the auctions that we study. Instead, we consider ϵ-*equilibria*. We are particularly interested in the ϵ-equilibria when $\epsilon \to 0$, which would suggest studying the limit set of ϵ-equilibria as $\epsilon \to 0$. However, it is easy to see that this defines exactly the set of pure equilibria again, which may be empty. Instead, we define the class of *limit equilibria* as follows:

Definition 130 (Limit Equilibria)**.** The set of *limit equilibria* of a game Γ is defined as $LE = \bigcap_{\epsilon \in \mathbb{R}_{>0}} \mathrm{Cl}(\epsilon\text{-}PE_\Gamma)$, where $\mathrm{Cl}(S)$ denotes the closure of set S, i.e., the union of S and its limit points.

Remark 131. Standard analysis arguments show that LE is always non-empty. It is moreover clear that whenever pure equilibria exist, they are also limit equilibria. Limit equilibria are further motivated by the following two observations. First, for each α-spiteful extension of an auction $\Gamma = (n, v, \mu, \preceq)$, for all $\alpha \in (-1, 0]^n$, the limit equilibria are precisely the bid vectors $s \in \Sigma$ such that there exists a tie breaking rule for the auction under which s is in fact a pure equilibrium.[3] Secondly, when $s \in \Sigma$ is a limit equilibrium, it is not necessarily a pure equilibrium, but it is still possible to obtain an ϵ-equilibrium for arbitrarily small $\epsilon > 0$, by perturbing s appropriately.

We measure the social welfare of a strategy profile of an auction by the sum of utilities of the players of the auction under that strategy profile, where (as is usual in the study of auctions) the auctioneer is considered a single-strategy player of the auction, for which its utility is the negated amount paid to the winning player.

Definition 132 (Social welfare function for spiteful procurement auctions)**.** Let $\alpha \in (-1, 0]^n$ and let Γ^α be the α-spiteful extension of an auction Γ. For $s \in \Sigma$, the *social welfare* $U_{\Gamma^\alpha}(s)$ of s in the game Γ^α is given by $-s_{w(s)} + \sum_{i \in [n]} u_i^\alpha(s)$ for first-price

[3]Notice the reversal of quantifiers which implies that this characterization does not contradict the fact that pure equilibria generally do not exist.

auctions, and by $-s_{r(s)} + \sum_{i\in[n]} u_i^\alpha(s)$. The subscript Γ^α in U_{Γ^α} will usually be dropped, in case no confusion can arise.

For auctions without spite, under this definition, the social welfare is equal to the negated valuation of the winner of the auction. Observe that this approach differs from what we did in Chapters 4 and 5: here, the spite levels are taken into account in our social welfare function, whereas in Chapters 4 and 5, the social welfare function considered is that of the original (spiteless/altruismless) game.

The social welfare function of this chapter thus measures the overall satisfaction of the players, in contrast to Chapters 4 and 5, where the purpose of the social welfare function was to provide an objective measure for the quality of each strategy profile *independent* of spite and altruism.

As we are interested in quantifying the loss in social welfare of the limit equilibria *as a consequence of spite between players*, for the case of limit equilibria we alter slightly the definitions of the price of anarchy and price of stability that we gave in Section 1.3.1.8, into the following:

Definition 133 (Worst-case and best-case price of spite)**.** Let $\alpha \in (-1,0]^n$, let LE be the set of limit equilibria of a given α-spiteful extension Γ^α of a first price or second price auction $\Gamma = (n, v, \mu, \preceq)$. Then, the *worst-case price of spite of* Γ^α is defined as

$$\sup\left\{\frac{U_{\Gamma^\alpha}(s)}{\sup\{U_\Gamma(s') \mid s' \in \Sigma\}} \;\middle|\; s \in LE\right\} = \sup\left\{\frac{U_{\Gamma^\alpha}(s)}{-\min\{v_i \mid i \in [n]\}} \;\middle|\; s \in LE\right\}.$$

and the *best-case price of spite of* Γ^α is defined as

$$\inf\left\{\frac{U_{\Gamma^\alpha}(s)}{\sup\{U_\Gamma(s') \mid s' \in \Sigma\}} \;\middle|\; s \in LE\right\} = \inf\left\{\frac{U_{\Gamma^\alpha}(s)}{-\min\{v_i \mid i \in [n]\}} \;\middle|\; s \in LE\right\}.$$

Thus, the price of spite measures the loss in social welfare of the auction compared to the optimal social welfare that can be attained when the players have no spite.

Our goal is to design auctions minimizing the worst-case and best-case price of spite. Our approach is to set the penalty multipliers μ_i so as to minimize these quantities, which can be considered a mechanism design problem in a limited design space. Our approach shares some similarities with the work of Archer and Tardos [2002], Karlin et al. [2005].

6.4 The Two-Player Case

In this section we study first price and second price auctions with two spiteful players. We begin by deriving necessary and sufficient conditions for a bid vector to be an ϵ-

equilibria. This allows us to precisely characterize limit equilibria and optimize the penalty multipliers.

For $(\alpha_1, \alpha_2) \in (-1, 0]^2$ and $\epsilon \in \mathbb{R}_{>0}$, we observe that a vector $(s_1, s_2) \in \Sigma$ is an ϵ-equilibrium of an α-spiteful extension of an auction Γ, if and only if it satisfies the following three conditions, where we assume for notational convenience and without loss of generality that player 1 is the winner under (s_1, s_2):

1. Player 1 cannot improve its utility by more than ϵ by changing its bid into something that is lower than $\mu_2 s_2/\mu_1$ (in which case player 1 would still win). In case Γ is a second price auction, this condition is true by definition. In case Γ is a first price auction, it is equivalent to $s_1 \geq \mu_2 s_2/\mu_1 - \epsilon$.

2. Player 1 cannot improve its utility by more than ϵ by changing its bid into something that is higher than $\mu_2 s_2/\mu_1$ (and thus losing to player 2). This condition is equivalent to $s_1 - v_1 \geq \alpha_1(s_2 - v_2) - \epsilon$ in case Γ is a first price auction, and $\mu_2 s_2/\mu_1 - v_1 \geq \alpha_1(s_2 - v_2) - \epsilon$ in case Γ is a second price auction.

3. Player 2 cannot improve by more than ϵ by changing its bid into something that is lower than $\mu_1 s_1/\mu_2$ (and thus becoming the winner of the auction). This condition is equivalent to $\mu_1 s_1/\mu_2 - v_2 \leq \alpha_2(s_1 - v_1) + \epsilon$ in case Γ is a first price auction and $\mu_1 s_1/\mu_2 - v_2 \leq \alpha_2(\mu_2 s_2/\mu_1 - v_1) + \epsilon$ in case Γ is a second price auctions.

Observe that in the above conditions, μ_1 and μ_2 occur only as part of the ratios μ_1/μ_2 and μ_2/μ_1. Thus, for ease of notation, we will henceforth assume that $\mu_1 = 1$, and write $\mu = \mu_2$.

For convenience, we define two constants, $\widetilde{s}_{1,2}$ and $\widetilde{s}_{2,1}$, that serve as "threshold" values in the results that follow:

$$\widetilde{s}_{1,2} = \frac{(1+\alpha_1)v_1 - \alpha_1 v_2}{(1+\alpha_1)\mu - \alpha_1}, \qquad \widetilde{s}_{2,1} = \frac{(1+\alpha_2)v_2 - \alpha_2 v_1}{(1+\alpha_2)/\mu - \alpha_2}.$$

For two spiteful players, for any $\epsilon \in \mathbb{R}_{>0}$ we first establish sufficient conditions for a bid vector (s_1, s_2) to be an ϵ-equilibrium.

Theorem 134 (ϵ-equilibria, sufficient conditions)**.** *Let $\epsilon \in \mathbb{R}_{>0}$, let $\alpha \in (-1, 0]^n$ and let Γ^α be a spiteful extension of an auction $\Gamma = (n, v, \mu, \preceq)$.*

1. *If $\mu\widetilde{s}_{1,2} \leq \widetilde{s}_{2,1}$, then define $\epsilon' = \epsilon/(1+\alpha_1)$ if Γ is a first price auction, and $\epsilon' = \epsilon/(-\alpha_2)$ if Γ is a second price auction.*

   A bid vector $(s_1, s_2) \in \Sigma$ is an ϵ-equilibrium of Γ^α if
   $s_1 \in [\mu s_2 - \epsilon', \mu s_2) \cap (-\infty, \widetilde{s}_{2,1}]$ and $\mu s_2 \in [\mu\widetilde{s}_{1,2}, \infty)$ (so player 1 wins).

2. *If* $\widetilde{s}_{2,1} \leq \mu\widetilde{s}_{1,2}$*, then define* $\epsilon' = \epsilon\mu/(1+\alpha_2)$ *if* Γ *is a first price auction, and* $\epsilon' = \epsilon\mu/(-\alpha_1)$ *if* Γ *is a second price auction.*

   A bid vector $(s_1, s_2) \in \Sigma$ *is an* ϵ*-equilibrium of* Γ^α *if* $s_1 \in (\mu s_2, \mu s_2 + \epsilon'] \cap [\widetilde{s}_{2,1}, \infty)$ *and* $\mu s_2 \in (-\infty, \mu\widetilde{s}_{1,2}]$ *(so player 2 wins).*

Proof. The two cases are symmetric, so we prove only the first one. Let $(s_1, s_2) \in \Sigma$ be a bid vector satisfying $s_1 \in [\mu s_2 - \epsilon', \mu s_2) \cap (-\infty, \widetilde{s}_{2,1}]$ and $\mu s_2 \in [\mu\widetilde{s}_{1,2}, \infty)$. We will prove that it is an ϵ-equilibrium. Consider the effects for player 1 of changing its bid to some other value s_1'. If $s_1' < s_1$, then for the second price auction, the utilities of the players do not change. For the first price auction, player 1 still wins, but obtains smaller utility. Thus, player 1 cannot improve by lowering its bid. Raising its bid to any value $s_1' \leq \mu s_2$ will not change the utilities in a second price auction, and will increase the utility in a first price auction by at most $(1+\alpha_1)\epsilon' = \epsilon$, because $s_1 \geq \mu s_2 - \epsilon'$ by assumption.

If $s_1' > \mu s_2$, then player 1 would cease to win. If Γ is a second price auction, the new utility is $\alpha_1(s_1'/\mu - v_2) \leq \alpha_1(s_2 - v_2)$, whereas it was previously $(1+\alpha_1)(\mu s_2 - v_1)$. If Γ is a first price auction, its new utility would be $\alpha_1(s_2 - v_2)$, whereas previously it was $(1+\alpha_1)(s_1 - v_1) \geq (1+\alpha_1)(\mu s_2 - \epsilon' - v_1) = (1+\alpha_1)(\mu s_2 - v_1) - \epsilon$. In both cases, the change is thus at most

$$-s_2(\mu + \alpha_1\mu - \alpha_1) + (1+\alpha_1)v_1 - \alpha_1 v_2 + \epsilon = (\mu + \alpha_1\mu - \alpha_1)(-s_2 + \mu\widetilde{s}_{1,2}) + \epsilon.$$

The first factor is non-negative because $\alpha_1 \geq -1$, and the second factor is non-positive by assumption on the bids. Thus, player 1 cannot improve its utility by more than ϵ by changing its bid.

Now consider player 2, changing its bid to some other value s_2'. If $s_2' > s_2$, then player 1 still wins, and obtains at least the same utility. If $s_2' \geq s_1/\mu$, then $s_2 - s_2' \leq \epsilon'/\mu$ by assumption, so the payment to player 1 decreases by at most $\mu\epsilon'/\mu = \epsilon'$, and the utility of player 2 thus increases by at most $-\alpha_2\epsilon' = \epsilon$.

If $s_2' < s_1/\mu$, then player 2 would win the auction. If Γ is a second price auction, its new utility is $(1+\alpha_2)(s_1/\mu - v_2)$, whereas it was previously $\alpha_2(\mu s_2 - v_1) \geq \alpha_2(s_1 + \epsilon' - v_1) = \alpha_2(s_1 - v_1) - \epsilon$. If Γ is a first price auction, its new utility would be $(1+\alpha_2)(s_2' - v_2) \leq (1+\alpha_2)(s_1/\mu - v_2)$, whereas previously, it was $\alpha_2(s_1 - v_1)$ In both cases, the change is thus at most

$$\frac{s_1}{\mu}(1 + \alpha_2 - \mu\alpha_2) - ((1+\alpha_2)v_2 - \alpha_2 v_1) + \epsilon = (1 + \alpha_2 - \mu\alpha_2)\Big(\frac{s_1}{\mu} - \frac{\widetilde{s}_{2,1}}{\mu}\Big) + \epsilon.$$

The first factor is non-negative because $\alpha_2 \geq -1$, and the second factor is non-positive by assumption on the bids. Thus, player 2 also cannot improve its utility by more than ϵ by changing its bid. So the bids are an ϵ-equilibrium. $\square$

From the theorem above, existence of ϵ-equilibria follows immediately, for all $\epsilon \in \mathbb{R}_{>0}$.

Likewise, we are able to derive a set of necessary conditions for a bid vector to be an ϵ-equilibrium, when $\epsilon \in \mathbb{R}_{>0}$ is small enough. These conditions are slightly weaker, but in the limit, i.e., when $\epsilon \to 0$, the sets coincide.

Theorem 135 (ϵ-equilibria necessary conditions)**.** *Let $\alpha \in (-1, 0]^n$ and let Γ^α be a spiteful extension of an auction $\Gamma = (n, v, \mu, \preceq)$. For a sufficiently small choice of $\epsilon \in \mathbb{R}_{>0}$, the following holds.*

1. *If $\mu\widetilde{s}_{1,2} \leq \widetilde{s}_{2,1}$, then let $\epsilon' = \epsilon/(1+\alpha_1)$ if Γ is a first price auction, and let $\epsilon' = \epsilon/(-\alpha_2)$ if Γ is a second price auction. Further, define $\epsilon^- = \epsilon/(1+\alpha_1-\alpha_1/\mu)$ and $\epsilon^+ = \epsilon/(1/\mu+\alpha_2/\mu-\alpha_2)$.*

   If (s_1, s_2) is an ϵ-equilibrium of Γ^α, then $s_1 \in [\mu s_2 - \epsilon', \mu s_2] \cap (-\infty, \widetilde{s}_{2,1} + \epsilon^+]$ and $\mu s_2 \in [\mu\widetilde{s}_{1,2} - \epsilon^-, \infty)$.

2. *If $\widetilde{s}_{2,1} \leq \mu\widetilde{s}_{1,2}$, then let $\epsilon' = \epsilon\mu/(1+\alpha_2)$ if Γ is a first price auction, and let $\epsilon' = \epsilon\mu/(-\alpha_1)$ if Γ is a second price auction. Further, define $\epsilon^- = \epsilon/(1/\mu+\alpha_2/\mu-\alpha_2)$ and $\epsilon^+ = \epsilon/(1+\alpha_1-\alpha_1/\mu)$.*

   If (s_1, s_2) is an ϵ-equilibrium of Γ^α, then $s_1 \in [\mu s_2, \mu s_2 + \epsilon'] \cap [\widetilde{s}_{2,1} - \epsilon^-, \infty)$ and $\mu s_2 \in (-\infty, \mu\widetilde{s}_{1,2} + \epsilon^+]$.

Proof. Because the cases are symmetric, we will only prove the first case. Assume $\mu\widetilde{s}_{1,2} \leq \widetilde{s}_{2,1}$, and let $(s_1, s_2) \in \Sigma$ be an ϵ-equilibrium of Γ^α. We first show that player 1 is the winner, both in case Γ is a first price auction and in case Γ is a second price auction. To derive a contradiction, assume that $s_1 > \mu s_2$, i.e., that player 2 wins the auction.

If Γ is a first price auction, and player 1 changes its bid to below μs_2, then its utility changes from $\alpha_1(s_2 - v_2)$ to at most $(1+\alpha_1)(\mu s_2 - v_1)$. The change in utility is thus at most

$$(1+\alpha_1)(\mu s_2 - v_1) - \alpha_1(s_2 - v_2) = \left(1 + \alpha_1 - \frac{\alpha_1}{\mu}\right)(\mu s_2 - \mu\widetilde{s}_{1,2}).$$

For ϵ small enough ($\epsilon \leq \widetilde{s}_{2,1} - \mu\widetilde{s}_{1,2}/2(1+\alpha_1-\alpha_1/\mu)$), we then get that $\widetilde{s}_{2,1} - \mu s_2 \geq (\widetilde{s}_{2,1} - \mu\widetilde{s}_{1,2})/2$. However, if player 2 changes its bid to something more than s_1/μ, its utility changes by

$$\alpha_2(s_1 - v_1) - (1+\alpha_2)(s_2 - v_2) = \frac{1}{\mu}(1 + \alpha_2 - \mu\alpha_2)(\widetilde{s}_{2,1} - \mu s_2) - \alpha_2(\mu s_2 - s_1)$$

$$= \frac{1}{\mu}(1+\alpha_2)(\widetilde{s}_{2,1} - \mu s_2) - \alpha_2(\widetilde{s}_{2,1} - s_1).$$

Therefore, for $s_1 \le \widetilde{s}_{2,1}$, and ϵ small enough ($\epsilon \le ((1+\alpha_2)/\mu)(\widetilde{s}_{2,1} - \mu\widetilde{s}_{1,2})/2$), the increase in utility exceeds ϵ. Finally, if $s_1 > \widetilde{s}_{2,1}$, then player 2 could change its bid to $\widetilde{s}_{2,1}/\mu$. Its utility then changes by $((1+\alpha_2)/\mu)(\widetilde{s}_{2,1} - \mu s_2)$, which is again greater than ϵ, for ϵ sufficiently small. This is a contradiction, so we conclude that player 1 is the winner.

If Γ is a second price auction, and player 1 changes its bid to just below μs_2, its utility changes from $\alpha_1(s_1/\mu - v_2) \le \alpha_1(s_2 - v_2)$ to $(1+\alpha_1)(\mu s_2 - v_1)$. As before, in an ϵ-equilibrium, with ϵ sufficiently small, this implies that $\widetilde{s}_{2,1} - \mu s_2 \ge (\widetilde{s}_{2,1} - \mu\widetilde{s}_{1,2})/2$. If player 1 changes its bid to just above μs_2, its utility changes by $-\alpha_1(s_1/\mu - s_2)$. Unless $s_1/\mu - s_2$ is small (say, $s_1/\mu - s_2 \le (\widetilde{s}_{2,1} - \mu\widetilde{s}_{1,2})/(4\mu)$), for ϵ small enough, this quantity will be larger than ϵ. Thus, it must be that $\widetilde{s}_{2,1} - s_1 \ge (\widetilde{s}_{2,1} - \mu\widetilde{s}_{1,2})/4$. Now, if player 2 changes its bid to just above s_1/μ, its utility changes by

$$\alpha_2(s_1 - v_1) - (1+\alpha_2)\left(\frac{s_1}{\mu} - v_2\right) = \frac{1}{\mu}(1+\alpha_2 - \mu\alpha_2)(\widetilde{s}_{2,1} - s_1),$$

which will be at least ϵ when ϵ is sufficiently small. A contradiction. Thus, we have shown that for both first price and second price auctions, at ϵ-equilibria (for $\epsilon \in \mathbb{R}_{>0}$ small enough), we have $s_1 \le \mu s_2$ (so player 1 wins).

Next, assume that $\delta = \mu s_2 - s_1 > \epsilon'$. In case Γ is a first price auction, player 1 could increase its bid by δ, and its utility by $(1+\alpha_1)\delta > (1+\alpha_1)\epsilon' = \epsilon$. In case Γ is a second price auction, player 2 could decrease its bid by δ/μ, thus increasing its utility by $-\alpha_2\delta > -\alpha_2\epsilon' = \epsilon$. Thus, we have proved that $s_1 \ge \mu s_2 - \epsilon'$.

Finally, if $\mu s_2 < \mu\widetilde{s}_{1,2} - \epsilon^-$, then by changing its bid to just above μs_2, player 1 could ensure that player 2 wins, and obtain utility $\alpha_1(s_2 - v_2)$, as opposed to the previous $(1+\alpha_1)(s_1 - v_1)$ when Γ is a first price auction, or $(1+\alpha_1)(\mu s_2 - v_1)$ when Γ is a second price auction. A similar calculation to the ones above (using for the first price auction that $s_2 \ge s_1/\mu$) then shows that the change in utility of player 1 is at least

$$\begin{aligned}(1+\alpha_1 - \alpha_1/\mu)(\mu\widetilde{s}_{1,2} - \mu s_2) \quad &\ge \quad (1+\alpha_1 - \alpha_1/\mu)\epsilon^- \\ &= \quad \epsilon,\end{aligned}$$

contradicting that the bids were at ϵ-equilibrium.

Similarly, if $s_1 > \widetilde{s}_{2,1} + \epsilon^+$, then by changing its bid to just below s_1/μ, player 2 could ensure to win, and obtain utility $(1+\alpha_2)(s_1/\mu - v_2)$, instead of $\alpha_2(s_1 - v_1)$ (when Γ is a first price auction) or $\alpha_2(s_2\mu - v_1)$ (when Γ is a second price auction). Again, a similar calculation (using $s_1 \le \mu s_2$ for the second price auction) gives that the change in utility of player 2 is at least

$$\begin{aligned}\left(\frac{1}{\mu}+\frac{\alpha_2}{\mu}-\alpha_2\right)(s_1-\mu\widetilde{s}_{1,2}) &\geq \left(\frac{1}{\mu}+\frac{\alpha_2}{\mu}-\alpha_2\right)\epsilon^+\\ &= \epsilon,\end{aligned}$$

again contradicting the assumption that (s_1, s_2) is an ϵ-equilibrium of Γ^α. This completes the proof. $\square$

Consider an auction $\Gamma = (2, v, \mu, \preceq)$ of two players and spite vector $\alpha \in (-1, 0]^2$. For any $\epsilon \in \mathbb{R}_{>0}$, let A_ϵ be the set of bid vectors satisfying the sufficient conditions of Theorem 134, and let B_ϵ be the set of bid vectors satisfying the necessary conditions of Theorem 135. Because $A_\epsilon \subseteq \epsilon\text{-}PE \subseteq B_\epsilon$, by Theorems 134 and 135, we also have $\mathrm{Cl}(A_\epsilon) \subseteq \mathrm{Cl}(N_\epsilon) \subseteq \mathrm{Cl}(B_\epsilon)$. Thus, $\bigcap_{\epsilon\in\mathbb{R}_{>0}} \mathrm{Cl}(A_\epsilon) \subseteq \mathcal{L}(\mathbf{c}) \subseteq \bigcap_{\epsilon>0} \mathrm{Cl}(B_\epsilon)$. But

$$\begin{aligned}\bigcap_{\epsilon\in\mathbb{R}_{>0}} \mathrm{Cl}(A_\epsilon) &= \bigcap_{\epsilon\in\mathbb{R}_{>0}} \mathrm{Cl}(B_\epsilon)\\ &= \{(s_1, s_2) \mid s_1 = \mu s_2 \in [\min\{\widetilde{s}_{2,1}, \mu\widetilde{s}_{1,2}\}, \max\{\widetilde{s}_{2,1}, \mu\widetilde{s}_{1,2}\}]\},\end{aligned}$$

both for first price and second price auctions. Thus, we have derived the following corollary:

Corollary 136. *Let $\alpha \in (-1, 0]^2$ and let $\Gamma = (2, v, \mu, \preceq]$ be a first price or second price auction. Then it holds that the limit equilibria of the α-spiteful extension of Γ are $\{(s_1, s_2) \mid s_1 = \mu s_2 \in [\min\{\widetilde{s}_{2,1}, \mu\widetilde{s}_{1,2}\}, \max\{\widetilde{s}_{2,1}, \mu\widetilde{s}_{1,2}\}]\}$.*

Conveniently, the fact that the limit equilibria are the same for both first price and second price auctions allows us to treat them jointly from now on (as we see that at a limit equilibrium, the social welfare of the first price auction is equal to the social welfare of the second price auction). We use Corollary 136 to bound the worst-case price of spite and best-case price of spite of both auction formats.

Theorem 137. *Let $\alpha \in (-1, 0]^2$ and let $\Gamma = (2, v, \mu, \preceq)$ be a first price or second price auction.*

1. *If $\mu\widetilde{s}_{1,2} \leq \widetilde{s}_{2,1}$, then the best-case price of spite of Γ^α is*

$$\frac{1}{\min\{v_1, v_2\}}\left(v_1 + \frac{\alpha_1(\alpha_1+\alpha_2)(\mu v_2 - v_1)}{\mu(1+\alpha_1)-\alpha_1}\right),$$

and the worst-case price of spite or Γ^α is

$$\frac{1}{\min\{v_1, v_2\}}\left(v_1 - \frac{(\alpha_1+\alpha_2)(1+\alpha_2)(\mu v_2 - v_1)}{(1+\alpha_2)-\mu\alpha_2}\right).$$

2. *If* $\widetilde{s}_{2,1} \leq \mu\widetilde{s}_{1,2}$*, then the best-case price of spite of* Γ^α *is*

$$\frac{1}{\min\{v_1, v_2\}}\left(v_2 + \frac{(\alpha_1+\alpha_2)\alpha_2(v_1-\mu v_2)}{(1+\alpha_2)-\mu\alpha_2}\right),$$

and the worst-case price of spite or Γ^α *is*

$$\frac{1}{\min\{v_1, v_2\}}\left(v_2 - \frac{(\alpha_1+\alpha_2)(1+\alpha_1)(v_1-\mu v_2)}{\mu(1+\alpha_1)-\alpha_1}\right).$$

Proof. We prove the first claim for first price auctions. Similar reasoning applies for the other three cases.

The social welfare of an outcome is the sum of all utilities, including the auctioneer's. The optimal social welfare is calculated for players without spite. As argued in Section 5.3, without spite, the optimal social welfare is $-\min\{v_1, v_2\}$.

All limit equilibria satisfy $s_1 = \mu s_2$. The auctioneer's utility is thus $-s_1$. The utility of player 1 (the winner) is $(1+\alpha_1)(s_1 - v_1)$, while player 2's utility is $\alpha_2(s_1 - v_1)$. Thus, the total utility is $-v_1 + (\alpha_1 + \alpha_2)(s_1 - v_1)$. Because $\alpha_1 + \alpha_2 \leq 0$, the welfare is maximized when s_1 is as small as possible, i.e., when $s_1 = \mu\widetilde{s}_{1,2}$. The welfare is minimized when s_1 is as large as possible, i.e., $s_1 = \widetilde{s}_{2,1}$. Substituting these two bids and simplifying now gives the claimed bound. □

Our next goal is to use Theorem 137 to choose the value of μ minimizing the worst-case price of spite or best-case price of spite. If both the spite vector α and valuation vector v are known, this question is neither interesting nor meaningful: a simple calculation shows that the auctioneer sets μ to either 0 or ∞, thus predetermining the winner. We therefore consider the question of how to optimize the price of spite using only knowledge of the spite levels α, so as to minimize the best-case and worst-case price of spite over all valuation vectors.

An immediate problem with this goal is that as $v_2/v_1 \to \infty$ or $v_1/v_2 \to \infty$, both the best-case and worst-case price of spite will diverge to infinity regardless of the choice of μ. This is not surprising: in the presence of spite, the losing player's utility will become more negative as the winning player's valuation increases. Therefore, we state our goal more precisely as minimizing the rate of divergence of the worst-case price of spite and best-case price of spite as a function of $\max(v_1/v_2, v_2/v_1)$. This motivates the following definitions.

Definition 138. Let $\rho \in \mathbb{R}_{>0}$. We define the set $\mathcal{C}_\rho$ as

$$\left\{(v_1, v_2) \in \mathbb{R}^2_{\geq 0} \;\middle|\; \max\left\{\frac{v_1}{v_2}, \frac{v_2}{v_1}\right\} \leq \rho\right\}.$$

Let $\alpha \in (-1,0]^2$. For $\mu \in \mathbb{R}^2_{\geq 0}$ and $\rho \in \mathbb{R}_{>0}$, denote by $BPoS_\alpha(\rho,\mu)$ and $WPoS_\alpha(\rho,\mu)$ the maximum best-case and worst-case price of spite among the α-spiteful extensions of the auctions $\{(2,v,\mu,\preceq) \mid v \in \mathcal{C}_\rho\}$. The penalty multipliers $\mu^B_\alpha(\rho)$ and $\mu^W_\alpha(\rho)$ are then defined as $\arg_\mu \min\{BPoS_\alpha(\rho,\mu) \mid \mu \in \mathbb{R}_{\geq 0}\}$ and $\arg_\mu \min\{WPoS_\alpha(\rho,\mu) \mid \mu \in \mathbb{R}_{\geq 0}\}$, respectively.

Using the above definition, we can characterize the optimal choice of μ, as a function of the spite levels α, as follows.

Theorem 139. *Let $\alpha \in (-1,0]^2$. If $\alpha_1 + \alpha_2 \geq -1$,*

$$\lim_{\rho\to\infty} \mu^B_\alpha(\rho) = \frac{\alpha_1 - \alpha_2 + \sqrt{(\alpha_1-\alpha_2)^2 + 4\alpha_1^2\alpha_2^2}}{2\alpha_1\alpha_2},$$

and

$$\lim_{\rho\to\infty} \mu^W_\alpha(\rho) = \frac{\alpha_1 - \alpha_2 + \sqrt{(\alpha_1-\alpha_2)^2 + 4(1-\alpha_2)^2(1-\alpha_1)^2}}{2(1-\alpha_1)(1-\alpha_2)}.$$

If $\alpha_1 + \alpha_2 < -1$,

$$\lim_{\rho\to\infty} \mu^B_\alpha(\rho) = \sqrt{\frac{\alpha_2(1+\alpha_1)}{\alpha_1(1+\alpha_2)}},$$

and

$$\lim_{\rho\to\infty} \mu^W_\alpha(\rho) = \sqrt{\frac{\alpha_2(1+\alpha_1)}{\alpha_1(1+\alpha_2)}}.$$

Proof. We will prove the claim for $\mu^B_\alpha(\rho)$ when $\alpha_1 + \alpha_2 \geq -1$. The other three cases are handled analogously.

Fix $\mu \in \mathbb{R}_{\geq 0}$. Let $\Gamma = (n,v,\mu,\preceq)$ be an auction such that $v_2/v_1 = \rho$, and ρ is sufficiently large, so that $v_1 \leq \mu v_2$.

By straightforward manipulations, the condition $\mu\widetilde{s}_{1,2} \leq \widetilde{s}_{2,1}$ is equivalent to $(1 + \alpha_1 + \alpha_2)(v_1 - \mu v_2) \leq 0$. Since we are in the case that $\alpha_1 + \alpha_2 \geq -1$, this means that $\mu\widetilde{s}_{1,2} \leq \widetilde{s}_{2,1}$ is equivalent to $v_1 \leq \mu v_2$ (so the player with smaller scaled valuation wins).

Therefore, $\mu\widetilde{s}_{1,2} \leq \widetilde{s}_{2,1}$ holds, and the first case of Theorem 137 applies. Hence, the best-case price of spite of the α-spiteful extension of Γ is $1 + \frac{\alpha_1(\alpha_1+\alpha_2)(\mu\rho-1)}{\mu+\mu\alpha_1-\alpha_1}$.

Suppose on the other hand that $\Gamma = (n,v,\mu,\preceq)$ is an auction such that $v_1/v_2 = \rho$, and ρ is sufficiently large, so that $v_1 > \mu v_2$. By reasoning similar to the above, we conclude that $\mu\widetilde{s}_{1,2} > \widetilde{s}_{2,1}$ holds, and the second case of Theorem 137 applies, and the best-case price of spite of the α-spiteful extension of Γ is $1 + \frac{\alpha_1(\alpha_1+\alpha_2)(\mu\rho-1)}{\mu+\mu\alpha_1-\alpha_1}$

This expression is decreasing in μ, while the expression in the case $v_2/v_1 = \rho$ is increasing in μ. Thus, the optimal value $\mu_\rho^B(\alpha)$ must be the value of μ where the two expressions are equal. Solving the resulting quadratic equation and letting $\rho \to \infty$ proves the claim. $\square$

A fairly straightforward calculation shows that in all cases, the optimum penalty multiplier $\mu_\rho^B(\alpha)$ for player 2 is increasing in $|\alpha_2|$ and decreasing in $|\alpha_1|$. Thus, players with higher spite levels are penalized by the auction in the sense that their payments are lower.

6.5 Auctions with n Bidders

In this section we extend some of our results to auctions with more than two players. The structure of the set of ϵ-equilibria in an auction with n players, $n > 2$, is in general more complex than in the two player case, and therefore we are not able to give reasonably matching sets of necessary and sufficient conditions, as we did for the two player case. However, it turns out that we are still able to show that limit equilibria exist for any α-spiteful extension of any first price and second price auction with any number n of players, and any spite vector $\alpha \in (-1, 0]^n$. We will even see that it is possible to compute the set of all limit equilibria in polynomial time.

It turns out that the number of limit equilibria in a first price and second price is always infinite. Nevertheless, we will see that the structure of the set of limit equilibria is reasonably simple: In an α-spiteful extension of an auction $\Gamma = (n, v, \mu, \preceq)$, with $\alpha \in (-1, 0]^n$, the set of limit equilibria is a union of at most n^2 sets, with the property that these sets are described by quadruples (w, r, a, b) that consist of two players $w, r \in [n]$ and two real numbers $a, b \in \mathbb{R}_{\geq 0}$, $a \leq b$. An equilibrium in such a set is one in which w and r both have the same scaled bid (i.e., $\mu_w s_w = \mu_r s_r$), r bids any value $s_r \in [a, b]$, and every remaining player i bids at least $(\mu_r/\mu_i)s_r$. We present a polynomial time algorithm that outputs all these quadruples.

In the exposition that follows, the values $\widetilde{s}_{i,j}$ (for every pair of players $i, j \in [n], i \neq j$) generalize the roles of $\widetilde{s}_{1,2}$ and $\widetilde{s}_{2,1}$ that we defined for the two player case:

$$\widetilde{s}_{i,j} = \frac{(1+\alpha_i)v_i - \alpha_i v_j}{(1+\alpha_i)(\mu_j/\mu_i) - \alpha_i}.$$

For the special case of two players, we observed that for a sufficiently small choice of $\epsilon > 0$, there is a unique winner among the set of ϵ-equilibria (except for the case that $\mu\widetilde{s}_{1,2} = \widetilde{s}_{2,1}$). This property does not necessarily hold anymore in three player auctions, as the following example shows.

Example 140. Consider the α-spiteful extension of the auction $(3, v, \mu, \prec)$ in which $\mu_1 = \mu_2 = \mu_3 = 1$, and $v_1 = 1.5, v_2 = 1, v_3 = 4$, and $\alpha_1 = -0.1, \alpha_2 = -0.6, \alpha_3 = -0.8$. It is not hard to verify that the values $\widetilde{s}_{i,j}$ are all distinct, and that the bid vectors $s = (s_1 = 1.3, s_2 = 1.3, s_3 = 2)$ and $s' = (s'_1 = 3, s'_2 = 1.6, s'_3 = 1.6)$ are limit equilibria. With a tie-breaking rule that prefers the higher-numbered players, the winner is player 2 in the former limit equilibrium (with $u_1^{\alpha}(s) = -0.03, u_2^{\alpha}(s) = 0.12, u_3^{\alpha}(s) = -0.24$), and the winner is player 3 in the latter equilibrium (with $u_1^{\alpha}(s') = -0.24, u_2^{\alpha}(s') = 1.44, u_3^{\alpha}(s') = -0.48$).

We now give our characterization of limit equilibria of auctions with n players. For two numbers $a, b \in \mathbb{R}$, we use the notation $[a, b]$ to denote the usual closed interval of reals if $a \leq b$. If $a > b$, then $[a, b]$ denotes the empty set.

Theorem 141. *Suppose that Γ^{α} is an α-altruistic extension of an auction $\Gamma = (n, v, \mu, \preceq)$, for $\alpha \in (-1, 0]^n$. The set of limit equilibria of Γ^{α} (irrespective of whether Γ is a first price or second price auction) is non-empty and is given by*

$$LE = \bigcup_{(w,r)\in[n]^2 : w \neq r} LE_{w,r}, \tag{6.2}$$

where $LE_{w,r}$ denotes the set of limit equilibria for player pair $(w, r) \in [n]^2, w \neq r$, with the property that in every limit equilibrium of $LE_{w,r}$ player w is a player with the lowest scaled bid, and r is a player with the lowest scaled bid among the set of players in $[n]\backslash\{w\}$. This set is given by:

$$LE_{w,s} = \Big\{(s_1, \ldots, s_n) \;\Big|\; \tag{6.3}$$

$$s_w = \frac{\mu_r}{\mu_w} s_r, \tag{6.4}$$

$$s_r \in \left[\widetilde{s}_{w,r}, \min_{i\in[n]\backslash\{w\}} \widetilde{s}_{i,w} \frac{\mu_w}{\mu_r}\right], \tag{6.5}$$

$$\forall i \in [n]\backslash\{w, r\} : s_i \geq \frac{\mu_r}{\mu_i} b_r \Big\}. \tag{6.6}$$

In words, (6.4) is the requirement that the scaled bid of w is equal to the second-lowest scaled bid. Equation (6.5) is the requirement that the scaled bid of r is in some specific interval that is determined by the values $\widetilde{s}_{i,j}$, where $(i, j) \in [n]^2, i \neq j$, which may be an empty interval (but is, according to the theorem, not empty for at least one of the sets $LE_{i,j}, (i, j) \in [n]^2$). Equation (6.6) is the requirement that the scaled bids of all remaining players are at least the scaled bid of player r.

Proof of Theorem 141. Fix ϵ, let $w, r \in [n]$ be two arbitrary players and suppose that $(s_1, \ldots, s_n)$ is an ϵ-equilibrium of Γ^α in which w wins (so w has the lowest scaled bid) and r is a player not equal to w, with a lowest scaled bid among the players in $[n] \setminus \{w\}$, so $\mu_w s_w \leq \mu_r s_r \leq \mu_i s_i$ for $i \in [n] \setminus \{w, r\}$. No player can improve its utility by more than ϵ by changing its strategy. This is equivalent to saying that the following constraints are satisfied: (i) w can not improve its utility by more than ϵ by changing its scaled bid to a bid of at least $\mu_r s_r$, thereby becoming the player with the second lowest scaled bid and letting r become the winner, (ii) no player $i \in [n] \setminus \{w\}$ can improve its utility by more than ϵ by setting its scaled bid to a bid if at least $\mu_w s_w$, (iii) no player $i \in [n] \setminus \{w\}$ can improve by more than ϵ by changing its scaled bid to just below $\mu_w s_w$.

These constraints can be formulated as a set of strict and non-strict inequalities both for first price and second price auctions (the strictness of such an inequality being determined by the tie-breaking rule), so the set of ϵ-equilibria (in which player w wins and player r is a player with the lowest scaled bid among players $[n] \setminus \{w\}$) is an intersection of a finite number of open and closed halfspaces. The closure of this set is then the polyhedron obtained by making all strict inequalities of this system non-strict. For second price auctions, the resulting inequalities are as follows:

$$\alpha_w(s_r - v_r) - (1 + \alpha_w)\Big(\frac{\mu_r s_r}{\mu_w} - v_w\Big) \leq \epsilon, \tag{6.7}$$

$$\forall i \in [n] \setminus \{w\} : \alpha_i(s_w - v_w) - \alpha_i\Big(\frac{\mu_r s_r}{\mu_w} - v_w\Big) \leq \epsilon, \tag{6.8}$$

$$\forall i \in [n] \setminus \{w\} : (1 + \alpha_i)\Big(\frac{\mu_w s_w}{\mu_i} - v_i\Big) - \alpha_i\Big(\frac{\mu_r s_r}{\mu_w} - v_w\Big) \leq \epsilon, \tag{6.9}$$

$$\forall i \in [n] \setminus \{w, r\} : \mu_r s_r \leq \mu_i s_i, \tag{6.10}$$

$$\mu_w s_w \leq \mu_r s_r. \tag{6.11}$$

By setting ϵ to 0, we obtain the polyhedron of limit equilibria in which w has the lowest scaled bid and r is a player with lowest scaled bid among the players $[n] \setminus \{w\}$. We show next that this polyhedron is equivalent to LE_{wr}, by rewriting the system (6.7–6.11) for $\epsilon = 0$ into the equivalent system (6.3–6.6): Divide (6.8) by α_i (which is a negative number), to obtain that $s_w \geq (\mu_r/\mu_w)s_r$, and combine this with (6.11) to obtain (6.4). Replace s_r with $(\mu_r/\mu_w)s_r$ in (6.9). Rewriting (6.9) and (6.7) into the constraints

$$s_r \geq \widetilde{s}_{w,r}, \qquad \forall i \in [n] \setminus \{w\} : s_r \leq \frac{\mu_w}{\mu_r}\widetilde{s}_{i,w}, \tag{6.12}$$

is then straightforward. It is now obvious that (6.12) is equivalent to (6.5). Finally, (6.10) corresponds to (6.6). We conclude that the set of limit equilibria of a second price auction is given by (6.2). For a first price auction, we arrive at exactly the same set of constraints when we use similar reasoning. We omit this here.

It remains to be shown that LE is not empty. Suppose for contradiction that LE is empty. Let $\ell : [n] \to [n]$ be any mapping such that, for $i \in [n]$, $\ell(i)$ is a player in the set $\arg_{i'} \min_{i' \in [n]} \widetilde{s}_{i',i}$. From (6.5) and the assumption that $LE = \varnothing$ we derive that for all pairs of players (i, i') it holds that

$$\widetilde{s}_{i,i'} > \widetilde{s}_{\ell(i),i}(\mu_i/\mu_{i'}). \tag{6.13}$$

Now pick an arbitrary player $i \in [n]$ and consider the procedure of iteratively generating the sequence of players $\ell(i), \ell^2(i), \ell^3(i), \ldots$, where $\ell^k(i), k \in \mathbb{N}$ denotes the repeated composition of ℓ with itself, k times, on i. Stop this procedure at the first iteration $k \in \mathbb{N}$ where a player $\ell^k(p)$ is encountered that has already been generated, at iteration $j \in \mathbb{N}, j < k$ say. So, $\ell^j(i) = \ell^k(i)$. As there is only a finite number of n players, this procedure certainly terminates. From (6.13) we obtain the following sequence of strict inequalities, giving rise to a contradiction:

$$\begin{aligned}
\widetilde{s}_{\ell^{j+1}(p),\ell^j(p)} &> \widetilde{s}_{\ell^{j+2}(p),\ell^{j+1}(p)} \frac{\mu_{\ell^{j+1}(p)}}{\mu_{\ell^j(p)}} \\
&> \widetilde{s}_{\ell^{j+3}(p),\ell^{j+2}(p)} \frac{\mu_{\ell^{j+2}(p)}}{\mu_{\ell^{j+1}(p)}} \frac{\mu_{\ell^{j+1}(p)}}{\mu_{\ell^j(p)}} \\
&> \cdots \\
&> \widetilde{s}_{\ell^{k+1}(p),\ell^k(p)} \frac{\mu_{\ell^k(p)} \cdots \mu_{\ell^{j+1}(p)}}{\mu_{\ell^{k-1}(p)} \cdots \mu_{\ell^j(p)}} \\
&= \widetilde{s}_{\ell^{j+1}(p),\ell^j(p)}.
\end{aligned}$$

So we conclude that our assumption that $LE = \varnothing$ was incorrect. □

An immediate corollary of Theorem 141 is that the number of limit equilibria is infinite: the set of limit equilibria is non-empty and a union of sets that are either empty or are of infinite cardinality.

By using expressions (6.2–6.6), we can straightforwardly output the set of limit equilibria in a compact and efficient way:

Corollary 142. *There is a polynomial time algorithm that takes as input a description of a first price or second price auction $\Gamma = (n, v, \mu, \preceq)$, and a spite vector $\alpha \in (-1, 0]^n$, and outputs the set of limit equilibria LE of the α-spiteful extension of Γ represented as a list of quadruples $((w_1, r_1, a_1, b_1), \ldots, (w_q, r_q, a_q, b_q)), q \le n^2$, where $a_i \le b_i$ for $i \in [q]$. This list represents LE in the sense that*

$$LE = \bigcup_{i \in [q]} \Big\{ (s_1, \ldots, s_n) \,\Big|$$

$$\frac{\mu_{w_i}}{\mu_{r_i}} s_{w_i} = s_{r_i}, s_{r_i} \in [a_i, b_i], (\forall j \in [n] \backslash \{w_i, r_i\}) \big(s_j \geq \frac{\mu_r}{\mu_j} s_r\big) \Big\}.$$

Theorem 141 immediately gives us expressions for the worst-case price of spite and best-case price of spite of both first price auctions and second price auctions.

Corollary 143. *Let* $\Gamma = (n, v, \mu, \preceq)$ *be a first price or second price auction, and let* $\alpha \in (-1, 0]^n$.

- *For any two players* $w, r \in [n], w \neq r$*, let*

$$B_{w,r} = \frac{1}{\min\{v_i \mid i \in [n]\}} \left(v_w + \frac{(v_r \mu_r / \mu_w - v_w) \alpha_w \sum_{i \in [n]} \alpha_i}{(1 + \alpha_w) \mu_r / \mu_w - \alpha_w} \right).$$

  The best-case price of spite of the α*-spiteful extension of* Γ *is*

$$\min \left\{ B_{w,r} \;\middle|\; (w, r) \in [n]^2, w \neq r, \widetilde{s}_{w,r} \leq \min_{i \in [n]} \frac{\widetilde{s}_{i,w} \mu_w}{\mu_r} \right\}.$$

- *For any two players* $w, t \in [n], w \neq t$*, let*

$$B_{w,t} = \frac{1}{\min\{v_i \mid i \in [n]\}} \left(v_w + \frac{(1 + \alpha_t)(v_w \mu_w / \mu_t - v_t) \sum_{j \in [n]} \alpha_j}{(1 + \alpha_t) \mu_w / \mu_t - \alpha_t} \right).$$

  The worst-case price of spite of the α*-spiteful extension of* Γ *is*

$$\begin{aligned} \max \Big\{ W_{w,t} \quad \Big| \quad & (w, t) \in [n]^2, w \neq t, \\ & \exists r \in [n] \backslash \{w\} : \widetilde{s}_{w,r} \leq \widetilde{s}_{t,w} \frac{\mu_w}{\mu_r}, \\ & \forall i \in [n] \backslash \{w\} : \widetilde{s}_{i,w} \geq \widetilde{s}_{t,w} \Big\}. \end{aligned}$$

The above corollary follows from Theorem 141 by applying similar reasoning as in proof of Theorem 137.

6.6 Future Work

The most pressing question that remains to be answered is probably that of how to set the penalty multipliers so as to minimize the price of spite in general n-player procurement auctions.

Another obvious research direction would be to do the same type of analysis in case there are altruistic players, or a mixed population of altruistic and spiteful players. Yet another interesting research direction would be to examine the same type of questions, but focus on the revenue of the auction instead of the social welfare.

Finally, it would also make sense to go beyond single-item auctions and consider spite or altruism for more complex auctions, e.g., path auctions or “hiring-a-team auctions”. Other promising auction games to analyze are keyword auctions [Liang and Qi, 2007], as well as other various combinatorial settings.

Chapter 7

Finding Social Optima in Congestion Games with Externalities*

We continue our study of behavior induced by externalities between the players. In this chapter we are concerned with the computational problem of finding a strategy profile that optimizes the social welfare, for a generalization of the class of symmetric singleton congestion games in which it is possible for the players to express their externalities in a refined way. As for the preliminary knowledge discussed in the introduction chapter, the reader is advised to be familiar with Section 1.3 up to Section 1.3.1.1, Section 1.3.1.7, Section 1.3.1.10, and Section 1.3.1.11.

The difference between regular congestion games and the congestion games considered in this chapter, is that in this latter class of symmetric singleton congestion games, players maximize their utility instead of minimize their costs. Moreover, every player expresses for each facility and each sufficiently small set of other players, a value that represents how much the player values being on the facility together with the set of players. Formally, our games of interest are defined as follows.

Definition 144 (Symmetric singleton congestion games with r-externalities)**.** A *symmetric singleton congestion game with r-externalities*, for $r \in \mathbb{N}_{>0}$, is a utility maximization game (n, Σ, u), for which there is an $m \in \mathbb{N}_{>0}$ such that $\Sigma_i = [m]$, and the utility functions u are defined as follows. For all $i \in [n]$, $j \in [m]$, and $P \subseteq [n] \setminus \{i\}$,

*A part of the contents of this chapter has been published as De Keijzer and Schäfer [2012].

$|P| \leq r$, there is a number $v_{i,P,j} \in \mathbb{Q}$, such that for every strategy profile $s \in \Sigma$ and player $i \in [n]$,

$$u_i(s) = \sum_{\substack{P \subseteq [n] \setminus \{i\} : |P| \leq r, \\ (\forall i' \in P : s_{i'} = s_i)}} v_{i,P,s_i}.$$

Therefore, a symmetric singleton congestion game with r-externalities can be represented as a triple (n, m, v), where v is the vector containing all values $v_{i,P,j}$. For the case that $r = 1$ we will abuse notation and write $v_{i,i',j}$ instead of $v_{i,\{i\},j}$ for $i, i' \in [n], j \in [m]$. Also, for notational convenience we will sometimes write $v_{i,i,j}$ to denote $v_{i,\varnothing,j}$ for $i \in [n], j \in [m]$.

The elements of $[m]$ are called the *facilities*. When $P \subseteq [n]$ the value $v_{i,P,j}$ is called the *externality* that player $i \in [n]$ has for player set P on facility $j \in [m]$. Note that $v_{i,\varnothing,j}$ represents the intrinsic utility that player $i \in [n]$ obtains from facility $j \in [m]$. When the values v are unrestricted, we speak of a congestion game with *mixed externalities*. When v are non-negative, we speak of *non-negative externalities*.

Our focus will initially be on symmetric singleton congestion games with 1-externalities, and most of our results deal with the case of non-negative externalities. We first adopt a purely computational perspective rather than a game-theoretic perspective, and we consider the problem of optimizing the social welfare. Our social welfare function of choice here is the regular sum-of-utilities function (1.4). In particular, we prove that finding the optimal strategy profile is NP-hard even for very special cases, and focus on approximation algorithms. We then derive, for $r = 1$ a 2-approximation algorithm that works by rounding an optimal solution of a natural LP formulation of the problem. This rounding procedure is nontrivial because it needs to take care of the dependencies between the players resulting from the pairwise externalities. We also show that this is essentially the best possible rounding algorithm (with respect to the approximation factor) by showing that the integrality gap of the LP is close to 2. Adaptations of our rounding algorithm enable us to derive approximation algorithms for several generalizations of the problem. Most notably, we obtain an $(r+1)$-approximation for the class of symmetric singleton congestion games with r-externalities. Further, for $r = 2$ we derive a 2-approximation for the *non-symmetric* case, i.e., where the strategy sets of the players are not necessarily all equal. Moreover, we obtain a $3/2$-approximation algorithm when these sets are of size 2, i.e., each player has access to only two facilities.

Remark 145. As all games considered in this chapter are symmetric singleton congestion games with r-externalities, we will for convenience abbreviate this to simply *congestion games with r-externalities*. Secondly, when we talk in this chapter about the *social welfare* of a strategy profile for a congestion game, we always refer to the sum-of-utilities social welfare (1.4).

The second problem we consider is an algorithmic mechanism design problem: we consider a setting where a central authority assigns the players to the facilities. The valuations and externalities of the players are considered private in this setting, and they have to be reported to the central authority by the players themselves. This may cause the players to misreport their valuations out of self-interest. We are interested in resolving the question of how to find (in polynomial time) an allocation of the players to the facilities, taking into account that the players may be misreporting their valuation.

7.1 Background

Ever since their introduction in 1973, congestion games have been the subject of intensive research in game theory and, more recently, in algorithmic game theory. Most of these studies adopt a distributed viewpoint and focus on issues like the existence and inefficiency of equilibria, and the computational complexity of finding such equilibria, etc. (see, e.g., Nisan et al. [2007] for an overview). Much less attention has been given to the study of congestion games from a centralized viewpoint.

Studying these congestion games from a centralized viewpoint is important in situations where a centralized authority has influence over the players in the game. Also, adopting a centralized perspective may help in acquiring insights about the decentralized setting: if it is hard to find an (approximate) optimum or near-optimum in the centralized case where all the players are completely coordinated, it certainly will be hard for the players to reach such a solution in the decentralized case, where besides lack of coordinated computation, additional issues related to selfishness and stability arise. Lastly, we believe that studying this optimization problem is interesting for its own sake, as it can be seen as a generalization of various fundamental optimization problems.

We know of only two related articles [Blumrosen and Dobzinski, 2006, Chakrabarty et al., 2005] that study player-specific congestion games from an optimization perspective. Both works assume that the players are *anonymous* [Blumrosen and Dobzinski, 2006] in the sense that the utility function of a player only depends on the number of players using the chosen facility, but not on the identities of these players.

The assumption that all players are anonymous is overly simplistic in many situations. Therefore, our congestion games variant extends the player-specific congestion game model of Milchtaich [1996] to incorporate *non-anonymous* players.

There are various other papers that study congestion games with negative or non-negative externalities. For example, negative externalities are studied in routing [Roughgarden, 2005], scheduling and load balancing [Awerbuch et al., 1995]. Non-negative externalities are studied in the context of cost sharing [Feigenbaum et al., 2001], facility location [Anshelevich et al., 2008] and negotiations [Conitzer and Sandholm, 2004, 2005].

Meyers and Schulz [2012] studied the complexity of finding a minimum cost solution in congestion games (according to Rosenthal's classical congestion game model [Rosenthal, 1973]). They study several variants of the problem and prove NP-hardness results, as well as inapproximability results for some cases and polynomial time computability results for some other cases.

Chakrabarty et al. [2005] were the first to study player-specific congestion games from a centralized optimization perspective. The authors study the cost-minimization variant of the problem where each player has a non-negative and non-decreasing cost function associated with each facility.[1] They show that computing an assignment of minimum total cost is NP-hard. The authors also derive some positive results for certain special cases of the problem (see Chakrabarty et al. [2005] for details).

Most related to the work discussed in this chapter is Blumrosen and Dobzinski [2006]. They study the problem of welfare maximization in player-specific congestion games with non-negative utility functions. Among other results, they give NP-hardness and inapproximability results for non-negative and negative externalities. They also provide a randomized 18-approximation algorithm for arbitrary (non-negative) utility functions.

For $r = 1$, the problem of computing a welfare maximizing allocation of players to facilities for our variant of congestion games can also be interpreted as the following graph coloring problem: We are given a complete undirected graph on vertex set $[n]$, and we are given a set of colors $[m]$. Every edge (i, i') (including self-loops) has a weight $w_{i,i',e} = v_{i,i',e} + v_{i',i,e}$ for each color $e \in [m]$. The goal is to assign a color to every node such that the total weight of all *monochromatic* edges, i.e., edges of which both endpoints have the same color, is maximized. The weight of a monochromatic edge (i, i') is defined as $w_{i,i',e}$, where e is the color of the endpoints. The minimization variant of this problem with identical weights $w_{i,i',e} = w_{i,i'}$ for all $e \in [m]$ and every edge (i, i') is also known as the generalized graph coloring problem [Kolen and Lenstra, 1995], graph k-partitioning [Kann et al., 1997], and k-min cluster [Sahni and Gonzalez, 1976].

Auctions with r-Restricted Complements

Strongly related to the class of congestion games we study here, and especially the mechanism design problem we consider, is the class of *auctions with r-restricted com-*

[1] Equivalently, the utility functions are assumed to be non-positive and non-increasing.

plements. These are games that have been studied in Abraham et al. [2012], in a mechanism design context.

In such auctions, there is a set of bidders $[m]$ and a set of items $[n]$ (we switch the roles of n and m here for reasons that will later become clear). Each player $i \in [m]$ has a non-negative rational valuation $w_{i,S} \in \mathbb{R}$ for every subset $S \subseteq [n]$ of at most r items. We denote the vector of all such valuations by w. The players report their valuations to an auction mechanism, and the auction mechanism subsequently decides which items to give to each bidder. There is only one copy of each item, and an item cannot be split: it is allocated to a player entirely or not at all. Secondly, a price is charged by the auction to each player. Suppose that a player $i \in [n]$ gets allocated by the auction the set of items $S_i \subseteq [n]$ and gets charged a price of p_i. The utility of player i is then defined as $\sum_{S \subseteq S_i} w_{i,S} - p_i$. In determining the social welfare of a given outcome of the auction (i.e., an allocation of the items to the players), we use the sum-of-utilities social welfare function (1.4), with the auctioneer considered as a single-strategy player for which its utility is the total amount of money paid by the other player (as is usual in auction games). Note that this means that the social welfare is therefore determined by the allocation of the items to the players, and not by the prices charged.

It can be seen that the problem of finding a social welfare maximizing strategy profile for congestion games with r-restricted complements, is equivalent to the problem of finding a social welfare maximizing allocation for congestion games with $(r-1)$-externalities: given a congestion game with r-externalities $\Gamma = (n, m, v)$, we can construct an auction Γ' with $(r+1)$-restricted complements on bidders $[m]$, items $[n]$, and valuations $w_{i,S} = \sum_{j \in S} v_{j,S \setminus \{j\},i}$ for all $i \in [m], j \in [n], S \subseteq [n], |S| \leq r$. It is then straightforward to see that the social welfare for Γ of a given strategy profile $s \in \Sigma$ equals the social welfare for Γ' of s, when s is considered as an allocation of the items $[n]$ to the players $[m]$, in the obvious way. Similarly, auctions with r-restricted complements can be converted to congestion games with $(r-1)$-externalities. It follows that any (randomized) approximation algorithm for the welfare maximization problem for auctions with r-restricted complements implies an approximation algorithm with the same approximation guarantee for the welfare maximization problem for congestion games with $(r-1)$-externalities, and vice versa.

7.2 Contributions and Outline

We consider in Section 7.4 the welfare maximization problem for the case of mixed 1-externalities and show that it is strongly NP-hard and $n^{1-\varepsilon}$-inapproximable for every $\varepsilon > 0$, even for $m = 2$ facilities. We also give a polynomial-time algorithm that solves the problem when the number of players is constant.

In light of this inapproximability result, we then focus in Section 7.5 on the problem of computing an optimal assignment for generalized congestion games with non-negative 1-externalities. We derive a polynomial time algorithm that solves this case for $m = 2$ facilities. We show that the problem gets strongly NP-hard for $m \geq 3$ facilities and therefore focus on approximation algorithms.

Thereto, we first study in Section 7.6 the polytope of the relaxation of a LP that will turn out to be useful in developing approximation algorithms for the problem. We give a characterization of the vertices of this polytope.

We derive in Section 7.7 a deterministic polynomial time 2-approximation algorithm for the case of non-negative 1-externalities with an arbitrary number of facilities. This algorithm computes an optimal solution to a natural LP relaxation of the problem and then iteratively rounds this solution to an integer solution, thereby losing at most a factor 2 in the value of the social welfare. We also show that the integrality gap of the underlying LP is close to 2 and therefore we suspect that the approximation factor of our algorithm is the best rounding algorithm possible.

The rounding procedure is non-trivial because it needs to take care of the dependencies between the players, resulting from the pairwise non-negative externalities. The key of our analysis is a probabilistic argument showing that these dependencies can always be resolved in each iteration such that the social welfare does not decrease by too much. We believe that this approach might be applicable to similar problems and is therefore of independent interest.

Our approach is flexible enough to extend the algorithm to more general settings. We do this in Section 7.9. One such generalization is to the non-symmetric version where the strategy sets of the players are not necessarily the same. We show that our 2-approximation algorithm can be adapted to the case where the facilities available to each player are restricted. We also obtain an improved $3/2$-approximation algorithm when every player is restricted to two facilities. The proof of the $3/2$-approximation factor crucially exploits a characterization of the extreme point solutions of the LP relaxation. We also extend our rounding algorithm to the case of r-externalities, and derive an $(r + 1)$-approximation algorithm.

We also settle a question left open by Blumrosen and Dobzinski [2006]. The authors showed that an optimal assignment can be computed efficiently for symmetric singleton congestion games with 1-externalities $\Gamma = (n, m, v)$ in case $v_{i,i',j} = v_{i,i'',j}$ and $v_{i,i,j} = 0$ for all $i, i', i'' \in [n], i \neq i', i \neq i''$ and $j \in [m]$, can be computed efficiently. We show that this problem becomes NP-hard when the constraint $v_{i,i,j} = 0, i \in [n], j \in [m]$ is dropped.

Lastly, in Section 7.10, we consider the problem of obtaining a truthful mechanism for our social welfare optimization problem, that runs in polynomial time and attains a good approximation factor. It requires some technical preliminary knowledge about mechanism design to state accurately our contribution towards solving this problem, so we refer to Section 7.3 for a more precise explanation of our contribution.

As for the relation of this work to auctions with r-restricted complements [Abraham et al., 2012], we note the following. As said in the previous section, all of our results on the welfare maximization problem for congestion games with r-externalities can be translated in a direct fashion to auctions with r-restricted complements. Independently in Abraham et al. [2012], a randomized r-approximation algorithm is given for the social welfare maximization for auctions with r-restricted complements. This randomized algorithm is also applicable to the congestion games with $(r-1)$-externalities). Interestingly, despite that the research of [Abraham et al., 2012] and the research discussed in this chapter has been carried out independently, the $(r+1)$-approximation algorithm presented in this chapter, for congestion games with r-externalities, turns out to be essentially a derandomization of the algorithm in Abraham et al. [2012]. Besides this approximation algorithm for the welfare maximization problem, Abraham et al. [2012] also provide some mechanism design related results, and propose a *truthful* $(1+\epsilon)$-approximation algorithm for a special type of auction with r-restricted complements, as well as a truthful mechanism for the general case that achieves an $O(\log^r m)$-approximation guarantee. These results can be translated and used in a direct manner, so that they apply to our congestion games studied here. Likewise, it is also straightforward to translate our mechanism design result of Section 7.10 to the setting of auctions with r-restricted complements.

7.3 Preliminaries on Mechanism Design

The result we present in Section 7.10 requires some preliminary knowledge on mechanism design. We present these preliminaries in this section.

Formally, an abstract *mechanism design setting* is a triple $F = (n, V, S)$ where $[n]$ is a set of players, S is a set of *outcomes*, and $V = (V_1, \dots, V_n)$, where V_i is called the *type set* of player i, for $i \in [n]$. A type set V_i consists of functions from S to $\mathbb{R}_{\geq 0}$. The functions in V_i are called *types*. A *deterministic (direct revelation) mechanism* is a function from $V = \times_{i \in [n]} V_i$ to $S \times \mathbb{R}^n$. Given a *type profile* $v \in V$, $(M(v))_1$ is referred to as the *outcome of the mechanism (w.r.t. v)*, and $(M(v))_{i+1}$ is referred to as player i's *payment* (w.r.t. v). A mechanism together with a type profile $v \in V$ induces a utility maximization game with players $[n]$, where player i's strategy set is V_i, and its utility given a strategy profile $\hat{v} \in V$ is defined as $u_i(\hat{v}) = v_i((M(\hat{v}))_1) - (M(\hat{v}))_{i+1}$. A *randomized mechanism* is a probability distribution over deterministic mechanisms. Just as for deterministic mechanisms, a randomized mechanism M together with a type profile $v \in V$ induces a game with players $[n]$, where player i's strategy set is V_i, and its utility given a strategy profile $\hat{v} \in V$ is defined as $u_i(\hat{v}) = \mathbf{E}[v_i(M(\hat{v})_1) - M(\hat{v})_{i+1}]$.

A randomized mechanism M may satisfy various desirable properties, some of which we list below.

- M is called *truthful* (also called *incentive compatible* or *strategy-proof*) *in expectation* if for all $v \in V$, v_i is a dominant strategy for player $i \in [n]$ in the game induced by (M, v). This means that for every type profile $\hat{v} \in V$, it holds that $\mathbf{E}[u_i(\hat{v})] \leq \mathbf{E}[u_i(v_i, \hat{v}_{-i})]$.
- M satisfies *universally non-negative payments* if for every deterministic mechanism M' in its support, the image of M' is a subset of $S \times \mathbb{R}^n_{\geq 0}$.
- M satisfies *universal individual rationality* if for every deterministic mechanism M' in its support, $u_i(v)$ is nonnegative in game (M', v) for all $v \in V$.
- Let $f : S \times V \to \mathbb{R}_{\geq 0}$ be a function. M is said to *approximate f to factor k* if it holds that $\mathbf{E}[f(M(v)_1, v)] \geq \max\{f(s, v) \mid s \in S\}/k$ for all $v \in V$.

We refer to a set of mechanism design settings as a *mechanism design problem.* A task that is often dealt with in the research area of (algorithmic) mechanism design, is for a given mechanism design problem, to define for each of its mechanism design settings a mechanism that satisfies the four properties above, and at the same time optimizes a given social welfare function, or at least approximates it up to some good factor. Given that we have satisfactory mechanisms for a mechanism design problem of the type just described, we additionally require to be able to compute the outcome of these mechanisms within time polynomial in the size of some natural representation of the mechanism design setting.

More precisely: we would like to have a randomized algorithm A for our mechanism design problem $\mathcal{F}$ that takes as input a mechanism design setting $F = (n, V, S) \in \mathcal{F}$ (where $V = (V_1, \ldots, V_n)$) and a reported type $\hat{v}_i \in V_i$ for all $i \in [n]$, such that A outputs an allocation and a payment for each player. A is required to run in expected polynomial time, and A is required to have the property that for every mechanism design setting $(n, V, S) \in \mathcal{F}$ that it can receive as input, it holds that the distribution of the output of A is the distribution of the output of a randomized mechanism that satisfies truthfulness in expectation, universally non-negative payments, universal individual rationality, and approximates the social welfare to an as good as possible factor. Because we look at mechanisms from the algorithmic perspective in this chapter, we will abuse terminology and sometimes interchange the concept of an *algorithm* with that of a *mechanism*.

In our case, we are facing congestion games with r-externalities, so the mechanism design problems we deal with constitute all mechanism design settings (n, V, S) for which there is a congestion game with r-externalities $\Gamma = (n, m, v)$ such that S is the set of strategy profiles of Γ. The type set V_i consists of the vectors of externalities in $\mathbb{R}$ that player $i \in [n]$ may have and report. Therefore, we are in essence interested in an algorithm/mechanism that runs in expected polynomial time and satisfies truthfulness in expectation, universal non-negative payments, universal individual rationality, and approximates the social welfare up to a good factor.

In [Abraham et al., 2012], for auctions with r-restricted complements an algorithm was proposed that satisfies all requirements, achieving an approximation factor of $O(\log^r m)$. This mechanism is easily adapted to the case of congestion games with $(r-1)$-externalities.

We present in Section 7.10 in this chapter a truthful $(r+2)$-approximation algorithm for congestion games with non-negative r-externalities. While this is a big improvement over the $O(\log^r m)$ algorithm of [Abraham et al., 2012], it should be noted that this improvement is due to that we take an additional liberty in the set of solutions that the mechanism is allowed to output: our mechanism requires the ability to output a solution in which some of the externalities are *disabled*. This means that the mechanism is able to allocate a set of players on the same facility, and set some of their externalities to 0. It can be interpreted as preventing some subsets of players to interact with each other, even if they are on the same facility. This prevents them from acquiring the benefits they would normally obtain from being together on the same facility.

7.4 Mixed 1-Externalities

We start off by studying the problem of optimizing the social welfare in congestion with mixed 1-externalities. It turns out that this problem is highly inapproximable, even for 2 facilities. Consider the following optimization problem:

Name: MAX-2FAC-CG-MIX-1EXT
Input: A description of a congestion game with mixed 1-externalities $\Gamma = (n, m, v)$.
Objective: Find a strategy profile $s \in \Sigma$ for Γ, that maximizes the social welfare $U(s)$.

The following result shows that it is impossible to find for this problem within polynomial time a strategy profile that attains a social welfare that is within a constant factor of the optimal social welfare.

Theorem 146. MAX-2FAC-CG-MIX-EXT *is strongly* NP*-hard and is not approximable to within a factor of* $n^{1-\epsilon}$ *in polynomial time, for all* $\epsilon > 0$*, assuming that* $\mathsf{P} \neq \mathsf{NP}$.

Proof. We give a polynomial time approximation preserving reduction from the strongly NP-hard optimization problem MAXCLIQUE.

Name: MAXCLIQUE
Input: A description of a graph $G = (V, E)$.
Objective: Find a clique in G with the maximum number of vertices among all cliques of G.

It is known that if $\mathsf{P} \neq \mathsf{NP}$, there does not exist an algorithm that approximates MAX-CLIQUE within a factor of $n^{1-\epsilon}$ for any $\epsilon > 0$ [Håstad, 1999, Zuckerman, 2007]. Given an instance $G = (V, E)$ of MAXCLIQUE, assume without loss of generality that $V = [n]$. We construct the congestion game with mixed 1-externalities $f(G) = (n, m, v)$ (i.e., an instance of MAX-2FAC-CG-MIX-EXT) as follows: Set

- $v_{i,i,1} = 1$ for all $i \in V$,
- $v_{i,i',1} = -2|V|$ for all $i, i' \in V$ such that $i \neq i'$ and $\{i, i'\} \notin E$,
- $v_{i,i',2} = 0$ for all $i, i' \in V$.

Let $C \subseteq V$ be a clique of G. Then there is a strategy profile for $f(G)$ that attains a social welfare of $|C|$: Let the players in C choose facility 1 as their strategy and let all other players choose facility 2 as their strategy. On the other hand, let $s \in \Sigma$ be a strategy profile of $f(G)$ that attains a social welfare of $U(s) \geq 0$, then there is a clique of size $U(s)$ in G: For two vertices i, i' for which $s_i = s_{i'} = 1$ the edge $\{i, i'\}$ is in E, otherwise $U(s)$ would be negative because $v_{i,i',1}$ is $-2|V|$. The set of vertices i such that $s_i = 1$ is therefore a clique in G.

We conclude that there is a one to one correspondence between cliques of G and strategy profiles that attain a non-negative social welfare in the congestion game $f(G)$, and the social welfare of a strategy profile corresponding to a clique C is equal to $|C|$. It follows that f is an approximation factor preserving reduction from MAXCLIQUE to MAX-2FAC-CG-MIX-EXT. □

The above result shows inapproximability when we fix the number of facilities. When we instead fix the number of players, it turns out to be rather easy to derive a polynomial time algorithm, as the next result shows.

Proposition 147. *The problem of computing the optimal strategy profile in a congestion game with mixed* 1*-externalities and a fixed number of players, can be solved in polynomial time.*

Proof. Suppose that we are given a congestion game with mixed 1-externalities Γ, and a partition $\mathcal{P}$ of the player set. Let $\Gamma^{\mathcal{P}}$ be the game obtained from Γ by setting $v_{i,i',j} = 0$ for all $i, i'' \in [n], j \in [m]$ for which it holds that i and i' are in distinct sets of $\mathcal{P}$ (and otherwise $v_{i,i',j}$ remains unaltered). Now consider the following problem $p(\mathcal{P})$: find the strategy profile of $\Gamma^{\mathcal{P}}$ that maximizes the social welfare, among the strategy profiles of $\Gamma^{\mathcal{P}}$ for which it holds that for all sets $P \in \mathcal{P}$, the players in P are allocated to the same facility. It is clear that this problem can be solved in polynomial time: just allocate the players in a set $P \in \mathcal{P}$ to

$$\arg_j \max \left\{ \sum_{i,i' \in P} v_{i,i',j} \mid j \in [m] \right\}.$$

Clearly, there is a partition $\mathcal{P}$ of $[n]$ such that the optimal strategy profile and optimal social welfare for Γ is the same as the optimal strategy profile and optimal social welfare for $\Gamma^{\mathcal{P}}$. Thus, it suffices to solve the problem $p(\mathcal{P})$ for all partitions $\mathcal{P}$ of $[n]$. This number of partitions is a constant known as the *nth Bell number* $\mathcal{B}_n$, and this proves the claim. □

7.5 Non-Negative 1-Externalities

Due to the inapproximability result for finding a social welfare maximizing strategy profile for congestion games with mixed 1-externalities, we will focus on the case that all externalities are non-negative in the remainder of this chapter. This section focuses on optimizing the social welfare in congestion games with non-negative 1-externalities. Central to our study of this problem will be the following integer LP for the problem.

Let $\Gamma = (n, m, v)$. Then the variables of the LP below are interpreted as follows: Every feasible solution x of the LP corresponds to a strategy profile $s(x) \in \Sigma$ of Γ. For $i, i' \in [n]$, $i \neq i'$ and $j \in [m]$, the variable $x_{i,j}$ is the $(0,1)$-variable that indicates whether player i is assigned to facility j under $s(x)$, and $x_{\{i,i'\},j}$ is the $(0,1)$-variable that indicates whether both players i and i' are assigned to facility j under $s(x)$. There are thus $m(n^2 + n)/2$ variables in total.

$$\max \sum_{j \in [m]} \left(\sum_{i \in [n] \setminus \{1\}} \sum_{i' \in [i-1]} (v_{i,i',j} + v_{i',i,j}) x_{\{i,i'\},j} + \sum_{i \in [n]} v_{i,i,j} x_{i,j} \right) \tag{7.1}$$

$$\text{s.t.} \sum_{j \in [m]} x_{i,j} = 1, \forall i \in [n] \tag{7.2}$$

$$x_{\{i,i'\},j} - x_{i,j} \leq 0, (\forall i, i' \in [n]), \forall j \in [m], \tag{7.3}$$
$$x_{\{i,i'\},j} \in \{0,1\}, (\forall i, i' \in [n]), \forall e \in [m]. \tag{7.4}$$

Because we assume that the externalities are all non-negative, it is clear that the optimal solution of the above LP corresponds to a social welfare optimizing strategy profile of Γ. The following result stands in stark contrast with Theorem 146.

Theorem 148. *Suppose we are given a congestion game Γ with positive externalities, with 2 facilities. Then it is possible to find an optimal strategy profile for Γ in polynomial time.*

Proof. We prove this by showing that the coefficient matrix of (7.1) is *totally unimodular*, in case $m = 2$. A matrix is said to be totally unimodular iff all of its square submatrices have determinant 0, 1 or -1. Let A be a p by q matrix. Then the following facts are well-known:

1. If A is totally unimodular, then for all $b \in \mathbb{Z}^q$, the vertices of the polyhedron $\{x \mid Ax \leq b\}$ are integer vectors.
2. If every column of A has at most two non-zero entries, each non-zero entry is -1 or 1, and these two values occur at most once per column, then A is totally unimodular.
3. If A is totally unimodular, then any matrix obtained from A by one of the following operations is also totally unimodular: 1.) taking the transpose, 2.) swapping two columns, 3.) deleting a column, 4.) introducing a column that is a copy of an existing column, 5.) introducing a column that is a zero vector or unit vector, 6.) multiplying a column by -1.

We show that the coefficient matrix for constraints (7.2) and (7.3) is totally unimodular. It then follows from point 3 in the above enumeration that the relaxation of this integer LP also has a totally unimodular coefficient matrix; completing the proof.

Let A be the coefficient matrix obtained from (7.2) and (7.3). Let A' be the matrix obtained from A by multiplying for all $i, i' \in [n]$ the column corresponding to variable $x_{\{i,i'\},1}$ by -1, and subsequently taking the transpose. Clearly, A' satisfies point 2 of the enumeration above, so A' is totally unimodular. By point 3, and the operations required to obtain A from A', it follows that A is totally unimodular as well. $\square$

Unfortunately, it turns out that for congestion games with non-negative 1-externalities with more than 2 facilities, the problem is strongly NP-hard. Consider the following decision problem:

Name: 3FAC-CG-NN-EXT
Input: A description of a congestion game with non-negative 1-externalities $\Gamma = (n, 3, v)$ with three facilities and a number $c \in \mathbb{Q}$.
Question: Is there a strategy profile $s \in \Sigma$ such that the social welfare $U(s) \geq c$?

Theorem 149. 3FAC-CG-NN-EXT *is strongly* NP*-complete.*

Proof. We give a polynomial time reduction from the *exact 3-satisfiability problem* (X3SAT) to an intermediate special case, X3SAT-PRIME, and finally we give a polynomial time reduction X3SAT-PRIME to 3FAC-CG-NN-EXT. Let us first define X3SAT and X3SAT-PRIME:

Name: X3SAT
Input: Sets $X = \{x_1, \ldots, x_p\}$, and $\mathcal{C} = \{C_1, \ldots, C_q\}\rangle$; X is called the set *variables* $\mathcal{C}$ is called the set of *clauses*. A clause $C_j \in \mathcal{C}$ is a vector of three *literals*. A literal is a variable or a *negation* of a variable (written $\neg x_i$, for $x_i \in X$).
Question: Is there an *exact truth assignment*? This means: does there exist a partition of X in two sets, T and F, such that for each clause $C_j \in \mathcal{C}$, precisely one literal of C_j is set to *true*, i.e., exactly one literal in C_j is in the set $T \cup \{\neg x_i \mid x_i \in F\}$?

X3SAT-PRIME is defined in exactly the same way as X3SAT, except that now the input satisfies three additional properties:

Property 1: each literal occurs in at least one clause.

Property 2: for all $x_i \in X$, and for all $k \in [3]$, the set $\{\ell \mid \exists C_j \in \mathcal{C} : \ell \in (C_j)_k\}$ does not contain both x_i and $\neg x_i$. In words: If a literal occurs in the kth position of a clause, then its negation does not occur in the kth position of any clause.

Property 3: for all $x_i \in X$, the sets $\{k \mid \exists C_j \in \mathcal{C} : x_i \in (C_j)_k\}$ and $\{k \mid \exists C_j \in \mathcal{C} : \neg x_i \in (C_j)_k)\}$ both have cardinality 1. In words: if a literal occurs on the kth position of a clause, then that literal does not appear on position j of any other clause, for $j \in [3]\backslash\{k\}$.

We are given a X3SAT instance $I = (X = \{x_1, \ldots, x_p\}, \mathcal{C} = \{C_1, \ldots, C_q\})$. A polynomial time reduction to X3SAT-PRIME then works as follows:

Suppose that I violates Property 1: there is an $i \in [p]$ such that x_i or $\neg x_i$ does not occur in any clause. Assume without loss of generality that $\neg x_i$ does not occur, and that x_i occurs at position 1 of some clause. We then introduce two variables d_1, d_2, and the clauses $(x_i, \neg x_i, d_1), (d_2, \neg d_1, d_1), (d_2, \neg d_2, d_1)$. It is not hard to see that the resulting instance is equivalent to I, but has 1 less variable that violates Property 1. Also note that this process does not introduce any variables that violate properties 2 and 3. We can repeat this process until there are no literals left that violate Property 1.

Suppose now that I satisfies Property 1, but not Property 2: there exists a variable x_i that is in position k of some clause, and its negation $\neg x_i$ is on the same position of another clause. We assume without loss of generality that $k = 1$. Then we replace all occurrences of $\neg x_i$ on the first position of a clause by negations of a new variable: $\neg x_i'$. We then introduce new variables $d_1, \ldots, d_5$ and the clauses $(x_i, d_1, d_2), (d_3, d_1, x_i'), (d_4, \neg d_4, d_2), (d_3, d_5, \neg d_5)$. Note that by the last two clauses that we added, d_3 and d_2 are false, in any exact truth assignment. The first clause added then forces d_1 to take the opposite value of x_i, and the second clause therefore forces x_i' to take the same value as x_i. So this new instance is equivalent to I, but with one less (variable,position)-pair that violates Property 2. Note that this procedure does not introduce any variables that violate Property 3. The procedure does however introduce some variables that violate Property 1, but this problem can be solved by re-running the above procedure for removing variables that violate Property 1.

Lastly, suppose that our instance I satisfies Properties 1 and 2, but not Property 3. Then there is a variable x_i such that either x_i or $\neg x_i$ occurs at more than 1 position among the clauses. Assume without loss of generality that literal x_i occurs at positions 1 and 2. We introduce a new variable x_i', and replace all occurrences of x_i on the second position of a clause, with x_i'. We also introduce the variables d_1 and d_2. and add the clauses $(x_i, \neg x_i, d_1), (d_2, \neg d_2, d_1)$. Note that in this new instance, under any exact truth assignment, d_i will be false, due to the second newly added clause. Therefore, the first newly added clause makes sure that x_i will take on the same value as x_i'. This new instance is thus equivalent to I, and in this new instance there is one less (literal,position,position)-pair that violates Property 3. This procedure does introduce some variables that do not satisfy Properties 1 and 2. But one can run the first two procedures to remedy this problem.

We conclude that it is possible to convert an X3SAT instance to an equivalent X3SAT-PRIME instance in polynomial time, using the three procedures described above.

Next, we reduce an instance $I = (X = \{x_1, \ldots, x_p\}, \mathcal{C} = \{C_1, \ldots, C_q\})$ of X3SAT-PRIME to an instance $f(I) = (n, 3, v)$ of 3FAC-CG-NN-EXT, as follows. We set $n = q + p$. Clause C_j corresponds to player j and variable i corresponds to player $q + i$, in a sense that will be explained later. For $i \in [p]$, let k_i and $k_{\neg i}$ be the numbers such that x_i occurs in position k_i of the clauses, and $\neg x_i$ occurs in position $k_{\neg i}$ of the clauses, and moreover, define $M = q+p+1$. We set values $v_{i,i',j}$ to 0 for all $i, i' \in [n]$, $j \in [3]$, except for the following ones:

- For all $i \in [p]$, the values $v_{q+i,q+i,k_i}$ and $v_{q+i,q+i,k_{\neg i}}$ are set to M.

- For all $i \in [p]$, $j \in [q]$ such that literal x_i occurs in C_j, the value v_{q+i,j,k_i} is set to 1.
- For all $i \in [p]$, $j \in [q]$ such that literal x_i occurs in C_j, the value $v_{q+i,j,k_{\neg i}}$ is set to 1.
- Let ℓ be a literal, and define $\#\ell$ as the number of clauses in $\mathcal{C}$ in which ℓ occurs. For all $k \in [3]$ and $j_1, j_2 \in [q]$ such that the same literal ℓ occurs in both C_{j_1} and C_{j_2}, the value $v_{j_1,j_2,k}$ is set to $1/(\#\ell)^2$. In other words, the total amount of 1 is divided equally among all values $v_{j_1,j_2,k}$ such that ℓ occurs in both C_{j_1} and C_{j_2}, for each literal ℓ.

Finally, the value c is set to $qM + q + p$.

There is a correspondence g from exact truth assignments for I to strategy profiles of $f(I)$: Let A be an exact truth assignment for I. If a literal x_i is true under A, this corresponds in $g(A)$ to player $q + i$ choosing facility k_i and player j choosing facility k_i as well, for all $j \in [q]$ such that literal x_i occurs in C_j. Likewise, if the literal $\neg x_i$ is true under A, this corresponds in $g(A)$ to player $q + i$ choosing facility $k_{\neg i}$ and player j choosing facility $k_{\neg i}$ as well, for all $j \in [q]$ such that literal $\neg x_i$ occurs in C_j.

It follows by construction that if A is an exact truth assignment for I, then $g(A)$ is a strategy profile for $f(I)$ that attains a social welfare of exactly $qM+q+p$. Note that in defining the correspondence g, and concluding that $g(A)$ attains a social welfare of $qM + q + p$ when A is an exact truth assignment, it is indeed essential that I satisfies Properties 1 to 3 above.

It remains to be shown that there is a strategy for $f(I)$ that attains a social welfare of at least $qM + p + q$ only if there is an exact truth assignment for I. First observe that if a player $q + i$, $i \in [p]$, chooses a facility other than k_i or $k_{\neg i}$, then at most $q - 1$ players will contribute M to the social welfare, hence the social welfare will not exceed $qM + q + p$, because of the large magnitude of M. So we may assume that all players $q + i$, $i \in [p]$ choose either k_i or $k_{\neg i}$ as their strategy. By construction, the maximum additional contribution to the social welfare that player $q + i$ can contribute, is equal to the total number of players $j \in [q]$ that choose the same facility as player $q + i$ such that x_i or $\neg x_i$ occurs in C_j. From this we conclude that the total contribution to the social welfare from the players $[n]\backslash[q]$ is at most $qM + q$. The maximum possible contribution to the social welfare by the players $j \in [q]$ is clearly at most q, and by construction this is attained iff for all $j \in [q]$, the chosen strategy k_j of player j is also chosen by all players j' such that $C_{j'}$ and C_j both have the same literal on position k. So the maximum possible social welfare is $qM + q + p$, and can only be attained by a strategy profile s that satisfies all the constraints that we just elicited. However, these constraints directly imply that s lies in the domain of g, and there is in that case an exact truth assignment A such that $g(A) = s$. □

It is still possible to approximate the problem within a factor of 2 from the optimal solution. We will show this in Section 7.7.

7.6 The Polytope of Feasible Fractional Solutions for 1-Externalities

We denote by $P(n, m)$ the polytope of feasible solutions of the relaxation of the integer LP (7.1). Denote by $d(n, m)$ the dimension of the points in $P(n, m)$. $P(n, m)$ is not integral for more than two facilities, but it is still possible to provide an interesting characterization of the vertices of $P(n, m)$. The material in this section is used in Sections 7.8 and 7.9.1, but is not used in the 2-approximation algorithm that we discuss in Section 7.7. In the discussion that follows, the subscripts used for the variables in the LP (7.1) will again be used here. We identify each such subscript with a different index in $[d(n, m)]$.

Let x be in $P(n, m)$. For $i, i' \in [n], i \neq i'$ and $j \in [m]$ we call the value $x_{\{i,i'\},j}$ *loose* if it is not equal to $x_{i,j}$ or $x_{i',j}$. Then it is clear that x does not have loose values if it is a vertex of $P(n, m)$.

For $j \in [m]$ and $w \in (0, 1]$, the *segment* $S(j, w)$ of x is the set of indices $\{(i, j) \in [d(n, m)] \mid x_{i,j} = w\} \cup \{(\{i, i'\}, j) \in [d(n, m)] \mid x_{\{i,i'\},j} = w\}$. Let $\mathcal{S} = \{S(j, w) \mid S(j, w) \neq \varnothing\}$ denote the set of non-empty segments. Note that $|\mathcal{S}|$ is finite (in fact: $|\mathcal{S}| \leq nm$).

The *segment hypergraph* $SH(x)$ of x, is the hypergraph $(\mathcal{S}, H, \ell)$. In this hypergraph, each segment in $\mathcal{S}$ is a vertex, $\ell : \mathcal{S} \to (0, 1]$ is the vertex labeling where $\ell(S(x, w)) = w$ for all $S(x, w) \in \mathcal{S}$, and H are the hyperedges $\{h_1, \ldots, h_n\}$, where $h_i = \{S(j, w) \in \mathcal{S} \mid x_{i,i,j} = w, j \in [m]\}$ for all $i \in [n]$.

In words: in the segment hypergraph, the segments are the vertices, each player corresponds to a hyperedge, and the vertices in this hyperedge are the segments that the player is assigned to (fractionally or integrally).

Example 150. Consider an instance of a congestion game with non-negative externalities with facility set [3] and player set [3]. Suppose that $v(i, i', j) = 1$ when $i, i', j \in [3]$ are all distinct, and 0 otherwise. Then, in the optimal solution of the LP relaxation of (7.1), the variables $\{x_{\{1,2\},3}, x_{\{1,3\},2}, x_{\{2,3\},1}, x_{1,2}, x_{1,3}, x_{2,1}, x_{2,3}, x_{3,1}, x_{3,2}\}$ are set to $1/2$, and the remaining variables are set to 0. There is thus one segment per facility,

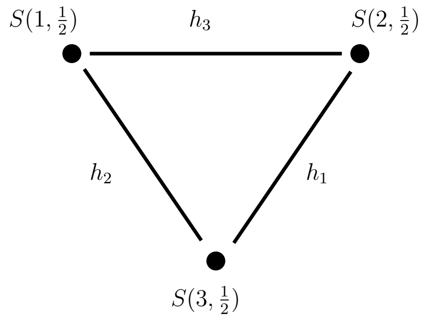

Figure 7.1: Example of a segment hypergraph of the optimal fractional solution of a congestion game instance with 3 players and 3 facilities.

and each player is assigned to 2 segments. The segment hypergraph is therefore the graph depicted in Figure 150.

The *incidence matrix* of a hypergraph (V, H) is defined as the $|V| \times |H|$ matrix where each row corresponds to a vertex of V and each column corresponds to a hyperedge of H. For $v \in V$ and $h \in H$, the value on the row of v and the column of h is 1 if $v \in h$ and 0 otherwise.

We can characterize the vertices of $P(n, m)$ as follows:

Theorem 151. *$x \in P(n, m)$ is a vertex of $P(n, m)$ if and only if it has no loose values, and the incidence matrix of $SH(x)$ has (full) rank $|\mathcal{S}|$.*

Proof. ($\Rightarrow$) Suppose that x is a vertex of $P(n, m)$. Then, as argued above, x obviously has no loose values. Let $SH(x) = (\mathcal{S}, H, \ell)$. It is clear that the following system of equations

$$\sum_{S \in h} y_S = 1 \qquad \forall h \in H \tag{7.5}$$

is satisfied when we set $y_S := \ell(S)$ for all $S \in \mathcal{S}$. It suffices to prove that there is no other solution to (7.5), because from the fact that there is a unique solution, it follows that the coefficient matrix of the system (7.5) is of (full) rank $|\mathcal{S}|$. This coefficient matrix is the transpose of the incidence matrix of $SH(x)$, so the incidence matrix has rank $|\mathcal{S}|$.

To see that the solution y for (7.5) corresponding to $\ell(S)$ is unique, suppose for contradiction that there is another solution y'. Now we have that for all $h \in H$, $\sum_{S \in h} y_S = \sum_{S \in h} y'_S = 1$. Next, define the vector $z = (z_S)_{S \in \mathcal{S}}$ where $z_S = y_S - y'_S$

for all $S \in \mathcal{S}$. Observe that for all $\epsilon > 0$, ϵz is not a zero vector and ϵz satisfies $\sum_{S \in h} \epsilon z_S = 0 = \sum_{S \in h} -\epsilon z_S$ for all $h \in H$. Choose ϵ to be any positive number smaller than $\min\{y_S - y_T, y_S \mid S, T \in \mathcal{S}\}$. Since all entries in y are strictly positive $y - \epsilon z$ and $y + \epsilon z$ are both solutions to (7.5). Moreover, if we now define the corresponding vertex labelings $\ell' : \mathcal{S} \to \mathbb{R}$ and $\ell'' : \mathcal{S} \to \mathbb{R}$ as $\ell'(S) = \ell(S) + \epsilon z_S$ and $\ell''(S) = \ell(S) - \epsilon z_S$, then it is straightforward to verify that the vertex-labeled hypergraphs $(\mathcal{S}, H, \ell')$ and $(\mathcal{S}, H, \ell'')$ are the segment hypergraphs of two points $x', x'' \in P(n, m)$ such that $x'/2 + x''/2 = x$. So x is a convex combination of two points in $P(n, m)$, contradicting the fact that x is a vertex of $P(n, m)$.

($\Leftarrow$) Let x be a point in $P(n, m)$ that is not a vertex and does not have loose values. We show that there is another point in $P(n, m)$ with a structurally equal segment hypergraph (i.e., has a segment hypergraph that differs only from $SH(x)$ in its vertex labeling). This shows that there is more than one solution to (7.5), which implies that the incidence matrix of $SH(x)$ is not of full rank. Because x is not a vertex of $P(n, m)$, there exists a non-zero vector $z \in \mathbb{R}^{d(n,m)}$ such that $x + z, x - z \in P(n, m)$ and $x + z$ and $x - z$ only differ from x at positions where x is fractional (i.e., non-$(0, 1)$), and such that $x + z, x - z$ is fractional on all positions where x is fractional. Define $\epsilon > 0$ as $\epsilon = (1/8) \min\{w \mid \exists j \in [m] \mid S(e, w) \in \mathcal{S}\} \cup \{|w_1 - w_2| \mid \exists j \in [m] : S(j, w_2), S(j, w_1) \in \mathcal{S}\}$. In words: ϵ is the minimum of: 1.) one quarter of the minimum difference between the value of any pair of segments that is on the same facility and 2.) the minimum value of a segment.

By convexity of $P(n, m)$, it must be that $x + \epsilon z, x - \epsilon z \in P(n, m)$. Observe that if two indices are in the same segment of $x + \epsilon z$, then they are in the same segment of $x - \epsilon z$, since $x - \epsilon z = (x + \epsilon z) - 2\epsilon z$. Also, if two indices are in different segments of $x + \epsilon z$, then they are in different segments of $x - \epsilon z$. This follows from the observation that $x + z$ and $x - z$ differ at each position by a sufficiently small amount $\epsilon/4$ (if this difference is large, then problems arise: it could be that two distinct segments of $x + z$ become one merged segment in $x - z$, and vice versa).

We conclude from this that the segment hypergraphs of $x + \epsilon z$ and $x - \epsilon z$ are structurally equal (i.e., the hypergraphs differ only in their vertex labeling). Consider now the point $y = x + \epsilon z/2$ and $z' = z/2$. Then $x = y - \epsilon z'$, and $y + \epsilon z', y - \epsilon z' \in P(n, m)$. Furthermore, $y + \epsilon z'$ and $y - \epsilon z'$ satisfy the above properties: $y + \epsilon z'$ and $y - \epsilon z'$ only differ from y at positions where y is fractional (i.e., non-$(0, 1)$), and $y + \epsilon z', y - \epsilon z'$ is fractional on all positions where y is fractional. By repeating the argument above, we thus conclude that $x = y - \epsilon z'$ has the same segment hypergraph as $y + \epsilon z'$. $\square$

It remains an open question to characterize combinatorially the class of hypergraphs of which the rank of its incidence matrix equals its number of vertices. For the special case of graphs, we do have such a characterization:

Theorem 152. *The rank of the incidence matrix of a graph $G = (V, E)$ equals $|V|$ if and only if all of its maximal connected components are non-bipartite.*

Proof. ($\Rightarrow$) Suppose that a maximal connected component $C = (V', E')$ of the graph is bipartite. Consider the system of linear equations

$$\sum_{v \in e} x_v = 1 \qquad \forall e \in E. \tag{7.6}$$

This system has an obvious solution x: assign $1/2$ to x_v for all $v \in V$. Because C is bipartite, there is another solution x' to this system: assign $1/2$ to x_v for all $v \notin V'$ and assign 1 to x_v for all $x \in V'$ on one side of the bipartition. Because there solution is not unique, the coefficient matrix of the above system (7.6) is not full rank, hence at most $|V| - 1$. This coefficient matrix is the transpose of the incidence matrix of G, so the rank of the incidence matrix of G is at most $|V| - 1$.

($\Leftarrow$) Given that the rank of the incidence matrix of G equals $|V|$, the system of linear equations (7.6) has a unique solution. This must be the solution x where $x_v = 1/2$ for all $v \in V$. Let $C = (V', E')$ be a maximal connected component of the graph. If C is bipartite, we can partition its vertices into V_1' and V_2' accordingly. By taking solution x and changing x_v to 1 for $v \in V_1'$, and changing x_v to 0 for $v \in V_2'$, we now obtain another solution to (7.6). A contradiction. $\square$

7.7 A 2-Approximation Algorithm for 1-Externalities

We now state for congestion games with 1-externalities an algorithm that computes a strategy profile of which the social welfare lies within a factor 2 of the maximum social welfare.

Some notation is needed in order to present the algorithm clearly. First, we extend U to the domain of fractional solutions in the natural way (where a fractional solution is a feasible solution for LP (7.1). So we will use U, and the term *social welfare*, to refer to the value of the objective function of LP (7.1), for the feasible solutions of LP (7.1).

For a congestion game with 1-externalities with non-negative 1-externalities $\Gamma = (n, m, v)$, we define for $j \in [m]$ and $a \in \mathbb{Q}$ the (j, a)-*boosted game* as the game obtained from Γ by introducing a copy j' of facility j that all players value a times as much as the original facility j. Formally, the (j, a)-boosted game of Γ is the game $\Gamma' = (n, m + 1, v)$ where $v_{i,i',m+1} = av_{i,i',j}$ for all $i, i' \in [n]$. In the (j, a)-boosted game, we refer to the introduced facility $m + 1$ as *the boosted facility*.

The algorithm below, parametrized by $a \in \mathbb{Q}$, is a rounding algorithm that takes as its starting point the fractional optimum solution of the relaxation of LP (7.1) (i.e., the LP where the constraints (7.4) are replaced by $0 \leq x_{\{i,i'\},j} \leq 1 (\forall i, i' \in [n], \forall e \in [m])$, and iteratively picks a facility $j \in [m]$ and assigns a set of players to the boosted facility in the (j, a)-boosted game.

In order to state the algorithm clearly, we fix for each facility $j \in [m]$ a strict total order $\preceq_j$ on $[m]$. The order $\preceq_j$ orders the players in any way that respects their fractional assignment on j, i.e., $\preceq_e$ is any order that satisfies $i \prec_j i'$ whenever $x_{i,j} < x_{i',j}$, for all $i, i' \in [n]$. Using these strict total orders, for $j \in [m], k \in [n]$ we define $P(j, k)$ as the player set $P \subseteq [n]$ for which it holds that $|P| = k$ and $i \in P$ if $i' \in P$ and $i \succ_j i'$. Informally, $P(j, k)$ consists of the k players with the highest fractional assignments on facility j.

Finally, for a fractional solution x for Γ, we define the (j, a, k)-*boosted assignment* obtained from x, for $j \in [m], a \in \mathbb{Q}, k \in [n]$, as the fractional solution to the (j, a)-boosted game where players $P(j, k)$ are assigned integrally to the boosted facility, and the remaining players are assigned according to x.

The description of our algorithm, BOOST(a), is given in Algorithm 1.

We prove next that if we set $a = 2$, then BOOST(a) is a 2-approximation algorithm for our optimization problem.

Theorem 153. *Algorithm* BOOST(2) *is a deterministic polynomial time* 2*-approximation algorithm for computing a social welfare maximizing strategy profile for a congestion game with non-negative* 1*-externalities.*

Proof. If Step 2.1 of our algorithm is valid (i.e., for a feasible solution x' to LP (7.1), for Γ', it is always possible to find in polynomial time a facility $j \in [m]$ and a number $k \in [n]$ such that the $(j, 2, k)$-boosted assignment (obtained from x') for the $(j, 2)$-boosted game (obtained from Γ')attains a social welfare that is at least the social welfare of x' for Γ'.), then the algorithm is valid, and polynomial runtime is straightforward to check: For step 1, solving an LP to optimality can be done in polynomial time using the ellipsoid method or an interior point method. For Step 2, note that the facility j that gets picked each time is from the set $[m]$ of *original* facilities, and not from the facilities that are introduced to the game during the execution of the algorithm. Each of these m facilities has at most n players, so at most nm possibilities need to be considered in Step 2.1. In each iteration of the loop of the algorithm, at least 1 additional player gets be assigned to a boosted facility, hence removed from $[m]$, so there are at most n iterations. For Step 3, polynomial runtime is obvious.

Given that Step 2.1 is valid, it is also easy to prove that Algorithm 1 outputs a solution for which the social welfare is within a factor $1/2$ from the optimal social welfare: Let us inspect the solution x' at the beginning of Step 3. Every player $i \in [n]$

Algorithm 1 BOOST(a): An LP rounding algorithm for approximating the social welfare optimizing strategy profile for a congestion game with non-negative 1-externalities.

Input: A congestion game with non-negative 1-externalities $\Gamma = (n, m, v)$.
Output: A strategy profile $s \in \Sigma$ for Γ.
Begin

1. Solve the relaxation of LP (7.1) for congestion game Γ, and let x be the optimal fractional solution. Let $\Gamma' = \Gamma$ and let $x' = x$.

2. Repeat the following until x' is integral:

2.1. Find a facility $j \in [m]$ and a number $k \in [n]$ such that the social welfare of the (j, a, k)-boosted assignment is at least the social welfare of x'.

2.2. Let Γ'' be the (j, a)-boosted game obtained from Γ' and let x'' be the (j, a, k)-boosted assignment obtained from x'. Subsequently, set $x' := x''$ and set $\Gamma' = \Gamma''$.

3. For $i \in [n]$, let j_i be the facility of $[m]$ such that in x', player i is assigned to a boosted copy of j_i. Let x'' be the integral solution of LP (7.1) for Γ' obtained by assigning i integrally to the original non-boosted facility j_i, for all $i \in [n]$. Output the strategy profile for Γ that corresponds to x''.

End

is assigned to a copy j' of a facility $j \in [m]$ for which it holds that $v_{i,i',j'} = 2v_{i,i',j}$. Therefore, assigning all players on such facilities j' to the original facility j, decreases the social welfare by a factor of at most 2. If we denote by SOL the social welfare of the strategy profile found by Algorithm 1, and if we denote by OPT the optimum social welfare, then the fact that Algorithm 1 is a 2-approximation algorithm follows from the following sequence of inequalities:[2]

$$SOL = U_\Gamma(x'') \geq \frac{1}{2}U_{\Gamma'}(x') \geq \frac{1}{2}U_\Gamma(x) \geq \frac{1}{2}OPT.$$

Therefore, all that needs to be shown that Step 2.1 of the algorithm can always be executed. This is implied by the following lemma.

Lemma 154. *Suppose that x' is a feasible solution to LP (7.1) for a congestion game with non-negative 1-externalities $\Gamma' = (n, m', v)$. Assume that there is a number $m \leq m'$ such that for each facility $j' \in [m'] \setminus [m]$ it holds that $x_{\{i,i'\},j'} \in \{0,1\}$ for all $i, i' \in [n]$. Then there is a facility $j \in [m]$ and a number $k \in [n]$ such that the social welfare of the $(j,2,k)$-boosted assignment (obtained from x') for the $(j,2)$-boosted game (obtained from Γ') is at least the social welfare of x' for Γ'. Moreover, j and k can be found in polynomial time.*

For the proof of Lemma 154, we need the following technical result:

Lemma 155. *For $n \in \mathbb{N}_{\geq 1}$, let $a_1, a_2, \ldots, a_n \in \mathbb{R}_{\geq 0}$, be a non-increasing sequence of non-negative numbers with a_1 non-negative, and let $b_1, \ldots, b_n \in \mathbb{R}$. Suppose that $\sum_{i \in [n]} a_i b_i \geq 0$. Then, there is a $k \in [n]$ such that $\sum_{i \in [k]} b_i \geq 0$.*

Proof. Let $n' \in [n]$ be the highest index such that $a_{n'} > 0$. There are two cases: either there is a $k \in [n'-1]$ for which the claim holds, or there is not such a k. For the latter case, we show that the claim must hold for $k = n'$. It follows from the following derivation:

$$\begin{aligned} 0 \leq & \sum_{i \in [n']} a_i b_i \\ = & (a_1 - a_2) \sum_{i \in [1]} b_i \\ & \quad + (a_2 - a_3) \sum_{i \in [2]} b_i \\ & \quad + \cdots \end{aligned}$$

[2]The function U is subscripted with the game of which it is the social welfare function.

$$+(a_{n'-1} - a_{n'}) \sum_{i \in [n-1]}^{n-1} b_i$$
$$+a_{n'} \sum_{i \in [n']} b_i$$
$$\leq \quad a_{n'} \sum_{i \in [n']} b_i.$$

□

Proof of Lemma 154. It suffices to only show existence of the appropriate $j \in [m]$ and $k \in [n]$, as finding them in polynomial time can then simply be done by complete enumeration of all (j,k)-pairs (because there are only mn such pairs).

For $j \in [m]$ and $k \in [n]$, let $\Delta(j,k)$ denote the amount by which social welfare increases when comparing the $(j,2,k)$-boosted assignment (obtained from x') for the $(j,2)$-boosted game (obtained from Γ') to the $(j,2,k-1)$-boosted assignment (obtained from x') for the (j,k)-boosted game (obtained from Γ'). Let $p(j,k)$ be the sole player in $P(j,k)\backslash P(j,k-1)$. We can express $\Delta(j,k)$ as $\Delta^+(j,k) - \Delta^-(j,k)$, where $\Delta^+(j,k)$ is the increase in social welfare due to the additional utility on the boosted facility, and $\Delta^-(j,k)$ is the loss in utility due to setting the assignment for player $p(j,k)$ to 0 on all facilities in $[m]$. For notational convenience, we define for all $i,i' \in [n], j \in [m]$ the value $w_{\{i,i'\},j} = v_{i,i',j} + v_{i',i,j}$ and the value $w_{i,j} = v_{i,i,j}$. Then we can write $\Delta^+(j,k)$ and $\Delta^-(j,k)$ as follows:

$$\Delta^+(j,k) \quad = \quad 2w_{p(j,k),j} + \sum_{i' \in [n]: i' \succ_j p(j,k)} 2w_{\{p(j,k),i'\},j},$$
$$\Delta^-(j,k) \quad = \quad \sum_{j' \in [m]} \left(x_{i,j'} w_{p(j,k),j'} + \sum_{i' \in [n]: i' \prec_j p(j,k)} x_{\{p(j,k),i'\},j'} \right).$$

Clearly, if we move for some $j \in [m]$ and $k \in [n]$ the players $P(j,k)$ to the boosted facility j' in the $(j,2)$-boosted game (obtained from Γ'), then the change in utility is $\sum_{i \in [k]} \Delta(j,i)$. We therefore need to show that there is a facility $j \in [m]$ and a number $k \in [n]$ such that $\sum_{i \in [k]} \Delta(j,i) \geq 0$.

To show this, let σ be the distribution on $\{\Delta(j,k) \mid j \in [m], k \in [n]\}$ given by

$$\sigma(\Delta(j,k)) = \frac{x_{p(j,k),j}}{\sum_{j \in [m], i \in [n]} x_{p(j,i),j}} = \frac{x_{p(j,k),j}}{\sum_{j \in [m], i \in [n]} x_{i,j}}, \forall j \in [m], k \in [n].$$

We derive the following bound on the expectation of a random variable with distribution σ:

$$
\begin{aligned}
&\mathbf{E}_{X\sim\sigma}[X] = \sum_{j\in[m],k\in[n]} \mathbf{Pr}_{X\sim\sigma}[X=\Delta(j,k)]\Delta(j,k)\\
&= \frac{1}{\sum_{j\in[m],i\in[n]} x_{i,j}} \sum_{j\in[m],k\in[n]} x_{p(j,k),j}\left(2w_{p(j,k),j} + \sum_{i'\in[n]:i'\succ_j p(j,k)} 2w_{\{p(j,k),i'\},j}\right.\\
&\quad\left. - \sum_{j'\in[m]}\left(x_{p(j,k),j'}w_{p(j,k),j'} - \sum_{i'\in[n]:i'\prec_j p(j,k)} x_{\{p(j,k),i'\},j'}w_{\{p(j,k),i'\},j'}\right)\right)\\
&= \frac{1}{\sum_{j\in[m],i\in[n]} x_{i,j}}\left(\sum_{j\in[m],k\in[n]} 2x_{p(j,k),j}w_{p(j,k),j}\right.\\
&\quad + \sum_{j\in[m],k\in[n]}\sum_{i'\in[n]:i'\succ_j p(j,k)} 2x_{p(j,k),j}w_{\{p(j,k),i'\},j}\\
&\quad - \sum_{j'\in[m],k\in[n]}\left(\sum_{j\in[m]} x_{p(j,k),j}\right) x_{p(j,k),f}w_{p(j,k),j'}\\
&\quad\left. - \sum_{j\in[m],k\in[n]} x_{p(j,k),j}\sum_{j'\in[m]}\left(\sum_{i'\in[n]:i'\prec_j p(j,k)} x_{\{p(j,k),i'\},j'}w_{\{p(j,k),i'\},j'}\right)\right)\\
&= \frac{1}{\sum_{j\in[m],i\in[n]} x_{i,j}}\left(\sum_{j\in[m],i\in[n]} 2x_{i,j}w_{i,j} + \sum_{\{i,i'\}:i,i'\in[n],i\neq i'}\sum_{j\in[m]} 2x_{\{i,i'\},j}w_{\{i,i'\},j}\right.\\
&\quad\left. - \sum_{j\in[m],i\in[n]} x_{i,j}w_{i,j} - \sum_{j\in[m],i\in[n]} x_{i,j}\sum_{j'\in[m]}\sum_{i'\in[n]:i'\prec_j i} x_{\{i,i'\},j'}w_{\{i,i'\},j'}\right)\\
&= \frac{1}{\sum_{j\in[m],i\in[n]} x_{i,j}}\left(\sum_{j\in[m],i\in[n]} x_{i,j}w_{i,j} + \sum_{\{i,i'\}:i,i'\in[n],i\neq i'}\sum_{j\in[m]} 2x_{\{i,i'\},j}w_{\{i,i'\},j}\right.\\
&\quad\left. - \sum_{\{i,i'\}:i,i'\in[n],i\neq i'}\left(\sum_{j'\in[m]} x_{\{i,i'\},j'}w_{\{i,i'\},j'}\right)\left(\sum_{j\in[m]}\max\{x_{i,j},x_{i',j}\}\right)\right)
\end{aligned}
\tag{7.7}
$$

$$\geq \frac{1}{\sum_{j\in[m],i\in[n]} x_{i,j}} \sum_{j\in[m],i\in[n]} x_{i,j} w_{i,j}$$
$$\geq 0.$$

In the above derivation, we make use of the following facts that hold for all $j \in [m], i, i' \in [n]$:

1. $\sum_{j'\in[m]} x_{i,j'} = 1$ (for the fourth equality),
2. $x_{i,j} = \max\{x_{i,j}, x_{i',j}\}$ if $i \succ_j i'$ (for the fifth equality),
3. $\sum_{j'\in[m]} \max\{x_{i,j'}, x_{i',j'}\} \leq 2$ (for the first inequality).

We can express $\mathbf{E}_{X\sim\sigma}[X]$ as the sum of the m terms

$$\left\{ T_j = \sum_{k\in[n]} \Delta(j,k)\mathbf{Pr}_{X\sim\sigma}[X = \Delta(j,k)] \;\middle|\; j \in [m] \right\}.$$

Because the expectation is non-negative, it holds that T_j is non-negative for at least one $j \in [m]$. We take this facility j and apply Lemma 155 to T_j (take $\Delta(j,i)$ for b_i and $\mathbf{Pr}_{X\sim\sigma}[X = \Delta(j,i)]$ for a_i). We conclude that there is a number $k \in [n]$ such that $\sum_{i\in[k]} \Delta(j,k) \geq 0$. □

□

7.8 A Matching Integrality Gap

In this section, we show that the integrality gap of the relaxation of LP (7.1) is very close to 2. This means that with respect to the integrality gap, the algorithm $\text{BOOST}(2)$ is almost the best possible with respect to the approximation ratio.

In order to show that the integrality gap is close to 2, we describe a sequence of sequences of increasingly larger examples of congestion games. It can subsequently be verified computationally that the integrality gap is close to 2 for large instances of these examples. We believe that the integrality gap of this sequence of examples tends to 2 in the limit, but so far we lack the analytical tools to prove this.

The construction of our examples is inspired by the insights from Section 7.6. We fix two parameters m and k, and we consider the k-uniform hypergraph on $[m]$, denoted $G(m,k)$. We define the instance $I(m,k)$ in such a way that the optimal fractional solution of $I(m,k)$ has $G(m,k)$ as its segment hypergraph. This is achieved as follows:

we let $I(m,k)$ consist of m facilities and $\binom{m}{k}$ players. We identify each player with a distinct edge of $G(m,k)$. The externality $v_{i,i',j}$ is set to 1 if the hyperedges $i,i' \in \left[\binom{m}{k}\right]$ both contain facility $j \in [m]$, and 0 otherwise. The value $v_{i,i,j}$ is set to 1 if hyperedge $i \in \left[\binom{m}{k}\right]$ contains facility $j \in [m]$.

We define $OPT_{\text{frac}}(m,k)$ as the feasible solution of the relaxation of LP (7.1) corresponding to instance $I(m,k)$, where each player is assigned with value $1/k$ to each of its facilities (conform the hyperedges in $G(m,k)$). For each facility $j \in [m]$ there are

$$\frac{k}{m}\binom{m}{k} = \binom{m-1}{k-1}$$

players that have j in their hyperedge, so the social welfare of $OPT_{\text{frac}}(m,k)$ is

$$\frac{m}{k}\binom{m-1}{k-1}^2.$$

We prove below that the social welfare of the optimal integral solution $OPT_{\text{int}}(m,k)$ for this instance is $\sum_{\ell=k-1}^{m-1}\binom{\ell}{k-1}^2$. Subsequently, evaluating (by computer) the expression $OPT_{\text{frac}}(m,k)/OPT_{\text{int}}(m,k)$ for particular choices of m and k indicates that the integrality gap approaches $\frac{2k-1}{k}$ as m gets larger. The largest integrality gap that we computed explicitly is 1.972013, for $m = 5000$ and $k = 71$.

Proposition 156.

$$OPT_{int}(m,k) = \sum_{\ell=k-1}^{m-1}\binom{\ell}{k-1}^2.$$

Proof. We claim that the social welfare of $\sum_{\ell=k-1}^{m-1}\binom{\ell}{k-1}^2$ is attained by the strategy profile $s(m,k)$ that results from the following iterative procedure: In the first iteration, assign to facility m all $\binom{m-1}{k-1}$ players that have facility m in their hyperedge. Because of the valuations these players have for facility 1, they add

$$\binom{m-1}{k-1}^2 \tag{7.8}$$

to the social welfare. After assigning the first $\binom{m-1}{k-1}$ players to facility m, there are

$$\binom{m}{k} - \binom{m-1}{k-1} = \binom{m-1}{k}.$$

players left. Each of these remaining players is identified as a distinct hyperedge of $G(m-1,k)$. In the second iteration, we assign to facility $m-1$ all $\binom{m-2}{k-1}$ players having $m-1$ in their hyperedge. This contributes

$$\binom{m-2}{k-1}^2$$

to the social welfare, and $\binom{m-2}{k}$ players remain, each of which are identified with a hyperedge in $G(m-2,k)$. In general, in the ℓth iteration we assign to facility $m-\ell+1$ all $\binom{m-\ell}{k-1}$ players having $m-\ell+1$ in their hyperedge. This contributes

$$\binom{m-\ell}{k-1}^2$$

to the social welfare. The total social welfare of $s(m,k)$ is therefore

$$\sum_{\ell=k-1}^{m-1} \binom{\ell}{k-1}.$$

We prove, using induction on m, that the utility of $s(m,k)$ is optimal for any choice of k. Our base case is $m = k-1$ for which optimality trivially holds. As our induction hypothesis, assume that optimality holds for $m = \ell, \ell \in \mathbb{N}$. We prove that optimality must also hold for $m = \ell + 1$. Suppose for contradiction that optimality does not hold. Then, in the optimal solution, no facility $j \in [m]$ has all of its $\binom{m-1}{k-1}$ players (i.e., the players that have j in their hyperedge) assigned to it. Pick an arbitrary facility j, and note that $I(m,k)$ restricted to all players that do not have facility j in their hyperedge, is equivalent to the instance $I(m-1,k)$. By our induction hypothesis, $s(m,k)$ is optimal for this instance. Therefore, if there is a facility $j \in [m]$ in the optimal solution such that all players that have j in their hyperedge are assigned to j, then the optimal social welfare is the sum of

$$\binom{m-1}{k-1}^2$$

and the social welfare of $s(m,k)$. In other words, the social welfare would then be

$$\sum_{\ell=k-1}^{m-1} \binom{\ell}{k-1}, \tag{7.9}$$

as claimed.

Suppose therefore that there is no facility $j \in [m]$ in the optimal solution such that all players that have j in their hyperedge are assigned to j. It is clear that no player is assigned to a facility that is not in its hyperedge. Let $j' \in [m]$ be the facility with most players assigned to it. There must be a player $i \in [n]$ that has j' in its hyperedge and is assigned to a facility $j'' \in [m]$ distinct from j'. By reassigning i to j', the social welfare can not decrease: j' is the facility with most players assigned to it, so the increase in social welfare due to i getting assigned to j' is at least as much as the decrease in social welfare due to i leaving j''. By the same argument, one can assign all players that have j' in their hyperedge to j', without decreasing the social welfare. □

7.9 Variations on the Problem

We study in this section the social welfare optimization problem for two generalizations and one special case of the class of congestion games with non-negative 1-externalities. Section 7.9.1 concerns the generalization to asymmetric strategy sets, i.e., where the strategy sets of the players need not all be the same. Section 7.9.2 studies congestion games with r-externalities, for $r > 1$. In both cases, we see that it is possible to use the algorithm of Section 7.7 in order to obtain approximation algorithms for variations of our main congestion game problem. The special case we study concerns a class of externalities called *affine externalities* and is studied in Section 7.9.3.

7.9.1 Asymmetric Strategy Sets

The strategy sets of the players can be restricted without loss of generality. Up to this point, the set of strategies of each player have been assumed to be the entire set of facilities in the game. Instead, we can set for each player its to any subset of facilities. Let us call such a game that falls within this generalization, a *asymmetric congestion game with non-negative* 1*-externalities*. Such a game is represented by the quadruple (n, m, v, Σ), where v is the vector containing the values $v_{i,i',j}$ for $i, i' \in [n], j \in [m]$ and $\Sigma = \times_{i \in [n]} \Sigma_i$ are the strategy profiles of Γ. This generalization can be 2-approximated using an only slightly modified version of BOOST(2): the only change that needs to be made is that in Step 1 of the algorithm, the algorithm solves a a relaxation of a modified version of LP (7.1), which we refer to as the *asymmetric* version of LP (7.1): This modified LP has the same objective function as (7.1), and the constraint set consists of all constraints in (7.1), and the additional constraints that $x_{i,j} = 0$ for all $i \in [n], j \notin \Sigma_i$. It can be seen that the rounding procedure that follows after Step 1 of the algorithm then puts every player $i \in [n]$ on a facility $j \in [m]$

for which $x_{i,j} \neq 0$, so it produces a feasible integral solution, and the remainder of the proof of Theorem 153 is still valid.

Corollary 157. *There exists a deterministic polynomial time* 2*-approximation algorithm for computing a social welfare maximizing strategy profile for a congestion game with non-negative* 1*-externalities, even when the strategy sets of the players are asymmetric.*

For the special case that each player's strategy set is of size 2, we can improve the approximation factor.

Proposition 158. *There is a deterministic polynomial time* $(3/2)$*-approximation algorithm for computing a social welfare maximizing strategy profile for a congestion game with non-negative* 1*-externalities, if the strategy sets are asymmetric and all of cardinality* 2*.*

Proof. Consider the algorithm $\text{BOOST}(3/2)$ with the modification described above (i.e., in Step 1 of $\text{BOOST}(3/2)$, the algorithm solves the relaxation of the asymmetric version of LP (7.1) for input asymmetric congestion game with 1-externalities $\Gamma = (n, m, v, \Sigma)$). Analogous to the proof of Theorem 153, we conclude that this modified algorithm is a valid polynomial time (3/2)-approximation algorithm if it is true that Step 2.1 can always be executed. So it suffices to show that the following variation of Lemma 154 holds:

Lemma 159. *Suppose that* x' *is a feasible solution to the asymmetric version of LP (7.1) for an asymmetric congestion game with non-negative* 1*-externalities* $\Gamma' = (n, m', v, \Sigma)$*. Assume that there is a number* $m \leq m'$ *such that for each facility* $j' \in [m'] \setminus [m]$ *it holds that* $x_{\{i,i'\},j'} \in \{0, 1\}$ *for all* $i, i' \in [n]$*, and* $[m] \cap \Sigma_i \leq 2$ *for all* $i, i' \in [n]$*. Then there is a facility* $j \in [m]$ *and a number* $k \in [n]$ *such that the social welfare of the* $(j, 3/2, k)$*-boosted assignment (obtained from* x'*) for the* $(j, 3/2)$*-boosted game (obtained from* Γ'*) is at least the social welfare of* x' *for* Γ'*. Moreover,* j *and* k *can be found in polynomial time.*

The proof of Lemma 159 is completely analogous to the proof of Lemma 154, substituting the factor of 2 with the factor of 3/2: We define $\Delta(j, k)$ as the amount by which social welfare increases when comparing the $(j, 3/2, k)$-boosted assignment obtained from x' to the $(j, 3/2, k-1)$-boosted assignment obtained from x', and we show that $\mathbf{E}_{X \sim \sigma}[X]$ is non-negative, where σ is defined as in the proof of Lemma 154. Non-negativity of $\mathbf{E}_{X \sim \sigma}[X]$ follows by following the derivation as in the proof of Lemma 154 up to the point (7.7) (appropriately replacing the factor 2 by $3/2$ in doing so), and then proceeding as follows:

By the restriction on the strategy sets of the players, it follows that the segment hypergraph $SH(x')$ of x' consists of components that are either isolated vertices (with a hyperedge of cardinality 1), and graphs. The isolated vertices correspond to segments to which players are assigned integrally. From Theorems 151 and 152 it follows that all components of $SH(x')$ that are graphs are non-bipartite. Because non-bipartite graphs have odd cycles, the only way the vertices of this component in the segment hypergraph can be labeled is by assigning all vertices value $1/2$. Because all vertices in $SH(x')$ are either $1/2$ or 1, it must be that $x_{\{i,i'\},j} \in \{0, 1/2, 1\}$ for all $i, i' \in [n], j \in [m]$. This in turn implies that

$$\sum_{j\in[m]} \max\{x_{i,j}, x_{i',j}\} \leq \frac{3}{2} \qquad \forall i, i' \in [n].$$

Using the above inequality, we can thus bound $\mathbf{E}_{X\sim\sigma}[X]$ as follows, continuing from (7.7):

$$\begin{aligned}
\mathbf{E}_{X\sim\sigma}[X] &= \frac{1}{\sum_{j\in[m],i\in[n]} x_{i,j}} \Bigg(\frac{1}{2} \sum_{j\in[m],i\in[n]} x_{i,j} v_{i,j} \\
&\quad + \sum_{\{i,i'\},i,i'\in[n],i\neq i'} \sum_{j\in[m]} \frac{3}{2} x_{\{i,i'\},j} w_{\{i,i'\},j} \\
&\quad - \sum_{\{i,i'\},i\neq i'} \Bigg(\sum_{j'\in[m]} x_{\{i,i'\},j'} w_{\{i,i'\},j'} \Bigg) \Bigg(\sum_{j\in[m]} \max\{x_{i,j}, x_{i',j}\} \Bigg) \Bigg) \\
&\geq \frac{1}{\sum_{j\in[m],i\in[n]} x_{i,j}} \frac{1}{2} \sum_{j\in[m],i\in[n]} x_{i,j} v_{i,j} \\
&\geq 0.
\end{aligned}$$

The proof is completed when we now apply Lemma 155, in the same way as in the proof of Lemma 154. $\square$

Example 150 in Section 7.6 shows that the integrality gap of the modified LP for strategy sets of size 2, is $\frac{3}{2}$. This matches the approximation ratio of Proposition 158.

7.9.2 Externalities on Bigger Sets of Players

We now consider the problem of finding a social welfare maximizing strategy profile for a congestion games with non-negative r-externalities, for $r \geq 2$. We will show that a simple adaptation of the algorithm BOOST$(r+1)$ returns an $(r+1)$-approximate solution. The only change that needs to be made is that the relaxation of the following generalization of LP (7.1), which we refer to as the *generalized* variant of LP (7.1), is solved in Step 1 of BOOST$(r+1)$: Suppose we are given a congestion game with r-externalities $\Gamma = (n, m, v)$. Every feasible solution x of the generalized variant of LP (7.1) again corresponds to a strategy profile $s(x)$ of Γ. For each player $i \in [n]$ and facility $j \in [m]$, the generalized variant of LP (7.1) the variable $x_{i,j}$ again indicates whether $s_i(x) = j$ for $i \in [n], j \in [m]$. Moreover, for every facility $j \in [m]$ and set $P \subset [n], |P| \leq r+1$, there is a variable $x_{P,e}$ that indicates whether all players in P are assigned to j. For notational convenience, let $w_{P,j} = \sum_{i \in P} v_{i,P\setminus\{i\},j}$ for all $P \subseteq [n], |P| \leq r+1, j \in [m]$. The generalized variant of LP (7.1) reads as follows:

$$\max \sum_{j\in[m]} \sum_{P:P\subseteq[n],|P|\leq r+1} w_{P,j} x_{P,j} \tag{7.10}$$

$$\text{s.t.} \sum_{j\in[m]} x_{i,j} = 1, \forall i \in [n], \tag{7.11}$$

$$x_{P,j} - x_{i,j} \leq 0, \ (\forall P \subseteq [n] : |P| \leq r+1), \forall i \in P, \forall j \in [m], \tag{7.12}$$

$$x_{P,j} \in \{0,1\}, \ (\forall P \subseteq [n] : |P| \leq r+1), \forall j \in [m] \tag{7.13}$$

We prove next that the adapted version of BOOST$(r+1)$ is a valid polynomial time $(r+1)$-approximation algorithm.

Proposition 160. *There is a deterministic polynomial time $(r+1)$-approximation algorithm for computing a social welfare maximizing strategy profile for a congestion game with non-negative r-externalities.*

Proof. Just like the proof of Proposition 158, this proof is very similar to the proof of Theorem 153. When reasoning analogous to the proof of Theorem 153, we conclude that the adaptation of BOOST$(r+1)$ runs in polynomial time and correctly outputs a $(r+1)$-approximate solution if it holds that Step 2.1 of the algorithm can always be executed. To see that Step 2.1 can always be done, we follow the proof of Lemma 154, while substituting the factor of 2 with the factor of $r+1$: We define $\Delta(j,k)$ as the amount by which social welfare increases when comparing the $(j, r+1, k)$-boosted assignment obtained from x' to the $(j, r+1, k-1)$-boosted assignment obtained from x', and we show that $\mathbf{E}_{X\sim\sigma}[X]$ is non-negative, where σ is defined as in the proof of Lemma 154. Proving that $\mathbf{E}_{X\sim\sigma}[X]$ is non-negative is done following the derivation in the proof of Lemma 154 up to the point 7.7. We then proceed as follows:

$$\mathbf{E}_{X\sim\sigma}[X] = \frac{1}{\sum_{j\in[m],i\in[n]} x_{i,j}} \left(\sum_{j\in[m]} \sum_{P:P\subseteq[n],|P|\le r+1} (r+1)x_{P,j}w_{P,j} \right.$$
$$\left. - \sum_{P:P\subseteq[n],|P|\le r+1} \left(\sum_{j'\in[m]} x_{P,j'}w_{P,j'} \right) \left(\sum_{j\in[m]} \max\{x_{i,j} \mid i\in P\} \right) \right).$$

Now by the observation that

$$\sum_{j\in[m]} \max\{x_{i,j} \mid i\in P\} \le r+1 \qquad \forall P\subset[n] : |P|\le r+1,$$

we conclude $\mathbf{E}_{X\sim\sigma}[X] \ge 0$. We can thus apply Lemma 155 to $\mathbf{E}_{X\sim\sigma}[X]$, which completes the proof. □

7.9.3 A Special Case: Affine Externalities

In this section, we study a very special case of our problem: congestion games with non-negative *affine externalities*. These are the congestion games with non-negative 1-externalities $\Sigma(n,m,v)$ for which it holds that $v_{i,i',j} = v_{i,i'',j}$ for all $i,i',i''\in[n], i\ne i', i\ne i''$, and all $j\in[m]$. I.e., there are numbers $a_{i,j}, b_{i,j}\in\mathbb{Q}_{\ge 0}$ for $i\in[n], j\in[m]$, such that $u_i(s) = a_{i,s_i}|\{i' \mid s_{i'} = s_i\}| + b_{i,s_i}$ for all $s\in\Sigma$. We therefore represent Σ by the quadruple (n,m,a,b), were a is the vector containing the values $a_{i,j}, i\in[n], j\in[m]$, and b is the vector containing the values $b_{i,j}, i\in[n], j\in[m]$.

The motivation for studying this special case is that Blumrosen and Dobzinski [2006] show that if $b_{i,j} = 0$ for all $i\in[n], j\in[m]$, then the optimal solution can be found in polynomial time. Allowing $b_{i,j}$ to be non-zero is thus one of the simplest generalizations that comes to mind.

The decision version of the social welfare optimization problem, for this class of games, is defined as follows:

Name: CG-AFF-NN-EXT
Input: A description of a congestion game with non-negative affine externalities $\Gamma = (n,m,a,b)$ and a number $c\in\mathbb{Q}$.
Question: Is there a strategy profile $s\in\Sigma$ such that the social welfare $U(s)\ge c$?

We show that this decision problem is NP-complete. This contrasts with the polynomial time result of Blumrosen and Dobzinski [2006].

Theorem 161. CG-AFF-NN-EXT *is strongly* NP*-complete.*

Proof. Clearly, this problem is in NP. To show strong NP-hardness we describe a polynomial time reduction from the following problem, which is known to be strongly NP-complete [Garey and Johnson, 1979]:

Name: TRIANGLE-PARTITION
Input: A description of a graph $G = (V, E)$ where $n = |V|$ is a multiple of 3, and each $v \in V$ occurs in at least one triangle.
Question: Is there a *triangle partition* for G? I.e., a partition $\{V_1, \ldots, V_{n/3}\}$ of V such that for all $k \in [n/3]$, it holds that $|V_k| = 3$ and the subgraph of G induced by V_k is a triangle.

Given an instance $G = (V, E)$ of TRIANGLE-PARTITION, assume without loss of generality that $V = [n]$. We construct an instance $f(G) = (\Gamma, c)$ of AFFINE-CG-NN-EXT. The reduction works as follows: The player set of Γ is the set of vertices of G, i.e., $[n]$. For each triangle t in G, we introduce a separate facility j_t. This runs in polynomial time since there are $\Theta(n^3)$ triples of vertices in total in G, and for each triple we can easily check whether the triple forms a triangle in G. The coefficients of the utility functions are defined as follows: $a_{i,j_t} = 0$ and $b_{i,j_t} = 0$ when $i \in [n]$ does not occur in triangle t. Otherwise, $a_{i,j_t} = \epsilon$ and $b_{i,j_t} = 1$, where $\epsilon \in \mathbb{Q}_{>0}$ is any sufficiently small rational number. [3] Lastly, we set $c = n + 3n\epsilon$.

We claim that for Γ there exists a strategy profile with social welfare at least c iff there exists a triangle partition for G:

First let us assume that for $f(G)$ there exists a strategy profile $s \in \Sigma$ with social welfare at least c. Then each player $i \in [n]$ must choose under s a facility j_t for which $b_{i,j_t} = 1$. It follows that each facility is chosen by at most 3 players. A player that chooses under s a facility that only 1 or 2 players choose, contributes at most $1 + 2\epsilon$ to the social welfare. If such a player exists, then the social welfare is at most $(n-1)(1+3\epsilon)+1+2\epsilon$, which is not the case. Therefore, under s, every player chooses a facility with exactly three players assigned to it, and by the direct correspondence between facilities and triangles, as well as players and vertices, it follows that G has a triangle partition.

For the other direction, let us assume that there is a triangle partition $\{V_1, \ldots, V_{n/3}\}$ for G. Let $t(k)$ denote the triangle corresponding to V_k, for $k \in [n/3]$. Then let s be the strategy profile in which a player $i \in V_k$ chooses facility $j_{t(k)}$. It is easily seen that $U(s) = k + 3k\epsilon$. □

[3] To be precise, it suffices to choose ϵ as any number a/b less than $1/n^2$ such that $\log a$ and $\log b$ are polynomial in n.

One could use the algorithm $\text{BOOST}(2)$ of Section 7.7 to find a 2-approximate solution to this affine special case of the problem. However, it is easy to see that there is a much simpler and faster algorithm possible for this case, that works as follows: When given as input a congestion game with non-negative affine externalities $\Gamma = (n, m, a, b)$, let the strategy profile $s \in \Sigma$ be given by

$$s_i = \arg_j \max\{b_{i,j} \mid j \in [m]\}$$

for all $i \in [n]$, and let the strategy profile $s' \in \Sigma$ be given by

$$\arg_j \max\left\{\sum_{i' \in [n]} n a_{i',j} \mid j \in [m]\right\}$$

for all $i \in [n]$. Output s if $U(s) \geq U(s')$ and output s' otherwise.

This is a 2-approximation algorithm because it is shown in Proposition 4.3 of Blumrosen and Dobzinski [2006] that if $b_{i,j} = 0$ for all $i \in [n]$, $e \in [m]$, then the strategy profile s' is optimal. It is obvious that if $a_{i,j} = 0$ for all $i \in [n]$, $j \in [m]$, then the strategy profile s is optimal. Therefore, for the case that neither a nor b is necessarily 0, the maximum of $U(s)$ and $U(s')$ must be within a factor of 2 from the optimal social welfare.

7.10 Truthful Polynomial Time Mechanisms for Non-Negative Externalities

In this section, we construct a mechanism for congestion games with non-negative r-externalities, that is truthful in expectation, runs in polynomial time in expectation, always charges non-negative payments, and is always individually rational. On the negative side, the mechanism requires the ability to output a solution in which some of the externalities are *disabled*. This means that the mechanism is able to allocate a set of players P on the same facility j, but set some of the externalities $v_{i,P',j}, i \in P, P' \subseteq P \setminus \{i\}$ to 0. It can be interpreted as preventing some subsets of players to interact with each other, preventing them from acquiring the benefits they would normally obtain from being together on the same facility.

We emphasize that the specification the set S of externalities that is to be disabled is part of the set of outcomes that the algorithm may generate, and that the approximation guarantee of the algorithm is still with respect to the optimal outcome, which is clearly attained when none of the externalities are disabled. In other words, the approximation guarantee of the algorithm is still with respect to the original optimum, and we do not "cheat" by measuring the approximation factor with respect to the optimum solution attainable with the externalities S disabled.

Formally, the output of the mechanism, for a congestion game $\Gamma = (n, m, v)$ (when given as input the vectors of externalities that the players of Γ report), is therefore a pair (s, t), where s is a strategy profile, and t is a specification of which of the externalities of v need to be disabled.

We start by describing a randomized $(r + 2)$-approximation algorithm for the the problem of maximizing the social welfare. This algorithm does not necessarily output a strategy profile for the input game, but rather something that we refer to as an *extended* strategy profile: In an extended strategy profile, players are allowed to not be allocated to any facility. A player that is not allocated, receives utility 0. We extend the domain of the social welfare function U in the natural way to extended strategy profiles. Moreover, if x is a feasible $(0, 1)$-solution to LP (7.10), and we obtain x' from x by setting $x_{i,j}$ to 0 for any set of players $i \in P \subseteq [n], j \in S \subseteq [m]$, then note that x' is not necessarily a feasible $(0, 1)$-solution to LP (7.10). We still refer to x' as an *extended* feasible $(0, 1)$-solution to LP (7.10).

Despite that we already have Algorithm 1 for maximizing the social welfare, we need this weaker randomized algorithm in order to solve our target mechanism design problem. The algorithm is again based on rounding the optimal fractional solution to the relaxation of LP (7.10). Our rounding algorithm has the interesting property that we can derive an exact polynomial-time computable expression for the expected social welfare of the solution that it outputs, and for the expected utilities of the players. This property is crucial for our truthful mechanism to work.

Let x be a extended feasible $(0, 1)$-solution to LP (7.10) for a congestion game with r-externalities Γ. We denote by $s(x)$ the extended strategy profile where $s_i = j$ iff $x_{i,j} = 1$, for $i \in [n], j \in [m]$.

The description of our randomized algorithm is given in Algorithm 2. It is an adaptation of the randomized algorithm of Abraham et al. [2012]. This adaptation is necessary in order to get an exact and easy-to-compute expression on the expected social welfare of the solution that it outputs.

It is clear that Algorithm 2 returns a valid extended strategy profile, as all constraints of LP (7.10) are satisfied at the end of the execution of the algorithm.

We proceed with deriving the expected value of the social welfare of the extended strategy profile output by Algorithm 2.

Lemma 162. *Let x' be the extended $(0, 1)$-solution to LP (7.10) that Algorithm 2 has produced at the end of its execution, when given as input a congestion game with non-negative r-externalities $\Gamma = (n, m, v)$. Moreover, let x be the optimal fractional*

Algorithm 2 A randomized LP rounding algorithm for approximating the social welfare optimizing strategy profile for a congestion game with non-negative externalities.

Input: A description of a congestion game with with non-negative r-externalities $\Gamma = (n, m, v)$.
Output: An extended strategy profile s for Γ.
Begin

1. Solve the relaxation of LP (7.10) for congestion game Γ, and let x be the optimal fractional solution. Let $\Gamma' = \Gamma$ and let $\mathbf{x}' = x$.

2. Repeat the following until x' is integral:

2.1. Select a facility $j \in [m]$ and select a number $p \in [0, 1]$ uniformly at random.

2.2. For all $i \in [n]$: if $x_{i,j} \geq p$, and facility j *has not* been selected in a previous iteration, then round $x'_{i,j}$ to 1 and round $x'_{i,j'}$ to 0 for all $j' \in [m] \backslash \{j\}$. If $x_{i,j} \geq p$ and facility j *has* been selected in a previous iteration, then round $x'_{i,j'}$ to 0 for all $j' \in [m]$.

(In both cases, for all $P \subseteq [n], |P| \leq r + 1$, round $x_{P,j}$ to 1 if $x'_{i,j} = 1$ for all $i \in [P]$ and round $x_{P,j}$ to 0 if there exists a player $i \in P$ such that $x'_{i,j} = 0$.)

3. Output the extended allocation $s(x')$.
End

solution to the relaxation of LP (7.10) for Γ. *Then:*

$$\mathbf{Pr}[x'_{P,j} = 1] = \frac{x_{P,j}}{1 + \sum_{j' \in [m] \setminus \{j\}} \max\{x_{i,j'} \mid i \in P\}},$$

for all $j \in [m], P \subseteq [n], |P| \leq r+1$.

Proof. We use a similar style of notation as in Langberg et al. [2006] (where a similar algorithm is analyzed for a similar problem): For a set of players $P \subseteq [n]$ a facility $j \in [m]$, and an iteration $k \in \mathbb{N}_{>0}$ of the loop of Algorithm 2, we use $P \underset{k}{\rightarrow} j$ to denote the event that j gets selected and all of the players in P get rounded in iteration k. We use $* \underset{<k}{\not\rightarrow} j$ to denote the event that j does not get selected in an iteration before k. Likewise, we use $P \underset{<k}{\not\rightarrow} *$ to denote the event that for all $j \in [m], i \in P$, $x'_{i,j}$ does not get rounded in an iteration before k.

$$\begin{aligned}
\mathbf{Pr}[x'_{P,j} = 1] &= \sum_{k=1}^{\infty} \mathbf{Pr}\left[P \underset{k}{\rightarrow} j \cap * \underset{<k}{\not\rightarrow} j \cap \bigcap_{i \in P} \{j\} \underset{<k}{\not\rightarrow} *\right] \\
&= \sum_{k=1}^{\infty} \mathbf{Pr}\left[P \underset{k}{\rightarrow} j \,\middle|\, * \underset{<k}{\not\rightarrow} j \cap \bigcap_{i \in P} \{i\} \underset{<k}{\not\rightarrow} *\right] \mathbf{Pr}\left[* \underset{<k}{\not\rightarrow} j \cap \bigcap_{i \in P} \{i\} \underset{<k}{\not\rightarrow} *\right] \\
&= \frac{1}{m} \min\{x_{i,j} \mid i \in P\} \sum_{k=1}^{\infty} \mathbf{Pr}\left[* \underset{<k}{\not\rightarrow} j \cap \bigcap_{i \in P} \{i\} \underset{<k}{\not\rightarrow} *\right] \\
&= \frac{x_{P,i}}{m} \sum_{k=1}^{\infty} \mathbf{Pr}\left[* \underset{k}{\not\rightarrow} j \cap \bigcap_{i \in P} \{i\} \underset{k}{\not\rightarrow} * \,\middle|\, * \underset{<k-1}{\not\rightarrow} j \cap \bigcap_{i \in P} \{i\} \underset{<k-1}{\not\rightarrow} *\right] \\
&\quad \cdot \mathbf{Pr}\left[* \underset{<k-1}{\not\rightarrow} j \cap \bigcap_{i \in P} \{i\} \underset{<k-1}{\not\rightarrow} *\right] \\
&= \frac{x_{P,j}}{m} \sum_{k=1}^{\infty} \left(1 - \mathbf{Pr}\left[* \underset{k}{\rightarrow} j \cup \bigcup_{i \in P} \{i\} \underset{k}{\rightarrow} * \,\middle|\, * \underset{<k-1}{\not\rightarrow} j \cap \bigcap_{i \in P} \{i\} \underset{<k-1}{\not\rightarrow} *\right]\right) \\
&\quad \cdot \mathbf{Pr}\left[* \underset{<k-1}{\not\rightarrow} j \cap \bigcap_{i \in P} \{i\} \underset{<k-1}{\not\rightarrow} *\right] \\
&= \frac{x_{P,j}}{m} \sum_{k=1}^{\infty} \left(1 - \frac{1}{n}\left(1 + \sum_{j' \in [m] \setminus \{j\}} \max\{x_{i,j'} \mid i \in P\}\right)\right)
\end{aligned}$$

$$
\begin{aligned}
&\cdot \mathbf{Pr}\left[* \underset{<k-1}{\not\to} j \cap \bigcap_{i \in P}\{i\} \underset{<k-1}{\not\to} *\right] \\
&= \frac{x_{P,j}}{m} \sum_{k=1}^{\infty} \left(1 - \frac{1}{n}\left(1 + \sum_{j' \in [m] \setminus \{j\}} \max\{x_{i,j'} \mid i \in P\}\right)\right)^{k-1} \\
&= \frac{x_{P,j}}{1 + \sum_{j' \in [m] \setminus \{j\}} \max\{x_{i,j'} \mid i \in P\}}.
\end{aligned}
$$

The seventh equality in the above derivation follows by inductively applying to $\mathbf{Pr}[* \underset{<k-1}{\not\to} j \cap \bigcap_{i \in P}\{i\} \underset{<k-1}{\not\to} *]$ the same as we did to $\mathbf{Pr}[* \underset{<k}{\not\to} j \cap \bigcap_{i \in P}\{i\} \underset{<k}{\not\to} *]$ (conform to the steps taken in the third to the sixth equality). □

Lemma 163. *Let x' be the extended $(0,1)$-solution to LP (7.10) that Algorithm 2 has produced at the end of its execution, when given as input a congestion game with non-negative r-externalities $\Gamma = (n, m, v)$. Moreover, let x be the optimal fractional solution to the relaxation of LP (7.10) for Γ. Then, the expected social welfare $\mathbf{E}[U(s(x'))]$ of the extended strategy profile that Algorithm 2 outputs, can be expressed as follows.*

$$
\mathbf{E}[U(s(x'))] = \sum_{\substack{(P,j): j \in [m], \\ P \subseteq [n], \\ |P| \leq r+1}} \frac{x_{P,j} w_{P,j}}{1 + \sum_{j' \in [m] \setminus \{j\}} \max\{x_{i,j'} \mid i \in P\}},
$$

where $w_{P,j} = \sum_{i \in P} v_{i,P \setminus \{i\},j}$, for $P \subseteq [n], |P| \leq r+1$ and $i \in P$.

Proof. By linearity of expectation,

$$
\mathbf{E}[U(s(x'))] = \sum_{\substack{(P,j): j \in [m], \\ P \subseteq [n], \\ |P| \leq r+1}} w_{P,j} \mathbf{E}[x'_{P,j}] = \sum_{\substack{(P,j): j \in [m], \\ P \subseteq [n], \\ |P| \leq r+1}} w_{P,j} \mathbf{Pr}[x'_{P,j} = 1].
$$

The claim then follows from Lemma 162. □

Likewise, we obtain from Lemma 162 exact expressions on the expected utilities of the players.

Corollary 164. *Let x' be the $(0,1)$-solution to LP (7.10) that Algorithm 2 has produced at the end of its execution, when given as input a congestion game with r-externalities $\Gamma = (n, m, v)$. Moreover, let x be the optimal fractional solution to the relaxation* *of LP (7.10) for Γ. Then, the expected utility $\mathbf{E}[u_i(s(x'))]$ of a player $i \in [n]$, of the extended strategy profile that Algorithm 2 outputs, can be expressed as follows.*

$$\mathbf{E}[u_i(s(x'))] = \sum_{j\in[m]} \sum_{P\subseteq[n]\setminus\{i\}:|P|\leq r} \frac{x_{P\cup\{i\},j}v_{i,P,j}}{1+\sum_{j'\in[m]\setminus\{j\}} \max\{x_{j',i'} \mid i' \in P\cup\{i\}\}}.$$

Lastly, we bound the approximation ratio of this rounding procedure. The following important observation follows from the fact that for all $P \subseteq [n], |P| \leq r+1$, it holds that $\sum_{i\in P}\sum_{j\in[m]} x_{P,j} = r$ for any feasible solution x of the relaxation of LP (7.10) of any congestion game with r-externalities Γ.

Proposition 165. *Let $\Gamma = (n, m, v)$ be a congestion game with non-negative r-externalities and let x be a feasible solution to the relaxation of LP (7.10) of Γ. Then,*

$$1 + \sum_{j'\in[m]:j'\neq j} \max\{x_{i,j} \mid i \in P\} \leq r+2$$

for all $j \in [m]$ and $P \subseteq [n], |P| \leq r+1$.

It follows from this observation that:

Corollary 166. *Algorithm 2 is a randomized $(r+2)$-approximation algorithm for computing a social welfare maximizing strategy profile for a congestion game with non-negative r-externalities, that runs in expected polynomial time.*

In the above corollary, the expected polynomial runtime holds because the constraints (7.11) ensure that during any round of the algorithm, there is at least one pair $(i,j), i \in [n], j \in [m]$, such that $x_{i,j} \geq 1/n$. This implies that in each iteration, for all $i \in [n]$, if the variables $x_{i,j}, j \in [m]$ are not rounded to values in $\{0,1\}$ yet, they get rounded with probability at least $1/n^2$. So it takes in expectation not more than a polynomial number of iterations before all of the relevant variables are rounded to $(0,1)$-values.

Next, consider Algorithm 3, which uses Algorithm 2 as a subroutine.

In the theorems that follow, the social welfare function U and utility functions u, have been extended in the obvious way to pairs (s,t), where s is a strategy profile and t a specification of disabled externalities.

Theorem 167. *Let (s,t) be as generated by Algorithm 3, at the end of its execution, when given as input a congestion game with non-negative r-externalities $\Gamma = (n, m, v)$. Moreover, let x be the optimal fractional solution to the relaxation of LP*

Algorithm 3 A polynomial-time mechanism that approximates the maximum social welfare of a congestion game with r-externalities to within a factor of $r+2$.

Input: A description of a congestion game with with non-negative r-externalities $\Gamma = (n, m, v)$.
Output: A strategy profile s for Γ, together with a specification t of which externalities are disabled.
Begin

1. Let x be the optimal fractional solution to the relaxation of LP (7.10) of Γ. Let t be such that no externality is disabled according to t.

2. For all $i \in [n]$, for all $P \subseteq [n] \setminus \{i\}, |P| \leq r$, for all $j \in [m]$, disable externality $v_{i,P,j}$ with probability

$$\frac{1 + \sum_{j' \in [m] \setminus \{j\}} \max\{x_{j',i'} \mid i' \in P \cup \{i\}\}}{r+2},$$

i.e., update t on externality $v_{i,P,j}$ with the above probability.

3. Run Algorithm 2 on Γ. Let s be the extended strategy profile output by Algorithm 2.

4. Let P_u be the set of players that are not allocated to any facility under s. Allocate the players of P_u to an arbitrary facility. Update t such that all externalities of players in P_u are disabled according to t.

5. Output (s, t).
End

(7.10) for Γ. *Then:*

$$\mathbf{E}[u_i(s,t)] = \sum_{\substack{(P,j):j\in[m],\\ P\subseteq[n]\setminus\{i\},\\ |P|\leq r}} \frac{v_{i,P,j}x_{P\cup\{i\},j}}{r+2}$$

for $i \in [n]$*, and*

$$\mathbf{E}[U(s,t)] = \sum_{\substack{(P,j):j\in[m],\\ P\subseteq[n],\\ |P|\leq r+1}} \frac{x_{P,j}w_{P,j}}{r+2} = \frac{xw}{r+2}.$$

Proof. Let x' be the extended feasible $(0,1)$-solution to LP (7.10) that Algorithm 2 has produced at the end of its execution (within the execution of Algorithm 3. We combine Lemma 162 with the fact that for $i \in [n], j \in [m], P \subseteq [n] \setminus \{i\}, |P| \leq r$ each externality v_{i,P,s_i} is disabled under t with probability

$$p_{i,P,j} = \frac{1 + \sum_{j'\in[m]\setminus\{j\}} \max\{x_{j',i'} \mid i' \in P\cup\{i\}\}}{r+2}.$$

By linearity of expectation, the following holds for all $i \in [n]$ just after Algorithm 3 has executed step 3.

$$\begin{aligned}
\mathbf{E}[u_i(s,t)] &= \sum_{\substack{(P,j):j\in[m],\\ P\subseteq[n]\setminus\{i\},\\ |P|\leq r}} v_{i,P,j}\mathbf{E}[x'_{P\cup\{i\},j}]p_{i,P,j} \\
&= \sum_{\substack{(P,j):j\in[m],\\ P\subseteq[n]\setminus\{i\},\\ |P|\leq r}} v_{i,P,j}\mathbf{Pr}[x'_{P\cup\{i\},j} = 1]p_{i,P,j} \\
&= \sum_{\substack{(P,j):j\in[m],\\ P\subseteq[n]\setminus\{i\},\\ |P|\leq r}} \frac{v_{i,P,j}x_{P\cup\{i\},j}}{r+2},
\end{aligned}$$

where the last equality follows from Lemma 162. By taking the sum over the players utilities, we obtain that just after Algorithm 3 has executed step 3,

$$\mathbf{E}[U(s,t)] = \sum_{\substack{(P,j):j\in[m],\\ P\subseteq[n],\\ |P|\leq r+1}} \frac{x_{P,j}w_{P,j}}{r+2} = \frac{xw}{r+2}.$$

In step 4, the social welfare and utilities of the players are clearly not affected. This proves the claim. □

Due to Theorem 167, we arrive at our desired mechanism.

Theorem 168. *There exists a randomized mechanism that approximates within a factor* $(r+2)$ *the maximum social welfare for a congestion game with non-negative* r*-externalities* $\Gamma = (n, m, v)$*, for the extended range of solutions* (s, t)*, where* $s \in \Sigma$ *is a strategy profile of* Γ *and* t *is a specification of disabled externalities. This mechanism runs in expected polynomial time, and satisfies truthfulness in expectation, universal individual rationality, and universally non-negative payments.*

Proof. Algorithm 3 has the property that it is *maximum-in-distributional-range* (MIDR) with respect to U and the set of possible valuations that may be reported by the players. The MIDR property means that among the set of output distributions that the algorithm may produce (when considering the set of all possible valuations that the players may report), it always selects the distribution that maximizes the social welfare U in expectation. This fact follows from Theorem 167: on input $\Gamma = (n, m, v)$ the expected social welfare of the allocation that is output, is the social welfare of the fractional feasible solution x of the relaxation of LP (7.10) of Γ, scaled down by a constant factor of $r+2$. Because x is the fractional solution that attains the maximum social welfare among all fractional solutions, it follows that Algorithm 3 is MIDR. It is well-known [Dobzinski and Dughmi, 2009] that for MIDR algorithms, there exists a payment rule that turns the algorithm into a mechanism that satisfies individual rationality in expectation, universally non-negative payments, and truthfulness in expectation. The payment rule used for this is a well known one, called *Clarke's pivot rule* (see e.g. Nisan et al. [2007]): it charges a player $i \in [n]$ the value $\sum_{i' \in [n] \setminus \{i\}} \mathbf{E}[u_{i'}(s', t')] - \sum_{i' \in [n] \setminus \{i\}} \mathbf{E}[u_{i'}(s, t)]$, where (s', t') is the output of Algorithm 3 if player i would not be present.

We can compute $\sum_{i' \in [n] \setminus \{i\}} \mathbf{E}[u_{i'}(s, t)]$ in polynomial time through Theorem 167. We can compute $\sum_{i' \in [n] \setminus \{i\}} \mathbf{E}[u_{i'}(s', t')]$ in polynomial time by Theorem 167 in combination with solving the relaxation of LP (7.10) for the game Γ' obtained from Γ by setting all externalities of player i to zero. Therefore we can compute the appropriate payments in polynomial time.

The resulting mechanism, which always charges the Clarke payments described above irrespective of which randomized steps are taken, is clearly individually rational in expectation, because the payment charged to a player is less than that player's expected utility. This mechanism is not universally individually rational however, because the rounding procedure of the mechanism may generate outcomes for which a player's utility is lower than the payment charged. We use the following trick, also used e.g. by Lavi and Swamy [2005], that allows us to convert the property of individual rationality in expectation into universal individual rationality: Let p_i denote the Clarke payment just described. If a player i's utility for the outcome (s', t') generated by the mechanism is 0 (according to the vector of valuations it has reported), then the mechanism charges the player a payment of 0. Otherwise, the mechanism charges the player a payment of $u_i(s', t')p_i/\mathbf{E}[u_i(s, t)]$. This payment rule is universally individually rational because $p_i/\mathbf{E}[u_i(s, t)] \leq 1$. Moreover, note that the expected payment that is charged by the mechanism is exactly p_i, so that truthfulness in expectation is preserved. □

7.11 Future Work

The most pressing question that remains to be answered, is whether a mechanism exists that has the same properties as the one we presented in Section 7.10, but does not require the extended solution range (in which we are allowed to disable the externalities of the players arbitrarily).

A less ambitious goal would be to see if it is possible to obtain such results when weakening some of the other requirements we impose on the mechanism. Can we for example attain a polynomial time constant approximation mechanism if we implement truthfulness in ex-post equilibria instead of dominant strategies?

Other interesting research directions include the study of the same type of questions, but for *non-positive* externalities, instead of non-negative ones. Secondly, it might make sense to try to acquire some insights into the equilibria of the congestion games studied here. Do they exist, and how hard is it to compute one if they do? Moreover, what is the price of anarchy of these games? We have not considered these topics yet and want to address this in future research.

Chapter 8

Housing Markets with Indifferences: a Tale of Two Mechanisms*

We consider in this chapter mechanism design in *(Shapley-Scarf) housing markets*. We aim in this setting for mechanisms that are stable in the sense that they are resistant against cooperative behavior of the players. This property is referred to as *core stability* in the setting of housing markets. The core stability requirement results in yet another flavor of mechanism design, when we compare this to the mechanism design problems of the previous two chapters. We advise the reader to be familiar with the material of Sections 1.3 up to Section 1.3.1, Section 1.3.1.11, and optionally the definition of the *core* in Section 1.3.2.

In a housing market, there is a set of players and a set of houses. Each player owns a single house initially, and each player has certain preferences over the houses. The goal is to reallocate the houses to the players in a mutually beneficial and stable manner.

Definition 169 ((Shapley-Scarf) housing market, allocation)**.** A *(Shapley-Scarf) housing market* is a triple $H = (n, \omega, \preceq)$, where $\preceq = (\preceq_1, \ldots, \preceq_n)$ and $\omega : [n] \to [n]$. For $i \in [n]$, $\preceq_i$ is a complete transitive relation on $[n]$. This is interpreted as there being a set of players $[n]$ and a set of equally many houses $[n]$. An *allocation* for H is a bijection from $[n]$ to $[n]$. Note that ω is an allocation for H. The function ω describes the

*The contents of this chapter have been published as Aziz and De Keijzer [2012] *initial endowment*, i.e., $\omega(i)$ is the house owned by player $i \in [n]$. Lastly, $\preceq_i$ describes the *preferences* of $i \in [n]$, where player i prefers house $j \in [n]$ to house $j' \in [n]$ iff $j' \preceq j$.

Definition 170 (Housing market mechanism)**.** A *housing market mechanism* M is a mapping from housing markets $(n, \omega, \preceq)$ to allocations $f : [n] \to [n]$ of the agents to the houses.

Allocations generated by a housing market mechanism may satisfy certain properties.

Definition 171 ((Strict) core stable, Pareto optimal, individually rational)**.** Let $H = (n, \omega, \preceq)$ be a housing market, and let $f : [n] \to [n]$ be an allocation for H.

- f is said to be *individually rational* under H iff $f(i) \succeq_i \omega(i)$ for $i \in [n]$. Informally: no player $i \in [n]$ is strictly better off with its initial house $\omega(i)$. A housing market mechanism is said to be individually rational iff it always outputs an allocation that is individually rational under its input input housing market.
- A set of players $S \subseteq [n]$ is said to *block* f under H iff there exists an allocation y such that $f(i) = g(i)$ for all $i \in [n] \setminus S$ and $g(i) \succ_i f(i)$ for all $i \in S$. Informally: the players in S can trade their allocated houses among each other such that each player in S strictly prefers its new house to its allocated house.
- A set of players $S \subseteq [n]$ is said to *weakly block* f under H iff there exists an allocation y such that $f(i) = g(i)$ for all $i \in [n] \setminus S$ and $g(i) \succeq_i f(i)$ for all $i \in S$, and moreover there exists a player $i \in S$ such that $g(i) \succ_i f(i)$. Informally: the players in S can trade their allocated houses among each other such that each player in S obtains a house that is at least as good as its allocated house, and at least one player in S strictly prefers its new house to its allocated house.
- f is said to be *core stable* under H iff no coalition $S \subseteq [n]$ blocks f. A housing market mechanism is said to be *core selecting* iff it always outputs an allocation that is core stable under its input housing market, if such an allocation exists.
- f is said to be *strict core stable* under H iff no coalition $S \subseteq [n]$ weakly blocks f. A housing market mechanism is said to be *strict core selecting* under H iff it always outputs an allocation that is strict core stable under its input housing market, if such an allocation exists.

- f is said to be *Pareto optimal* under H iff $[n]$ does not block f. A housing market mechanism is said to be *Pareto optimal* iff it always outputs an allocation that is Pareto optimal under its input housing market.

A housing market mechanism M induces a game for each housing market $(n, \omega, \preceq)$ with player set $[n]$, where the strategy set of each player is the set of preference profiles on the house set. Given a strategy profile s of this game, a player achieves higher utility if $M(n, \omega, s)$ gives the player a more preferred house. We aim in this chapter toward housing market mechanisms that always output core stable, strict core stable, Pareto optimal, and individually rational allocations, even when the players have the ability to strategize by misreporting their preferences in this sense. In order to achieve this, we need our mechanism to satisfy *truthfulness*, which means the following in the context of housing markets.

Definition 172 (Truthfulness)**.** A housing market mechanism M is *truthful* iff for every housing market $H = (n, \omega, \preceq)$, every vector of preference profiles $s = (s_1, \ldots, s_n)$, and every player $i \in [n]$, it holds that $f(i) \succeq_i g(i)$, where $f = M(n, \omega, (\preceq_i, s_{-i}))$ and $g = M(n, \omega, s)$.

A final requirement that we impose is that our housing market mechanism is polynomial time computable.

The case where the preferences of a player are restricted to strict total orders has been examined and resolved by Shapley and Scarf [1974], where it is shown that a simple and elegant mechanism called *Gale's Top Trading Cycle (TTC)* mechanism is truthful and core selecting.

Alcalde-Unzu and Molis [2011] and Jaramillo and Manjunath [2011] independently examined the general case in which players can express indifferences among houses. They proposed two important families of mechanisms, known as *TTAS* and *TCR* respectively. We formulate in this chapter a family of mechanisms which not only includes TTAS and TCR, but also satisfies many desirable properties of both families. As a corollary, we show that TCR is strict core selecting. Finally, we settle an open question regarding the computational complexity of the TTAS mechanism. Our study also raises a number of interesting research questions.

8.1 Background

Housing markets are fundamental models of exchange economies of goods where the goods could range from dormitories to kidneys [Sönmez and Ünver, 2011].

The classic housing market (also called the Shapley-Scarf Market) consists of a set of players each of which owns a house and has strict preferences over the set of all houses. The goal is to redistribute the houses to the players in the most desirable fashion. Shapley and Scarf [1974] showed that a simple yet elegant mechanism called *Gale's Top Trading Cycle (TTC)* is truthful and finds an allocation which is in the core. TTC is based on multi-way exchanges of houses between players. Since the basic assumption in the model is that players have strict preferences over houses, TTC is also strict core selecting and therefore Pareto optimal.

Indifferences in preferences are not only a natural relaxation but are also a practical reality in many cases. Many new challenges arise in the presence of indifferences: core stability does not imply Pareto optimality; the strict core can be empty [Quint and Wako, 2004]; and issues related to truthfulness need to be re-examined. In spite of these challenges, Alcalde-Unzu and Molis [2011] and Jaramillo and Manjunath [2011] proposed mechanisms for housing markets with indifferences, that satisfy many nice properties. Alcalde-Unzu and Molis [2011] presented the *Top Trading Absorbing Sets (TTAS)* family of mechanisms which are truthful, core selecting (and therefore individually rational), Pareto optimal, and strict core selecting. Independently, Jaramillo and Manjunath [2011] came up with a different family of mechanisms called *Top Cycle Rules (TCR)* which are truthful, core selecting, and Pareto optimal. Whereas it was shown in [Jaramillo and Manjunath, 2011] that each TCR mechanism runs in polynomial time, the time complexity of TTAS was raised as an open problem in [Alcalde-Unzu and Molis, 2011].

8.2 Contributions and Outline

We first highlight the commonality of TCR and TTAS by describing a simple class of mechanisms called *Generalized Absorbing Top Trading Cycle (GATTC)* in Section 8.4, which encapsulates the TTAS and TCR families. It is proved that each GATTC mechanism is core selecting, strict core selecting, and Pareto optimal.

The TCR and TTAS mechanisms are described in Section 8.5, where it is also proved that TCR and TTAS are special cases of GATTC. As a corollary, TCR is strict core selecting. We note that whereas a GATTC mechanism satisfies a number of desirable properties, the truthfulness of a particular GATTC mechanism hinges critically on the order and way of choosing trading cycles.

Finally, we settle in Section 8.6 the computational complexity of TTAS. By simulating a binary counter, it is shown that a TTAS mechanism can take exponential time to terminate. This answers an open question raised by Alcalde-Unzu and Molis [2011].

8.3 Preliminaries

Desirable allocations of housing markets can be computed via a graph-theoretic approach.

Definition 173 (G(H))**.** For a housing market $H = (n, \omega, \preceq)$ we denote by $G(H)$ the directed graph $G(H) = (V, A)$ where $V = \{v_1, \ldots, v_n\} \cup \{h_1, \ldots, h_n\}$ and $A = \{(v_i, h_j) \mid i, j \in [n], \forall j' \in [n] : j \succeq_i j'\} \cup \{(h_j, v_i) \mid i, j \in [n], \omega(i) = j\}$. In words, there is a set of vertices corresponding to the player set and there is a set of vertices corresponding to the house set. Each player vertex points to its set of most preferred houses, and each house vertex points to its initial owner, conform ω.

For convenience, when discussing the graph $G(H)$ of a housing market H we will make light of the difference between *players* and *vertices corresponding to players*, i.e., we often refer to the vertices $\{v_1, \ldots, v_n\}$ as simply players. Likewise, we refer to $\{h_1, \ldots, h_n\}$ as houses. We say that a player $i \in [n]$ *points to* a house $j \in [n]$ in $G(H)$ iff (v_i, v_j) is an arc of $G(H)$. Analogously, we may say that houses *point to* players.

We introduce next some essential graph-theoretic notions.

Definition 174 (Absorbing set, symmetric pair, paired symmetric)**.** Let $G = (V, A)$ be a directed graph. An *absorbing set* $S \subseteq V$ of G is a strongly connected component for which there exists no arc in $A \cap (S \times (V \setminus S))$, i.e., there are no outgoing arcs. Two vertices $v, v' \in V$ constitute a *symmetric pair* of G iff $(v, v'), (v', v) \in A$, i.e., there are arcs between the two vertices going in both directions. An absorbing set S of G is *paired symmetric* if each vertex in S belongs to a symmetric pair of G.

8.4 GATTC

In this section, we formulate a simple family of housing market mechanisms called *Generalized Absorbing Top Trading Cycle (GATTC)* which is designed for housing markets with indifferences and extends not only TTC, but also includes the two families TTAS and TCR. It is based on multi-way exchanges of houses between players. We will show that GATTC satisfies many desirable properties of housing market mechanisms, such as being core selecting and Pareto optimal.

Before we describe GATTC, we will introduce the original TTC mechanism which is defined for the domain of housing markets with strict preferences, i.e., in which all players preferences are strict total orders. TTC works as follows. For a housing market H with strict preferences, we first construct the corresponding graph $G(H)$ as defined above. Then, we start from a player and walk arbitrarily along the arcs until a cycle is completed. A cycle is of course guaranteed to exist, as each vertex of $G(H)$ has positive out-degree. This cycle is removed from $G(H)$. Within the removed cycle C, each player (corresponding to a vertex in C) gets allocated the house (corresponding to the vertex in C) that it was pointing to. The graph $G(H)$ is then *adjusted* so that the remaining players point to their most preferred houses *among* the remaining houses in $G(H)$. The process is repeated until all the houses and players are deleted from the graph.[1]

For a housing market with indifferences, TTC can still be used to return a core selecting allocation: break ties of the preference relations of the players arbitrarily (turning the housing market into a housing market with strict preferences), and then run TTC. However such an allocation may not be Pareto optimal [see *e.g.*, Alcalde-Unzu and Molis, 2011, Jaramillo and Manjunath, 2011]. GATTC achieves Pareto optimality and is based on absorbing sets and the concept of a 'good cycle'.

Definition 175 (Good cycle, implementing a cycle)**.** Let $H = (n, \omega, \preceq)$ be a housing market. Let $G = (V, A)$ be any directed graph with vertex set $V = V_1 \cup V_2, V_1 \subseteq \{v_1, \dots v_n\}, V_2 \subseteq \{h_1, \dots, h_n\}$ with only arcs between V_1 and V_2. I.e., V is a graph on a subset of the houses and players of H. A *good cycle* is any cycle of G that contains at least one vertex that is not paired symmetric. Let C be a cycle of G. By *implementing* C we mean changing the graph G by removing all edges that point from V_2 to V_1 and introducing edges $\{(h_j, v_i) \mid (v_i, h_j) \in C\}$.

Algorithm 4 defines our class of GATTC mechanisms, by providing an algorithmic description of how an outcome of such a mechanism may be generated.

We stress that the choices that a GATTC mechanism makes in Steps 1.1 and 1.2 are allowed to be different each time the algorithm reaches these steps, within the same execution. The same holds for the number of times that Steps 1.1 and 1.2 are repeated, each time that Step 1 is executed.

Example 176. Consider a housing market $H = (n, \omega, \preceq)$ where ω is such that $\omega(i) = i$ for all $i \in [n]$, and preferences $\preceq$ are defined as follows.

player	1	2	3	4	5
preferences	2	3	$4, 5$	1	2
	1	2	3	5	4
				4	5

This table should be interpreted as that a house is more preferred by a player $i \in [n]$ iff it occurs on a higher row in column i of the table. If two houses j, j' occur on the same row in column i, it means that player i is indifferent between house j and j'.

[1]See Section 2.2 of [Sönmez and Ünver, 2011] for an elegant illustration of how TTC works.

Algorithm 4 A blueprint for mechanisms falling in the GATTC class. The concept of *adjusting a graph* is defined here in the same way as for the TTC mechanism.

Input: A housing market H.
Output: An allocation f for H.
Begin

Let $G = G(H)$. Let $f = \varnothing$. (At the end of execution, f will indeed be a function from $[n]$ to $[n]$. In this description of GATTC, we view f as a set of pairs, and we initialize f as the empty set.) Repeat the following until G is empty:

1. Repeat the following a finite number of times on G:

1.1. Either implement a non-good cycle (if G is not empty), or do nothing.

1.2. Either remove a paired symmetric absorbing set S of G (setting f to $f \cup \{(i, j) \mid i, j \in [n], v_i$ and h_j are symmetric pairs of the subgraph of G induced by $S\}$), and adjust G, or do nothing.

2. Repeatedly remove paired symmetric absorbing sets S of G from G (setting f to $f \cup \{(i, j) \mid i, j \in [n], v_i$ and h_j are symmetric pairs of the subgraph of G induced by $S\}$), and adjust G, until there are no paired symmetric absorbing sets left in G.

3. If G is not empty, implement a good cycle.

End

Given these preferences, if ties (i.e., indifferences in the preferences of the players) are broken in any way, TTC does not return a Pareto optimal allocation. However, GATTC mechanisms (including TTAS and TCR) will only return the following Pareto optimal allocations: $\{(1,2),(2,3),(3,5),(4,1),(5,4)\}$ or $\{(1,1),(2,3),(3,4),(4,5),(5,2)\}$. Figure 2 illustrates the first steps in the execution of a GATTC mechanism on this housing market.

We say that an algorithm *induces a* housing market mechanism iff it terminates and returns an allocation.

Theorem 177. *An algorithm in the class of GATTC algorithms induces a core selecting, and Pareto optimal housing market mechanism.*

Proof. We prove each property separately. We start by proving that an algorithm in the class of GATTC algorithms induces a housing market mechanism, i.e., it terminates and returns an allocation for its input housing market.

Terminates and returns allocation: From its description (Algorithm 4), due to the definition of implementation of a cycle, the out-degree of a house in G is always 1 at every step of the algoritm. So at the beginning of every step, G has the property that each vertex has positive out-degree. For non-empty graphs with this property, an absorbing set of cardinality greater than 1 is guaranteed to exist [Kalai and Schmeidler, 1977]. Therefore, if G is not empty, then at Step 1.1 there is guaranteed to be a cycle, and at Step 3 there is guaranteed to be a good cycle (because there is an absorbing set that is not paired symmetric). In each iteration (consisting of Steps 1, 2, and 3), if paired symmetric absorbing sets exist they are removed in Step 2.[2] Also, at least one good cycle is implemented in Step 3, which reduces the number of vertices that are not paired symmetric. Therefore, there is a maximum of $O(n)$ iterations until GATTC terminates. Since each removed house is allocated to the player it was last pointing to, GATTC returns an allocation for its input housing market.

Core selecting: When a player $i \in [n]$ is removed from the graph along with its allocated house $h \in [n]$, then h is among player i's set of most preferred houses, among those houses that are still in the graph. Therefore i cannot be in a blocking set S of players, when S consists only of players remaining in the graph. Thus, for each set S there is a player $i \in S$ for which there is no house among the houses allocated to S that i prefers more than the house allocated to i, so S is not a blocking coalitions.

[2] An absorbing set of a graph can be computed in linear time via the algorithm of Tarjan [1972].

Pareto optimal: Let S_k be the kth paired symmetric absorbing set that arises at some point in the GATTC mechanism (and is thus removed from the graph by the GATTC mechanism, and is included accordingly in the allocation produced by the GATTC mechanism). In any allocation x in which none of the players in S_1 are worse off than in the allocation produced by GATTC, the players in S_1 must be allocated to houses in S_1. Taking this as the base case, it follows by easy induction that in x, the players of S_k must be allocated to houses in the kth paired symmetric absorbing set. Next, suppose that i is a player in S_k for some k. Then no house in S_k is more preferred by i than the house that the GATTC mechanism assigns it to. It follows that no player is strictly better off in x than in the allocation produced by GATTC.

This completes the proof. □

Theorem 178. *GATTC is strict core selecting.*

Proof. We prove the statement by proving two claims.
Claim 179. *For any GATTC mechanism it holds that if at any point during the execution, there is an absorbing set A and an allocation in which each player in A gets allocated its most preferred house among the houses in A, then the GATTC mechanism will output an allocation in which every player in A gets allocated its most preferred house among the houses in A.*

Proof. Define an *inward set* as a set of vertices without arcs pointing outward from A. An absorbing set is by definition an inward set. We prove this claim for the more general notion of inward sets. Let A be an inward set that arises at some time point t during the execution of the GATTC mechanism, and assume that there exists an allocation in which each player in A gets allocated its most preferred house among the houses in A. If A eventually becomes paired symmetric, then every player in A surely gets a maximal house within A. Let us thus assume that A does not eventually become paired symmetric. Consider the first point in time t' where vertices are removed from A by the mechanism. This point t' exists because the mechanism terminates. All cycles that are implemented in between t and t' either lie completely inside A or completely outside A, because there are no arcs pointing from a vertex outside A to a vertex inside A. It follows that at point t', the removed paired symmetric absorbing set A' is a strict subset of A. Note that players in $A \setminus A'$ cannot get a house from within A' without some player in A' getting a house that it prefers strictly less. Hence, by the assumption that an allocation exists in which each player in A gets its most preferred house within A, it follows that players in $A \setminus A'$ can still all get a maximal house from within $A \setminus A'$.

The proof follows by induction; repeating the same argument on the inward set $A \setminus A'$ that arises when removing A' from the graph.

□

Claim 180. *The returned allocation f is in the strict core if and only if for each absorbing set A encountered in the algorithm, each player in A gets allocated a most preferred house in A.*

Proof. ($\Rightarrow$) Assume there is a player i that is in A, for which there exists a house h in A for which $h \succ_i f(i)$. Then any set of players in a cycle within A that contains player i, is a weakly blocking coalition. This contradicts that f is in the strict core.

($\Leftarrow$) Assume that each player i in A gets a most preferred house among the houses within A. Then i cannot be part of a blocking coalition. It can still be part of a weakly blocking coalition if a player i in A has a most preferred house h among the houses in the remaining graph, where h is outside A, and there exists a cycle of the form $(i, h, \dots, i)$. But this is not possible since A is absorbing. □

From the two claims, the theorem follows.

□

We also observe that on the domain of strict preferences, GATTC is equivalent to TTC. The reason is that implementation of any cycle results in a paired symmetric absorbing set which is then removed from the graph. Ma [1994] proved that for housing markets with strict preferences, a mechanism is strict core selecting if and only if it is individually rational, Pareto optimal, and truthful.

Theorem 181. *Not every GATTC mechanism is truthful.*

Proof sketch. Consider the following GATTC mechanism in which no non-good cycle is implemented, and every good cycle is found in the following way. Consider player $i \in [n]$, and house $j \in [n]$ for which $(v_i, h_j) \in E$, $(h_j, v_i) \notin E$, and v_i and h_j are in a strongly connected component of $G(H)$. Then, there exists a shortest path P from h_j to v_i. Find this path P by Dijkstra's shortest path algorithm. Path P gives us a good cycle (v_i, h_j, P, v_i).

For this subclass of GATTC, it can be shown that a player may have an incentive to lie about its preferences to obtain a better allocation. Informally, there exist instances of a housing market in which if a player $i \in [n]$ does not misreport its preferences, it may only get a third most preferred house. However, if i points to its second most preferred house h in the graph, it can manage to influence which good cycle is selected by the mechanism, and be included in that good cycle. Player i then gets allocated house j. □

8.5 TTAS and TCR

We now describe the two families of mechanisms in the literature — TTAS [Alcalde-Unzu and Molis, 2011] and TCR [Jaramillo and Manjunath, 2011] — designed for housing markets with indifferences. Both families of mechanisms are extensions of TTC. We will later show that both families are subclasses of GATTC.

TTAS

> *Let $H = (n, \omega, \preceq)$ be the input housing market. Fix a priority ranking of the houses; i.e., a complete, transitive and antisymmetric binary relation over $[n]$. Construct the graph $G(H)$, and run the following procedure on it (starting with $i = 1$, incrementing i every iteration) until no more players are remaining in the graph.*

Step i.

- *(i.1) Let each player remaining in the graph point to its most preferred set of houses among the houses remaining in the graph. Select the absorbing sets of this directed graph.*
- *(i.2) Consider the paired symmetric absorbing sets. Their players are allocated to the house that the players currently point to in the graph. These absorbing sets are removed from the graph.*
- *(i.3) Consider the remaining absorbing sets. Select for each player i a unique house to point to by using the following criterion: Let h be the house that points to player i in the graph that is remaining. Player i points only to the house that it prefers most among the houses remaining. Ties are broken by selecting among the candidate houses the one that comes after h in the priority order (if there is no such house, then select among the candidate houses the first house in the priority order).*
- *(i.4) Then, in this subgraph, there is necessarily at least one cycle and no two cycles intersect. Implement all cycles in this subgraph, but do not remove them from the graph.*

The algorithm terminates when no players and houses remain, and the outcome is the assignment formed during its execution.

TCR

Consider a priority ranking of the players; i.e., a complete, transitive and anti-symmetric binary relation over $[n]$. Construct the graph $G(H)$ and repeat the following three steps until no more players are left.

Departure: *A set of players S is chosen to "depart" if for each player $i \in S$ it holds that one of i's most preferred houses (among the houses remaining in the graph) points to i, and secondly, any house in the set of most preferred houses by S among the houses remaining in the graph, points to a player in S.*

Once a set S of players departs, each player in S is allocated the house pointing to it. The set S and the houses pointing to this set are removed from the graph.

There may be other sets of players that can depart. The process continues until there are no more sets that can depart. If the two conditions are not met by any set, then no set departs.

Pointing: *Each player in the graph is made to point to a particular player holding one of its most preferred houses (among the houses remaining in the graph), which is done as follows.*

Stage 1. *For each remaining player $i' \in [n]$ where i' is pointed at by the same house as in the previous iteration, each player $i \in [n]$ that pointed at i' in the previous step points to i' in the current step. This step does not apply in case it is the very first step.*

Stage 2. *Each player $i \in [n]$ with a unique most preferred house j (among the houses remaining in the graph) is made to point to the player i' for which it holds that j points to i'.*

Stage 3. *A player $i \in [n]$ is called* unsatisfied *iff none of its most preferred houses among the houses remaining in the graph points toward i. Each player i for which at least one of its most preferred houses (among the houses remaining in the graph) points to an unsatisfied player, is made to point to such an unsatisfied player that is pointed at by one of i's most preferred houses (among the houses remaining in the graph). If there are multiple players that i can point to, ties are broken by taking the player that appears first in the priority order that we fixed beforehand.*

Stage 4. *Each player $i \in [n]$ for which at least one of its most preferred houses (among the houses remaining in the graph) points to a satisfied player that points to an unsatisfied player, is made to point to such a satisfied player.*

Ties are broken by making i point to the satisfied player that points to the unsatisfied player that appears first in the priority order. If two or more such satisfied players point to the unsatisfied player with highest priority, ties are broken by making the player player point to the satisfied player with highest priority.

Stage … *And so on.*

Trading: *Since each player remaining in the graph points to another player, there is at least one cycle of players. For each cycle C of players in the graph, each house pointing to a player i in C is repointed to the player pointing to i.*

Note that TTAS and TCR mechanisms depend on the priority ordering over the houses and players respectively, and different priority rankings lead to different mechanisms. Thus, TTAS and TCR are *classes* of mechanisms rather than single mechanisms. Next, we show that TTAS and TCR are subclasses of GATTC, in which cycles are selected via the priority order over houses and players respectively.

Theorem 182. *GATTC generalizes both the TTAS and TCR families of mechanisms.*

Proof. (*GATTC generalizes TTAS*). Step i.2 of TTAS corresponds to repeatedly executing Step 1.2 (and skipping Step 1.1). After that, TTAS may implement a number of non-good cycles. This corresponds in GATTC to executing Step 1.1 (skipping Step 1.2). However, the proof of Proposition 1 in [Alcalde-Unzu and Molis, 2011] shows that TTAS can never perpetually implement non-good cycles: Either the graph becomes empty, or eventually a good cycle is found and implemented. So executing in TTAS Step i.2 to i.4 on iterations where a good cycle is implemented, corresponds to executing Steps 3 and 4 of GATTC.

(*GATTC generalizes TCR*). A TCR rule reduces to the GATTC mechanism if zero non-good cycles are implemented in Step 1 and if in Step 3 of GATTC, a good cycle is implemented in the particular way as outlined in the definition as TCR.[3] It is clear from Step 2 (pointing) of TCR that because of the way players are made to point, the cycles implemented in Step 3 involve at least one vertex that is not paired symmetric. Therefore the cycles implemented are good cycles.

□

In contrast to TTAS (which is known to be strict core selecting), it was not known whether TCR is also strict core selecting. As a corollary of Theorems 178 and 182, we obtain the following.

[3] Of course, in the description of TCR, players are made to point to other players rather than houses, but TCR is trivially rephrased so that in the trading step players are made to point to the houses instead of their owners, leading to the trading step being equivalent to the implementation of a good cycle.

Corollary 183. *Each TCR mechanism is strict core selecting.*

In the next section, we answer an open question concerning the running time of the TTAS mechanism.

8.6 Complexity of TTAS

An important property of TTAS is that if a player $i \in [n]$ is repointed at by a house $j \in [n]$ during the running of TTAS, but i and j are not yet deleted from the graph, then player i is guaranteed to be ultimately allocated a house $j' \in [n]$ where $j \sim_i j'$ [Lemma 1, Alcalde-Unzu and Molis, 2011]. Therefore, the number of symmetric pairs can only increase during the running of the algorithm although it may stay constant in a number of iterations. Alcalde-Unzu and Molis [2011] showed that despite a number of stages in which no obvious progress is being made, TTAS eventually terminates [Proposition 1, Alcalde-Unzu and Molis, 2011]. Although, we know that TTAS terminates and results in a proper allocation, the proof of [Proposition 1, Alcalde-Unzu and Molis, 2011] does not give insight into how many steps are taken before TTAS finishes.We will show the following.

Theorem 184. *There exists a family of housing markets $\{H_k = (n_k, \omega_k, \preceq^k) \mid k \in \mathbb{N}_{>0}\}$ with $n_k = 2k+1$, and corresponding priority rankings $\{R_k \mid k \in \mathbb{N}_{>0}\}$ such that if the TTAS mechanism receives input H_k and chooses R_k as its priority ranking in Step 0, then the TTAS mechanism runs for at least $2^k = 2^{(n_k-1)/2}$ steps until it terminates.*

This theorem shows thus that the TTAS mechanism, according to its current description, does not run in polynomial time. It still might be that for each instance, there is some priority ranking such that the TTAS mechanism runs in polynomial time, but then at least some additional details are needed in the description on how to choose the priority ranking. I.e., the algorithm described in Alcalde-Unzu and Molis [2011] is not sufficient to attain a polynomial running time.

Proof. For convenience, we name the houses of housing market H_k as $\{h_1, h'_1, h_2, h'_2, \ldots, h_k, h'_k, h_{k+1}\}$ and we name the players $\{v_1, v'_1, v_2, v'_2, \ldots, v_k, v'_k, v_{k+1}\}$. In the initial endowment ω_k, house h_j is assigned to player v_j for all $j \in [k+1]$, and house h'_j is assigned to player v'_j for all $j \in [k]$. The preference profile of player v_j, $j \in [k]$ is described by two equivalence classes: its class of most preferred houses is $\{h'_j, h_j, h_{j+1}\}$, and the remainder of the houses is in its other equivalence class, i.e., its class of least preferred houses. The preference profile of player v'_j, $j \in [k]$, is also described by two equivalence classes: its class of most preferred houses is $\{h_j, h'_j, h_1\}$ (so for $j = 1$, this set has cardinality 2), and the remainder of the houses are in the other equivalence class, i.e., its class of least preferred houses. The preference profile of player v_{k+1} is also described by two equivalence classes: Its class of most preferred houses is $\{h_1\}$, and the remainder of the houses is in its other equivalence class, i.e., its class of least preferred houses. The priority ranking R is $(h_1, h'_1, h_2, h'_2, \ldots, h_k, h'_k, h_{k+1})$.

The high level idea of this example is to simulate a binary counter. The graph that the TTAS mechanism maintains will contain a single absorbing set at every step: the entire graph. At every step except the last one, the only player that prevents the graph from being paired symmetric will be player v_{k+1}. We associate bit-strings of length k to the graphs that may arise in some of the steps of the TTAS algorithm: Let $b \in \{0,1\}^k$ be any bit-string of length k, then we define the graph G_b as the graph where for all j,

- v_j and v'_j all point to their set of most preferred houses,
- if $b_j = 0$, then h_j points to v_j and h'_j points to a'_j.
- if $b_j = 1$, then h_j points to v'_j and h'_j points to a_j.

We prove that for all bit-strings b of length k there is a step i_b such that the graph at the beginning of step i_b is equal to G_b. Because there are 2^k possible bit-strings, it then follows that there are at least 2^k steps before the algorithm terminates.

In order to understand what happens during the execution of the TTAS algorithm on an instance M_j, we recommend the reader to inspect the example of Figure 8.2, where the graph at the beginning of every step is shown when we run the TTAS mechanism on M_3.

Let us assume that at the beginning of step i of the execution of the TTAS mechanism, the graph is equal to G_b for some b. We can prove that G_b is strongly connected:

Claim 185. *For each length k bit-string b, G_b is strongly connected.*

Proof. We first show that there is a path from h_1 to every other vertex v.

If $b_1 = 0$, then h_1 points to v_1 and h'_1 points to a'_1. If $b_2 = 0$, then there exists a path $(h_1, v_1, h_2, v_2, h'_2, v'_2)$. If $b_2 = 1$, then there exists a path $(h_1, v_1, h_2, v_2, h'_2, v_2)$.

If $b_1 = 1$, then h_1 points to v'_1 and h'_1 points to v_1. If $b_2 = 0$, then there exists a path $(h_1, v'_1, h'_1, v_1, h_2, v_2, h'_2, v'_2)$. If $b_2 = 1$, then there exists a path $(h_1, v'_1, h'_1, v_1, h_2, v'_2, h'_2, v_2)$.

Therefore h_1 has a path to the following vertices: $v_1, v_2, h_1, h_2, v'_1, v'_2, h'_1, h'_2$.

Using the same argument, we can see that for each v_j, there is a path to v_{j+1}; for each v'_j, there is a path to v'_{j+1}; for each h_j there is a path to h_{j+1}; for each h'_j, there is a path to h'_{j+1}. Therefore, it holds that: From h_1, there is a path to each v_j for

$j \in [k+1]$; From h_1, there is a path to each v'_j for $j \in [k]$; From h_1, there is a path to each h_j for $j \in [k+1]$; and from h_1, there is a path to each h'_j for $j \in [k]$.

Similarly, it can be shown that from every vertex, there is a path to h_1. This completes the argument of the claim.

□

Therefore, G_b has only one absorbing set: the whole of G_b.

Also observe that for all b, G_b is not paired symmetric, because of player $k+1$. From this we conclude that if the graph at the beginning of a step i is equal to G_b, for some $b \in \{0,1\}^k$, then the TTAS mechanism does not terminate at step i, and the mechanism will certainly reach step $i+1$.

For some step i of the TTAS mechanism, and for every player $v \in [n]$, let S_v^i denote the set of most preferred houses of v that are ranked lower than the house assigned to v in step i. However, if this set is empty, then define S_v^i to be the set of most preferred houses of v. Let us assume that for step i, the following property holds, which we will call *Property A_i*: for every player $v \in [n]$, it holds that the set of most preferred houses of v that have been repointed to v the least number of times (including 0 times), is S_v^i.

We define a straightforward bijection $c : \{0,1\}^k \to [2^k - 1] \cup \{0\}$ as follows: bit-string b corresponds to the integer $\sum_{j=1}^{k} 2^{j-1} b_j$. We then see that the following happens:

Claim 186. *Let b be a bit-string of length k, suppose that i is a step in the TTAS mechanism such that the graph at step i is equal to G_b, and suppose that Property A_i holds.*

- *If $c(b)$ is even, then the graph at step $i+1$ of the TTAS algorithm is equal to G_{b+1}, and Property A_{i+1} holds.*
- *If $c(b)$ is odd and not equal to $2^k - 1$, then the graph at step $i+2$ of the TTAS algorithm is equal to G_{b+1}, and Property A_{i+2} holds.*

Proof. If $c(b)$ is even, it is easy to see that at the beginning of step $i+1$, the graph will be $G_{c^{-1}(c(b)+1)}$: the only cycle found in part 3 of step i is $(h_1, v_1, h'_1, v'_1, h_1)$. Any other cycles would have to make use of one of the arcs pointing toward h'_1, but that is not possible by the vertex-disjointness property of the cycles in the subgraph used at part 3 of step i. After augmenting G_b according to cycle $(h_1, v_1, h'_1, v'_1, h_1)$, it is easy to check that the graph is equal to G_{b+1}. Also, observe that Property A_{i+1} holds.

If $c(b)$ is odd and not equal to $2^k - 1$, then define j to be the largest index such that $b_{j'} = 1$ for all $j' \leq j$. Then, in part 3 of step i, the cycle $(h_1, v'_1, h'_1, v_1, h_2, v'_2, h'_2, v_2, \ldots, h_j, v'_j, h'_j, v_j, h_{j+1}, v_{j+1}, h'_{j+1}, v'_{j+1}, h_1)$ is found, and no other cycle is found, because otherwise h_1 would be in such a cycle: a contradiction. It is not hard to verify that property A_{i+1} holds, and the graph that now arises at the beginning of step $i + 1$ is again a single absorbing set that is not paired symmetric, because of v_{k+1}. Step $i + 2$ will therefore certainly be reached, and it can be verified by similar reasoning as before that again a single cycle is found in part 3 of step $i + 1$. This cycle is $(h_1, v'_{j+1}, h_{j+1}, v_j, h_j, v_{j-1}, h_{j-1}, v_{j-2}, h_{j-2}, \ldots, v_1, h_1)$. Augmenting the graph on this cycle makes the graph exactly equal to $G_{c^{-1}(c(b)+1)}$. Moreover, Property A_{i+2} holds. □

Property A_1 is certainly satisfied, and the graph at step 1 is $G_{000\ldots}$. By straightforward induction, using the claim above, it follows that for all bit-strings b of length k there is indeed a step i_b such that the graph at the beginning of step i_b is equal to G_b. □

8.7 Discussion

Properties	TTAS	TCR	GATTC
Core, Pareto optimal	✓	✓	**✔** Th. 177
Strict core (if non-empty)	✓	**✔** Cor. 183	**✔** Th. 178
Truthful	✓	✓	**✗** Th. 181
Polynomial time	**✗** Th. 184	✓	**✗** Th. 184

Table 8.1: Housing market mechanisms: new results are in a bolder font.

We analyzed and compared the two housing market mechanisms TTAS and TCR. Whereas it was shown that TTAS may take exponential time, TCR was shown to be strict core selecting just like TTAS. The new and old results are summarized in Table 8.1. Our abstraction from TTAS and TCR to GATTC helps identify the crucial higher level details and commonality of both TTAS and TCR. This leads to simple proofs for properties satisfied by any GATTC mechanism. Whereas core, strict core, and Pareto optimality are properties that can be fulfilled by any GATTC mechanism, additionally satisfying truthfulness requires subtlety in choosing which cycles are implemented in which order. This additional complexity leads to an exponential time lower bound in the case of TTAS and a difficulty in having a very simple description in the case of TCR.

Our study leads to a number of further research questions. It will be interesting to characterize the subset of GATTC mechanisms which are truthful or are both truthful and run in polynomial time. Another question is to see whether being a GATTC mechanism is a necessary condition to simultaneously achieve core stability, Pareto optimality and strict core stability. We have seen that all known housing market mechanisms which are core selecting and Pareto optimal are also strict core selecting (if the strict core is non-empty). This raises the question whether every housing market mechanism which is core selecting and Pareto optimal is also strict core selecting (if the strict core is non-empty).

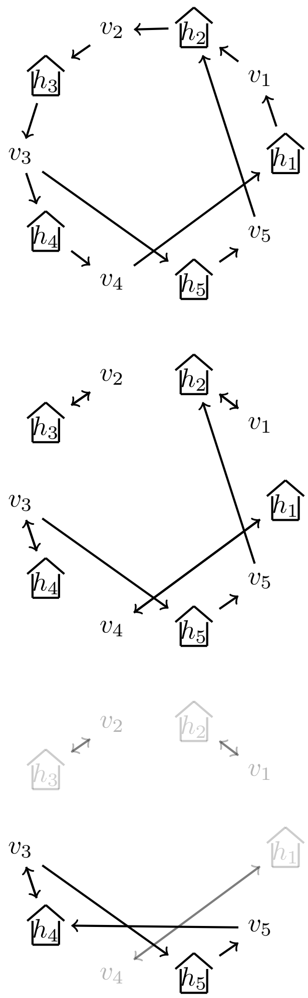

Figure 8.1: Illustration of the first steps of a GATTC mechanism applied to the housing market H in Example 176. The top figure shows the graph $G(H)$ as initialized. The algorithm proceeds by executing Step 1 zero times, removing no paired symmetric absorbing sets in Step 2 (as there are none), and implementing the cycle $(v_1, h_2, v_2, h_3, v_3, h_4, v_4, h_1, v_1)$ in Step 3. The graph after implementing this cycle is shown in the middle figure. Subsequently, the mechanism removes the paired symmetric absorbing sets, forcing v_5 to point to its second-most preferred houses, i.e., house h_4.

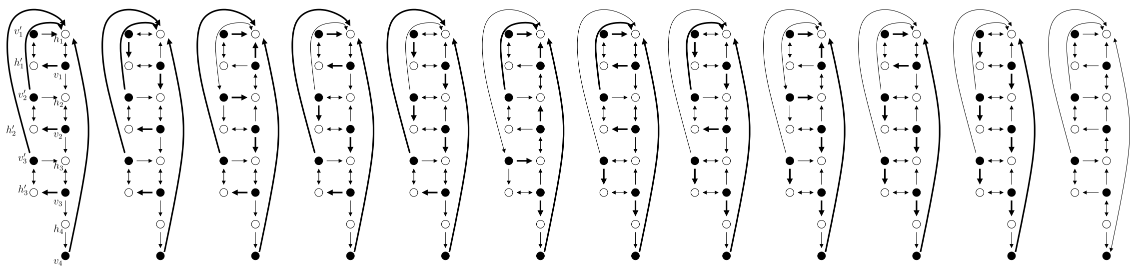

Figure 8.2: **(Illustrative example for the proof of Theorem 184.)** The graph at the beginning of every step of the TTAS mechanism when it is ran on the instance M_3. Black vertices represent players and white vertices represent houses. When an arc is drawn that has arrows pointing to both its vertices, say vertices a and b, then it stands for the presence of arcs (a, b) and (b, a) in the graph. At the graph for step 1, the names of the vertices are displayed. This is omitted for subsequent steps. In the last step it can be seen that the entire graph is paired symmetric. For every step i except the last one, an arc is displayed in bold in the graph of step i when that arc points from a player to a house and when that arc is included in the subgraph generated in part 3 of step i (the remaining arcs in this subgraph are all arcs pointing from houses to players). When in some step, the graph at the beginning of that step equals G_b for some $b \in \{0,1\}^k$, then this is indicated in the figure by the tag "(G_b)" after the step number.

Chapter 9

Complexity of coalition structure generation*

We deal in this chapter with the following problem: given a cooperative game (n, v) (see Definition 26), find a partition $\mathcal{P}$ of $[n]$ such that $\sum_{P \in \mathcal{P}} v(P)$ is maximized. This problem is referred to as the *coalition structure generation problem*, and is an important and widely studied problem in algorithmic cooperative game theory; especially for cooperative games that are known to not satisfy the superadditivity property. We will refer to the coalition structure generation problem as OPTCS. Computing optimal coalition structures is a natural problem in which the aim is to utilize resources in the most efficient manner.

OPTCS is relevant to certain optimization problems in strategic games as well: the main problem studied in Chapter 7 (i.e., finding an optimal assignment of the players to the facilities, in a singleton congestion game with externalities) can be straightforwardly seen as a coalition structure generation problem of a cooperative game. This chapter can thus be regarded as a more general study of such OPTCS problems.

We first study the complexity of OPTCS under a "black-box" model, where the algorithm has only query access to the given characteristic function. It turns out that for this model, a polynomial time algorithm exists that solves OPTCS for all cooperative games for which the *types* of the players are known, and this number of types is bounded by a constant. This notion of player type is formally defined later; one may think of a set of players having the same type iff permuting these players does not change the cooperative game. As a corollary of our result, we obtain a polynomial

*The contents of this chapter have been published as Aziz and De Keijzer [2011a,b] time algorithm to compute an optimal partition for three classes of games: *weighted voting games* with a constant number of weight values, *linear games* with a constant number of desirability classes, and *cooperative skill games* with a constant number of skills. However, on the negative side we show that an exponential number of queries to the characteristic function may be needed in order to find out which player has which type, even if there are only two types and the input cooperative game is simple (i.e., a monotone game of which the characteristic function maps to $\{0, 1\}$).

Lastly, we consider the coalition structure generation problem for various well-known cooperative games defined compactly on combinatorial domains. For these games, we characterize the complexity of computing an optimal coalition structure by presenting polynomial-time algorithms, approximation algorithms, or NP-hardness and inapproximability lower bounds.

The reader is recommended to study Sections 1.2 and 1.3.2 from the introduction chapter of this thesis, prior to reading this chapter.

9.1 Background

Cooperative games have been used to model various cooperative settings in operations research, artificial intelligence, and multiagent systems [see e.g, Bachrach and Rosenschein, 2008, 2009, Elkind et al., 2007]. The area of cooperative game theory, which studies (among other aspects) coalition formation, has seen considerable growth over the last few decades.

OPTCS has received attention in the artificial intelligence community where the focus has generally been on computing optimal coalition structures for general cooperative games [Michalak et al., 2010, Sandholm et al., 1999] without any combinatorial structure. Traditionally, the input considered is an oracle called a characteristic function which returns the value for any given coalition (in time polynomial in the number of players). In this setting, it is generally assumed that the value of a coalition does not depend on players that are not in the coalition.

Computing optimal coalition structures is a computationally hard task because the number of coalition structures that grows exponentially in the number of players, and in general it is necessary to inspect all coalition structures in order to find the optimal one. The total number of coalition structures for a player set of size n is $B_n \in \Theta(n^n)$ where B_n is the nth Bell number. Various algorithms have been developed in the last decade which attempt to satisfy many desirable criteria, e.g. outputting an optimal solution or a good approximation, the ability to prune, the *anytime* property, worst-case guarantees, distributed computation, etc. [Michalak et al., 2010, Rahwan et al., 2009, Sandholm et al., 1999, Service and Adams, 2010]. In all of the cases, the algorithms have a worst-case time complexity that is exponential in n.

Alternatively, cooperative games can be represented compactly on combinatorial domains where the valuation function is implicitly defined [Deng and Fang, 2008, Deng and Papadimitriou, 1994]. Numerous such classes of cooperative games have been the subject of research: weighted voting games [Elkind et al., 2007]; skill games [Bachrach and Rosenschein, 2008]; multiple weighted voting games [Aziz et al., 2010]; network flow games [Bachrach and Rosenschein, 2009]; spanning connectivity games [Aziz et al., 2009]; and matching games [Kern and Paulusma, 2003]. Apart from some exceptions (skill games [Bachrach et al., 2010b] and marginal contribution nets [Ohta et al., 2009]), most of the Algorithms research for these classes of games has been on computing stability-based solutions.

9.1.0.0.1 Contribution and Outline In Section 9.2, we define formally the coalition structure generation problem, OPTCS, and we introduce the various classes of cooperative games for which we study the OPTCS problem.

Subsequently, in Section 9.3, we consider the case where there is a fixed number of *types* of players, and it is known in advance which player has which type. We show that a polynomial time algorithm for OPTCS exists in this case, and as a corollary it holds that OPTCS is solvable in polynomial time for weighted voting games with a constant number of weight values, linear games with a constant number of desirability classes, and cooperative skill games with a constant number of skills. On the negative side, we show that in general an exponential number of queries to the characteristic function may be needed in order to derive the types of the players even in case the input cooperative game is *simple* (i.e., a monotone game for which the characteristic function maps to $\{0, 1\}$) and there are only 2 types. The latter holds even in expectation (when the algorithm is randomized), and is shown through an application of Yao's minimax principle.

Some specific cooperative games are studied in Sections 9.4 and 9.5: Simple games and *weighted voting games* are studied in Section 9.4: We present a 2-approximation algorithm for the case of weighted voting games and show that this approximation bound is the best possible. For general simple games, the degree of inapproximability deteriorates to $n^{1-\epsilon}$ (for any $\epsilon > 0$). The approximation and inapproximability results concerning weighted voting games may be of independent interest since they address a problem in the family of knapsack problems [Kellerer et al., 2004] which has not been studied before.

We examine in Section 9.5 some well-known cooperative games based on graphs, and characterize the complexity of computing the optimal coalition structures.

Finally, a table summarizing all results of this chapter is provided in section 9.6.

9.2 Preliminaries

In this section, we define several important classes of cooperative games and formally define our fundamental computational problem OPTCS.

9.2.1 Simple Games and Threshold Games

Definition 187 (Simple game, winning and losing coalition, minimal winning coalition)**.** A *simple game* is a monotone cooperative game (n, v) with $v : 2^{[n]} \to \{0, 1\}$ such that $v(\varnothing) = 0$ and $v([n]) = 1$. A coalition $P \subseteq [n]$ is said to be *winning* iff $v(P) = 1$ and *losing* iff $v(P) = 0$. A *minimal winning coalition* (MWC) of a simple game v is a winning coalition for which removing of any player makes the coalition losing. Likewise, one can define the set of *maximal losing coalitions* as the set of losing coalitions for which adding any player makes the coalition winning. A simple game can be represented by its set of minimal winning coalitions, and by its set of maximal losing coalitions.

For any monotone cooperative game, one can construct a corresponding threshold game [Aziz et al., 2010]. Threshold versions of cooperative games are common in the multiagent systems literature; see for instance [Bachrach and Rosenschein, 2009, Elkind et al., 2007].

Definition 188 (Threshold version of a cooperative game)**.** For a cooperative game (n, v) and a number $t \in \mathbb{R}$, the *corresponding t-threshold game* is defined as the cooperative game (n, v^t), where for all $P \subseteq [n]$,

$$v^t(P) = \begin{cases} 1 & \text{if } v(P) \geq t \\ 0 & \text{otherwise.} \end{cases}$$

It can easily be verified that if (n, v) is monotone, then for any threshold $t \leq v([n])$, the t-threshold version (n, v^t) is a simple game.

9.2.2 Cooperative Game Classes

We now review a number of specific classes of cooperative games. Here we adopt the convention that if CLASS denotes a particular class of games, we have T-CLASS refer to the class of threshold games corresponding to games in CLASS, i.e., for every threshold $t \in \mathbb{R}$, (n, v^t) is in T-CLASS if and only if (n, v) is in CLASS.

Weighted voting games are a widely studied class of monotone games (see e.g. Taylor and Zwicker [1999]).

Definition 189 (Weighted voting games)**.** A *weighted voting game (WVG)* is a simple game (n, v) for which there is a *quota* $q \in \mathbb{R}_{\geq 0}$ and a *weight* $w_i \in \mathbb{R}_{\geq 0}$ for each player $i \in [n]$ such that for all $P \subseteq [n]$,

$$v(P) = 1 \text{ if and only if } \sum_{i \in P} w_i \geq q.$$

The WVG with quota q and weights $w_1, \ldots, w_n$ is denoted by $[q; w_1, \ldots, w_n]$, where we commonly assume $w_i \geq w_{i+1}$ for $i \in [n-1]$.

A *multiple weighted voting game* (MWVG) is the simple game (n, v) for which there are WVGs $(n, v_1), \ldots, (n, v_m)$ such that for all $P \subseteq [n]$,

$$v(P) = 1 \text{ if and only if } v_k(P) = 1 \text{ for all } k \in [m].$$

The remaining classes of cooperative games we will study are defined on graphs. These are *spanning connectivity games*, *independent set games*, *matching games*, *network flow games*, and *induced subgraph games*, where either nodes or edges are controlled by players, and the value of a coalition of players depends on their ability to connect the graph, enable a bigger flow, or obtain a heavier matching or edge set.

Definition 190 (Spanning connectivity game [Aziz et al., 2009])**.** For each connected undirected graph $G = (V, E)$, we define the *spanning connectivity game (SCG)* on G as the simple game (m, v) where $m = |E|$ and each player corresponds to an edge in E, according to a bijection $b : [m] \to E$. For all $P \subseteq [m]$, $v(P) = 1$ if and only if there exists some $P' \subseteq P$ such that $(V, \{b(i) \mid i \in P'\})$ is a spanning tree.

Definition 191 (Independent set game [Deng and Fang, 2008])**.** For each connected undirected graph $G = ([n], E)$, we define the *independent set game (ISG)* on G as the game (n, v) such that for all $P \subseteq [n]$, $v(P)$ is the cardinality of the maximum independent set on the subgraph of G induced by P.

Definition 192 ((Unweighted) matching game [Kern and Paulusma, 2003])**.** Let $(G = ([n], E), w)$ be an undirected edge-weighted graph (i.e., $w : E \to \mathbb{R}$). The *(unweighted) matching game* corresponding to G is the cooperative game (n, v) such that for all $P \subseteq [n]$, the value $v(P)$ equals the weight of the maximum weighted matching of the subgraph induced by P.

Definition 193 (Induced subgraph game [Deng and Papadimitriou, 1994])**.** For an undirected edge-weighted graph $(G = ([n], E), w)$ (i.e., $w : E \to \mathbb{R}$), the *induced subgraph game (ISG)* corresponding to G is the cooperative game (n, v) such that for all $P \subseteq [n]$, $v(P)$ is the total weight of edges in the subgraph induced by P. In this chapter, we sometimes assume that the graph corresponding to a induced subgraph game has only positive edge weights and denote the class of such induced subgraph games by $\text{ISG}_{\geq 0}$. We denote the class of induced subgraph games where negative edge weights are allowed, by ISG. Note that for this latter general class of induced subgraph games, we allow the characteristic function v to map to negative numbers.

For the definition of network flow games, we need to introduce the notions of flow networks and network flows.

Definition 194 (Flow network, network flow, value of a flow)**.** A *flow network* $((V, E), c, s, t)$ consists of a directed graph (V, E), where $c : E \to \mathbb{R}_{\geq 0}$ is referred to as the *capacity function*, and $s \in V$ and $t \in V$ are referred to as the *source* and *sink* vertices, respectively. A *network flow* for $((V, E), c, s, t)$ is a function $f : E \to \mathbb{R}_{\geq 0}$, that *obeys* the capacity function (i.e., $f(e) \leq c(e)$ for all $e \in E$), as well as the condition that the total flow entering any vertex (other than s and t) equals the total flow leaving the vertex (i.e., for all $v \in V \setminus \{s, t\}$ it holds that $\sum_{e \in \text{in}(v)} f(e) = \sum_{e \in \text{out}(v)} f(e)$ where $\text{in}(v) = \{(v_1, v_2) \in E \mid v_2 = v\}$ and $\text{out}(v) = \{(v_1, v_2) \in E \mid v_1 = v\}$). The *value* of the network flow f is the amount flowing out of the source (i.e., $\sum_{e \in \text{out}(s)} f(e)$).

Definition 195 (Network flow game [Bachrach and Rosenschein, 2009])**.** Given a flow network $((V, E), c, s, t)$, the associated *network flow game (NFG)* is the cooperative game (n, v), where $m = |E|$ and each player corresponds to an edge in E, according to a bijection $b : [m] \to E$. For each $P \subseteq [m]$, the value $v(P)$ is the maximum value of a network flow f for $((V, E, c, s, t)$ such that $f(e) = 0$ for all $e \in \{b(i) \mid i \notin P\}$.

Definition 196 (Path cooperative games)**.** Let $G = ([n], E)$ be a directed graph, and let $s, t \in [n]$

- the *edge (s, t)-path cooperative game (EPCG)* corresponding to G is a simple game (m, v) such that $m = |E|$ and each player corresponds to an edge in E, according to a bijection $b : [m] \to E$. For each $P \subseteq [m]$, $v(P) = 1$ if and only if there exists an (s, t)-path in $(V, \{b(i) \mid i \in P\})$.
- the *vertex path cooperative game (VPCG)* corresponding to G is a simple game (n, v) such that for all $P \subseteq [n]$, $v(P) = 1$ if and only if there is an (s, t)-path in the subgraph of G induced by P.

Definition 197 (Cooperative skill games Bachrach and Rosenschein [2008])**.** A *cooperative skill domain* is a tuple (n, m, U, S, T, u) where $n, m \in \mathbb{N}_{\geq 1}$ and $[n]$ is the *player set*, $[m]$ is the *task set*, and U is a finite set of *skills*. S is an n-dimensional vector of subsets of U and $S_i, i \in [n]$, is referred to as the *skill set of player* i. T is a

m-dimensional vector of subsets of U, and $S_j, j \in [m]$, is referred to as the *set of skills required for task* j. The set of skills that a coalition $P \subseteq [n]$ has is $S(P) = \bigcup_{i \in P} S_i$. A coalition $P \subseteq [n]$ is said to *be able to perform* task j if $T_j \subseteq S(P)$. The set of tasks that coalition P is able to perform is $T(P) = \{j \mid T_j \subseteq S(P)\}$. The function $u : 2^T \to \mathbb{R}$ is referred to as the *task value function*, and is monotone. The *cooperative skill game (CSG)* corresponding to a given cooperative skill domain (n, m, U, S, T, u) is the game (n, v) such that for all $P \subseteq [n]$, $v(P) = u(T(P))$. A *weighted task cooperative skill game (WTCSG)* is a CSG corresponding to a cooperative skill domain (n, m, U, S, T, u) where each task $j \in T$ has an associated weight $w_j \in \mathbb{R}_{\geq 0}$ and the task value function is given by $u(T') = \sum_{j \in T'} w_j$. A threshold version of WTCSG can be defined according to Definition 188.

Definition 198 (Linear games Taylor and Zwicker [1999])**.** On a cooperative game (n, v), we define the *desirability relation* $\preceq_D$ as follows: we say that a player $i \in [n]$ is *more desirable than* a player $i' \in [n]$ ($i \succeq_D i'$) if for all coalitions $P \subseteq [n] \backslash \{i, i'\}$ it holds that $v(P \cup \{i\}) \geq v(P \cup \{i'\})$. The relations $\succ_D$ ("strictly more desirable"), $\sim_D$ ("equally desirable"), and $\preceq_D$ and $\prec_D$ ("(strictly) less desirable") are defined in the obvious way. *Linear games* are simple games with a complete desirability relation, i.e. every pair of players is comparable with respect to $\preceq_D$. Weighted voting games form a strict subclass of linear games. A linear game on players $[n]$ is *canonical* iff $\forall i, i' \in [n]$ such that $i < i'$ it holds that $i \succeq_D i'$. A *right-shift* of a coalition P, in a linear game (n, v), is a coalition that can be obtained by a sequence of replacements of players in P by less desirable players. A *left-shift* of a coalition P is defined analogously. Canonical linear games can be represented by listing their *shift-minimal winning coalitions*: minimal winning coalitions for which it holds that any right-shift is losing. Similarly they can be represented by listing their *shift-maximal losing coalitions*, defined analogously.

9.2.3 Problem Definition

We formally define *coalition structures* and OPTCS.

Definition 199 ((Optimal) coalition structure)**.** A *coalition structure* for a cooperative game (n, v) is a partition of $[n]$. The *value attained by* a coalition structure $\mathcal{P}$, denoted $v(\mathcal{P})$ (we overload notation), is defined as $\sum_{P \in \mathcal{P}} v(P)$. A coalition structure $\mathcal{P}$ is *optimal* for (n, v) when $v(\mathcal{P}) \geq v(\mathcal{P}')$ for every coalition structure $\mathcal{P}'$ for (n, v).

We consider the following computational problem.

Definition 200 (Problem OPTCS)**.** For any class of cooperative games X, the problem OPTCS(X) is as follows: given a cooperative game $(n, v) \in X$ in its natural representation, compute an optimal coalition structure. Throughout this chapter, it will be clear what is meant with the "natural representation", for each class X that we consider.

9.3 Games with Fixed Player Types

We study the problem of computing an optimal coalition structure for a game in the case that the number of *player types* is fixed. Shrot et al. [2010] considered player types and showed that some intractable problems become tractable when only dealing with a fixed number of player types. However, the coalition structure generation problem was not addressed in that paper.

Definition 201 (Strategic equivalence, player type)**.** For a cooperative game (n, v), we call two players $i, i' \in [n]$ *strategically equivalent* iff for every coalition $P \subseteq [n] \backslash \{i, i'\}$ it holds that $v(P \cup \{i\}) = v(P \cup \{i'\})$. When two players $i, i' \in [n]$ are strategically equivalent, we say that i and i' are of the same *player type*.

Definition 202 (Valid type-partition)**.** A *valid type-partition* for a game (n, v) is a partition $\mathcal{P}$ of $[n]$ such that for each player set $P \in \mathcal{P}$, all players in P are of the same player type.

Let $k \in \mathbb{N}_{\geq 1}$ be a given constant. We consider the computational problem where the goal is to compute an optimal coalition structure for a cooperative game (n, v), given as input a valid type partition $\mathcal{P}$ for (n, v), such that $|\mathcal{P}| \leq k$. In general, it is not easy to verify that a given partition for a simple game is a valid type-partition. But under the assumption that we are given a valid type partition of which the cardinality is at most k, and under the assumption that it is computationally easy to evaluate v, it turns out that an optimal coalition structure can be computed in polynomial time.

9.3.1 A General Algorithm

We show that there exists a polynomial time algorithm to compute an optimal coalition structure for any cooperative game when we are given a valid type partition with the number of player types bounded by a constant. Our algorithm utilizes dynamic programming to compute an optimal coalition structure provided there are a constant number of player types.

Theorem 203. *Let $k \in \mathbb{N}_{\geq 1}$. There exists an algorithm A that finds, given a game (n, v) and type partition $\mathcal{P}$, an optimal coalition structure for (n, v) in time polynomial in n, provided that v can be evaluated in polynomial time, the input type partition $\mathcal{P}$ is a* *valid type partition, and $|\mathcal{P}| \leq k$. More precisely, A runs in time $O((n+1)^k \cdot time(v) + (n+1)^{2k})$, where time$(v)$ is the time it takes to evaluate v.*

Proof. Let (n, v) be the input cooperative game, let $\mathcal{P} = \{T_1, \ldots, T_k\}$ be the input type-partition, and assume $\mathcal{P}$ is valid. We define *coalition types* as follows. For numbers $t_1, \ldots, t_k \in [n] \cup \{0\}$, the *coalition-type* $T(t_1, \ldots, t_k)$ is the set of coalitions $\{P \subseteq [n] \mid \forall i \in \{1, \ldots, k\} : |P \cap T_i| = t_i\}$. In words: coalitions in coalition-type $T(t_1, \ldots, t_k)$ have t_i players of type T_i, for all $i \in [k]$. Note that v maps all coalitions of the same coalition type to the same value.

First, our algorithm computes a table V of characteristic function values for each coalition type. In order to do this we need to query v at most $(n+1)^k$ times, since $t_i \in [n] \cup \{0\}$ for all $i \in [k]$. Let time(v) denote the time it takes to query v, then computing V takes $O(n^k \cdot \text{time}(v))$ time.

We proceed with a dynamic programming approach in order to find an optimal coalition structure: Let $f(a_1, \ldots, a_k)$ be the optimal value attained by an optimal coalition structure on the cooperative game (n, v) restricted to a subset of players $N' \in \{N' \mid \forall i \in \{1, \ldots, k\} : |N' \cap T_i| = a_i\}$. Note that $f(a_1, \ldots, a_k)$ is indeed well-defined since it is independent of our particular choice of N'. We are interested in computing $f(|T_1|, \ldots, |T_k|)$. By γ, we signify the set of type-partitions $\mathcal{P}$ such that $v(\mathcal{P}) = v(\varnothing)$.

Since $v(\varnothing) = 0$, the following recursive definition of $f(a_1, \ldots, a_k)$ holds:

$$f(a_1, \ldots, a_k) = \begin{cases} 0 & \text{if } a_i = 0 \text{ for all } i \in [k], \\ \max\{f(a_1 - b_1, \ldots, a_1 - b_k) + v(b_1, \ldots, b_k) \\ \qquad \mid \forall i \in [k] : b_i \leq a_i\} & \text{otherwise.} \end{cases} \tag{9.1}$$

The recursive definition of $f(a_1, \ldots, a_k)$ directly implies a dynamic programming algorithm. The dynamic programming approach works by filling in a $|T_1| \times \cdots \times |T_k|$ table Q, where the value of $f(a_1, \ldots a_k)$ is stored at entry $Q[a_1, \ldots, a_k]$. Once the table has been computed, $f(|T_1|, \ldots, |T_k|)$ is returned. The entries of Q are filled in according to (9.1). In order to utilize (9.1), "lower" entries are filled in first, i.e. $Q[a_1, \ldots, a_k]$ is filled in before $Q[a'_1, \ldots, a'_k]$ if $a_i \leq a'_i$ for $1 \leq i \leq k$. Evaluating (9.1) then takes $O((n+1)^k)$ time (due to the "otherwise"-case of (9.1), where the maximum of a set of at most $(n+1)^k$ elements needs to be computed). There are $O((n+1)^k)$ entries to be computed, so the algorithm runs in $O(n^k \cdot \text{time}(v) + (n+1)^{2k})$ time.

It is straightforward to extend this algorithm so that it (instead of outputting only the optimal value) also computes and outputs an actual coalition structure that attains the optimal value. To do so, maintain another table $|T_1| \times \cdots \times |T_k|$ table R. At each point in time that some entry of Q is computed, say $Q[a_1, \ldots, a_k]$, now we also fill in $R[a_1, \ldots, a_k]$. $R[a_1, \ldots, a_k]$ contains a description of a set $\mathcal{P}$ of coalitions such that $\sum_{P \in \mathcal{P}} v(P) = f(a_1, \ldots, a_n)$ and $\bigcup \mathcal{P} \in T(a_1, \ldots, a_k)$. It suffices to describe $\mathcal{P}$ by simply listing the type of each $P \in \mathcal{P}$, and it is straightforward to verify that we can set $R(a_1, \ldots, a_k)$ to $\varnothing$ if $(a_1, \ldots, a_k) \in \gamma$, and otherwise we set $R(a_1, \ldots, a_k)$ to $(P(a_1 - b_1, \ldots, a_1 - b_k), (b_1, \ldots, b_k))$, where $(b_1, \ldots, b_k)$ is the argument in the $\max$-expression of (9.1). □

9.3.2 Difficulty of Finding Types

The polynomial time algorithm given in the proof of Theorem 203 relies on the promise that the input type partition is valid. A natural question is now whether it is also possible to efficiently compute the type partition of a game in polynomial time when given only the weaker promise that the number of player types is constant k. We answer this question negatively. For randomized algorithms, we show high communication complexity is necessary, i.e., we show that an exponential amount of information is needed from the characteristic function v when we are given no information on the structure of the characteristic function and we rely only on querying v. In fact, the theorem states that this is the case even when v is simple and $k = 2$. It should be noted that this result also holds for deterministic algorithms, since they are a special case of randomized algorithms. Despite this negative result, we show in Section 9.3.3 that we can do better for some subclasses of cooperative games, where we are provided with information on the structure of function v.

Theorem 204. *Any randomized algorithm that computes a player type-partition when given as input a monotone simple game (n, v) that has 2 player types, requires at least $\Theta(\frac{2^n}{\sqrt{n}})$ queries to v.*

Proof. We use Yao's minimax principle Yao [1977], which states that the expected cost of a randomized algorithm on a given problem's worst-case instances is at least the lowest expected cost among all deterministic algorithms that run on any fixed probability distribution over the problem instances.

Consider the following distribution over the input, where the player set is $[n]$ and n is even, the number of player types is always $k = 2$, and the given cooperative game (n, v) is a simple game. Valuation v is drawn uniformly at random from the set $V = \{v_P \mid P \subset [n], |P| = n/2\}$ where in v_P, we call P the *critical coalition*. Function v_P is specified as follows:

- $v_P(D) = 0$ when $|D| < n/2$;
- $v_P(D) = 1$ when $|D| > n/2$;

- $v_P(D) = 1$ when $D = P$, i.e. D is the critical coalition;
- $v_P(D) = 0$ otherwise.

Observe that there are exactly two player types in any of the games in V: for v_P, the type partition is $(P, [n]\backslash P)$. Also observe that because v is drawn from V, for coalitions P of size $\frac{n}{2}$, $v(P) = 1$ with probability $1/\binom{n}{n/2}$, because v is drawn uniformly at random from V.

Now let us consider an arbitrary deterministic algorithm A that computes the type-partition for instances in this input distribution by queries to v. Let P be the critical coalition of $n/2$ players such that $v(P) = 1$. A will have to query $v(P)$ in order to know which characteristic function of V has been drawn, and thus determine the type partition correctly. Let $Q(v)$ be the sequence of queries to v that A generates. Let $\mathcal{Q}'(v)$ be the subsequence obtained by removing from $\mathcal{Q}(v)$ all queries $v(D)$ such that $|D| \neq n/2$ and all queries that occur after $v(P)$. Because A is deterministic, the query sequence of A is the same among all instances up to querying the critical coalition, since the critical coalitions are the only points in which the characteristic functions of V differ from each other. Therefore the expected length of $\mathcal{Q}'(v)$ is $\binom{n}{n/2}/2$. Because A was chosen arbitrarily, we conclude that also the most efficient deterministic algorithm is expected to make at least $\binom{n}{n/2}/2 = \Theta(\frac{2^n}{\sqrt{n}})$ queries to v, and the theorem now follows from Yao's minimax principle. □

Shrot et al. [2010] showed that checking whether two players are of the same type is NP-hard for the cooperative games defined in Conitzer and Sandholm [2006], but those games are such that even computing the value of a coalition is NP-hard. It is possible to say something stronger.

Proposition 205. *There exists a representation for a class of cooperative games for which checking whether two players are of the same type is* coNP*-complete even though the value of each coalition can be computed in polynomial time under this representation.*

Proof. A coalition $P \subseteq [n] \setminus \{i, j\}$ such that $v(P \cup \{i\}) \neq v(P \cup \{j\})$ is a polynomial time certificate for membership in coNP. Also, it is well known that checking whether two players in a WVG have the same Banzhaf index is coNP-hard [Matsui and Matsui, 2000].[1] Since two players in a WVG are of the same type if and only if they have same the Banzhaf index, the claim follows. □

[1] The *Banzhaf index* of a player i is defined as the ratio of $|\{S \subseteq [n] \setminus \{i\} \mid v(S) = 0, v(S \cup \{i\}) = 1\}|$ and 2^{n-1}.

9.3.3 Applications of Theorem 203

Theorem 204 and Proposition 205 indicate that finding player types is in general a difficult task. Despite these negative results, Theorem 203 still applies to all classes of games and many natural settings where the type partition is implicitly or explicitly evident:

Corollary 206. *There exists a polynomial time algorithm that solves* OPTCS(WVG) *in case it holds that in the input game (given in the form of a quota and a weight for each player), the number of distinct weights is constant. There exists a polynomial time algorithm that solves* OPTCS(MWVG) *in case it holds that in the input game (given in the form of a vector of quotas and a vector of weights for each player), the number of distinct weight vectors for the players is constant.*

Proof. When two players have the same weight (in the case of WVGs) or weight vectors (in the case of MWVGs), they are strategically equivalent. Therefore we can type-partition the players according to their weights or weight vectors and apply Theorem 203. □

There exists a polynomial-time algorithm for computing the desirability classes, when given the list of shift-minimal winning coalitions of a linear game [Aziz, 2008]. This immediately yields the following corollary:

Corollary 207. *In the following cases, there exists a polynomial-time algorithm that computes an optimal coalition structure for linear games with a constant number of desirability classes:*

- *The input game is represented as a list of (shift-)minimal winning coalitions;*
- *The input game is represented as a list of (shift-)maximal losing coalitions.*

Bachrach et al. [2010b] proved that OPTCS(CSG) is polynomial time solvable if the number of tasks is constant and a certain associated graph, named the “skill graph”, has bounded treewidth. As a corollary of Theorem 203, we obtain a complementing positive result that applies to all of the cooperative skill games defined in [Bachrach and Rosenschein, 2008].

Corollary 208. *There exists a polynomial time algorithm that computes an optimal coalition structure for WTCSGs and T-WTCSGs with at most a fixed number of skills.*

Proof. Assume that the number of skills is bounded by a constant k'. Then there is a maximum of $2^{k'}$ player types. A polynomial-time algorithm that computes an optimal coalition structure now follows from Theorem 203. □

9.4 Weighted Voting Games and Simple Games

In this section, we examine weighted voting games (WVGs) and, more generally, simple games. Weighted voting games are cooperative games widely used in multiagent systems and artificial intelligence. We have already seen that there exists a polynomial time algorithm to compute an optimal coalition structure for WVGs with a constant number of weight values. We show that if the number of weight values is not a constant, then the problem becomes strongly NP-hard.

In the following propositions, we use the computational problem k-PARTITION. An instance of the problem k-PARTITION is a set of n integer weights $A = \{a_1, \dots, a_n\}$ and the question is whether it is possible to partition A, into k sets $P_1, \dots, P_k \subseteq A$ such that for $i, j \in [k]$, $P_i \cap P_j = \varnothing$ and $\bigcup_{i \in [k]} P_i = A$ and for all $i \in [k]$, $\sum_{a_j \in A_i} a_j = \sum_{j \in [n]} a_j / k$.

Proposition 209. *For a WVG, checking whether there is a coalition structure that attains a value of k or more is* NP*-complete.*

Proof. We prove this by a reduction from an instance of the classical NP-hard PARTITION problem to checking whether a coalition structure in a WVG gets value at least 2.

Without loss of generality, assume that $W = \sum_{a_i \in A} a_i$ is a multiple of k. Given an instance of k-PARTITION, $I = \{a_1, \dots, a_k\}$, we can transform it to a WVG $v = [q; w_1, \dots, w_k]$ where $w_i = a_i$ for all $i \in \{1, \dots, k\}$ and $q = W/k$. Then the answer to I is yes if and only if there exists a coalition structure $\mathcal{P}$ for v such that $v(\mathcal{P}) = k$. □

Since 3-PARTITION is strongly NP-complete, it follows that OPTCS(WVG) is strongly NP-hard. This stands in contrast with other results concerning WVGs where computation becomes easy when the weights are encoded in unary [Matsui and Matsui, 2000]. Proposition 209 does not discourage us from seeking an approximation algorithm for WVGs. We show that there exists a polynomial time 2-approximation algorithm.

Proposition 210. *There exists a polynomial time* 2*-approximation algorithm for* OPTCS(WVG).

Proof. Consider the following algorithm: Let $[q; w_1, \dots, w_n]$ be the input. We assume without loss of generality that $w_i \leq q$ for all $i \in [n]$. The algorithm first sets $p[0] := 0$, and then computes for some number c the values $p[1], \dots, p[c]$ using the rule

$$p[i] := \begin{cases} n & \text{if } \sum_{k=p[i-1]+1}^{n} w_k < q, \\ \min\{j \mid \sum_{k=p[i-1]+1}^{j} w_k \geq q, (p[i-1]+1) \leq j \leq n\} & \text{otherwise,} \end{cases} \tag{9.2}$$

where c is taken such that $p[c] = n$. The algorithm outputs the coalition structure $\{P_1, \ldots, P_c\}$, where for $i \in [c]$, $P_i = \{p[i-1]+1, \ldots, p[i]\}$.

Observe that the coalitions P_1 to P_{c-1} are all winning and P_c is not necessarily winning, so the value of the computed coalition structure is at least $c-1$. By our assumption, the total weight of any of the coalitions $P_1, \ldots, P_{c-1}$ is less than $2q$, and the total weight of P_c is less than q. Therefore, the total weight of $[n]$ is strictly less than $q(2c-1)$, so the optimal value is at most $2c-2 = 2(c-1)$. This is two times the value of the coalition structure computed by the algorithm. $\square$

A tight example for the algorithm described in the proof of Theorem 210 would be $[q; q-\epsilon, q-\epsilon, \epsilon, \epsilon]$, where q is a fixed constant and ϵ is any positive real number strictly less than $q/2$. On this input, the algorithm outputs a coalition structure that attains a total value of 1, while clearly the optimal coalition structure attains a value of 2. The following proposition shows that there does not exist a better polynomial time approximation algorithm under the assumption that $\mathsf{P} \neq \mathsf{NP}$.

Proposition 211. *Let* $\alpha < 2$. *Unless* $\mathsf{P} = \mathsf{NP}$, *there exists no polynomial time algorithm that computes an* α*-approximate optimal coalition structure for a WVG.*

Proof. We would be able to solve the NP-complete problem 2-PARTITION in polynomial time if there existed a (< 2)-optimal polynomial-time approximation algorithm for OPTCS(WVG). We could reduce a 2-PARTITION instance $(w_1, \ldots, w_n)$ to a weighted voting game $[q; w_1, \ldots, w_n]$ where $q = \sum_{i \in [n]} w_i/2$. Because the sum of all weights of the players is $2q$, a (< 2)-optimal approximation algorithm would output an optimal coalition structure when provided with this instance. The output coalition structure directly corresponds to a solution of the original 2-PARTITION instance, in case it exists. Otherwise, the value attained by the output coalition structure is 1. $\square$

Simple games that are not necessarily weighted, and are represented by the list of minimal winning coalitions, are even harder to approximate.

Proposition 212. OPTCS(MWC), *i.e.* OPTCS *for simple games represented as a list of minimal winning coalitions, cannot be approximated within a factor of* $n^{1-\epsilon}$ *in polynomial time, unless* $\mathsf{P} = \mathsf{NP}$.

Proof. This can be shown by means of a reduction from an instance of the classical NP-hard maximum clique (MAXCLIQUE) problem. It is known that MAXCLIQUE cannot be approximated within any constant factor [Maffioli and Galbiati, 2000].

Consider the instance I of MAXCLIQUE represented by an undirected graph $G_I = (V, E)$. Transform I into instance $I' = (|V|(|V|-1)/2, W^m)$ of OPTCS(MWC) in the following way. Define $N = \{\{v, v'\} \mid v \in V, v' \in V\}$ to be all subsets of

V of cardinality 2, and let each player of I' correspond to a subset in N, according to a bijection $b : [|V|(|V|-1)/2] \to N$. Next, for $i \in V$ define $P_i = \{\{i,j\} \in V \mid \{i,j\} \notin E\}$ and set $W^m = \{b^{-1}(P_i) \mid i \in V\}$. Now, two coalitions $b^{-1}(P_i)$ and $b^{-1}(P_{i'})$ are disjoint if and only if $\{i,i'\} \in E$. Then the maximum clique size is greater than or equal to k if and only if there is a coalition structure for I' that attains value k. Assume that there exists a polynomial time algorithm which computes a coalition structure $\mathcal{P}$ that gets a value that lies within a constant factor α of the maximum possible value k. Then we can use $\mathcal{P}$ to get a constant factor approximation solution to instance I in polynomial time in the following way. Assume without loss of generality that $\mathcal{P}$ contains only minimal winning coalitions, and consider the set of vertices $\{i \mid b^{-1}(P_i) \in \mathcal{P}\}$. Since for $P_i, P_{i'} \in \mathcal{P}$, P_i and $P_{i'}$ are disjoint, then we know that $(i,i') \in E$. Therefore the vertices $\{i \mid b^{-1}(P_i) \in \mathcal{P}\}$ form a clique of size k/α. □

9.5 Games on Graphs

Numerous classes of cooperative games are based on graphs. We characterize the complexity of OPTCS for many of these classes in the section. We first turn our attention to one such class for which the computation of cooperative game solutions is well studied [Deng and Papadimitriou, 1994]. We see that that OPTCS is computationally hard in general for induced subgraph games:

Proposition 213. *For the general class of induced subgraph games ISG, the problem* OPTCS *is strongly* NP*-hard.*

Proof. We prove this by presenting a reduction from the strongly NP-hard problem MAXCUT, where we are given a (non-negative) edge-weighted undirected graph and the goal is to find a cut of the graph such that the total weight of the edges crossing the cut is maximized. Consider an instance $I = (G = ([n], E), w)$ of MAXCUT (where $w : E \to \mathbb{R}_{\geq 0}$), and non-negative weights $w(i,i')$ for each edge $(i,i') \in E$. Let $W = \sum_{(i,i') \in E} w(i,i')$ and for a partition $\mathcal{P}$ of $[n]$ into two parts, define $\mathcal{P}(i)$ as the set of vertices in the same part as vertex i. We show that if there is a polynomial time algorithm that computes an optimal coalition structure, then there is a polynomial time algorithm for MAXCUT.

There exists a polynomial-time reduction that reduces I to an instance $I' = (([n+2], E'), w')$ of OPTCS(ISG) where $E' = E \cup \{\{n+1, i\} \mid i \in [n]\} \cup \{\{n+2, i\} \mid i \in [n]\} \cup \{\{n+1, n+2\}\}$. The weight function w' is defined as follows: $w'(\{i,i'\}) = -w(\{i,i'\})$ if $i, i' \in [n]$, $w'(\{i,i'\}) = W + 1$ if $i \in \{n+1, n+2\}$ and $i' \in [n]$, $w'(\{n+1, n+2\}) = -(|V|+1)W$.

We now show that a solution to instance I' of OPTCS(ISG) can be be used to solve instance I of MAXCUT. Assume that $\mathcal{P}'$ is an optimal coalition structure for I'. Then we know that $\mathcal{P}$ is of the form $\{\{n+1, A'\}, \{n+2, B'\}\}$ where (A', B') is a partition of $[n]$. We also know that $\sum_{\{i,i'\}\subseteq[n]:i\notin P(i')} w'(\{i,i'\})$ is minimized in $\mathcal{P}'$. Therefore, removing $n+1$ and $n+2$ from $\mathcal{P}$ yields a maximum weight cut for I. □

Observation 214. It is clear that for ISG$_{\geq 0}$, the coalition structure containing only the set of all players is the optimal coalition structure.

We now present some positive results concerning OPTCS for other games on graphs:

Proposition 215. OPTCS*(SCG) can be solved in polynomial time.*

Proof. OPTCS(SCG) is equivalent to computing the maximum number of edge disjoint spanning trees. This problem is solvable in $O(m^2)$ [Roskind and Tarjan, 1985]. □

Proposition 216. *For EPCGs and VPCGs,* OPTCS *can be solved in polynomial time.*

Proof. The problems are equivalent to computing the maximum number of edge disjoint and vertex disjoint (s,t)-paths respectively. There are well-known algorithms to compute them. For example, the maximum number of edge disjoint (s,t)-paths is equal to the max-flow value of the graph in which each edge has unit capacity.

The problem of maximizing the number of of vertex disjoint paths can be reduced to maximizing the number of of vertex disjoint paths in the following way: duplicate each vertex (apart from s and t) with one getting all ingoing edges, and the other getting all the outgoing edges, and an internal edge between them. □

Proposition 217. *The coalition structure containing only the coalition of all players is an optimal coalition structure for NFGs and matching games.*

Proof. We first prove the claim for NFGs. Assume there is a partition $\mathcal{P}$ of the edges which achieves a total value of x. This means that the sum of the maximum values of the flows of the flow networks $((V, E'), c, s, t), E' \in \mathcal{P}$ totals x. Since $\mathcal{P}$ is a partition, the sum of a set of feasible flows of the flow networks $((V, E'), c, s, t), E' \in \mathcal{P}$ is a feasible flow of $((V, E), c, s, t)$. Therefore, the sum of max-valued flows of $((V, E'), c, s, t)$ is also a feasible flow of $((V, E), c, s, t)$, and has a value equal to the sum of the values of these flows, which is x. This means that the set of all players is an optimal coalition structure.

Next, we prove the claim for matching games. Let $G = ([n], E)$ be a graph on which a matching game is defined. Assume there is a partition $\mathcal{P} = \{V_1, \ldots, V_k\}$ of the vertices that attains a value of x. Let the maximum weighted matching of the graph of G induced by V_i be m_i. Then we know that $\sum_{i\in[k]} m_i = x$. Since each member of $\mathcal{P}$ is mutually exclusive, for any $V_i, V_j \in \mathcal{P}$, the matchings in the subgraphs of G induced by V_i and V_j respectively, do not intersect. Now, consider the partition $\mathcal{P}' = \{[n]\}$ which consists of the set of all vertices. Then V has a value of at least x because the union of the maximum matchings of the subgraphs of G induced by $V_1, \ldots, V_k$, is a matching of G. This implies that that $v(\mathcal{P}') \geq x$. Therefore, the coalition structure consisting of only the set of all players attains a value that is at least the value attained by any other coalition structure. □

On the other hand, the threshold versions of certain games are computationally harder to solve because of their similarity to WVGs [Aziz et al., 2010]. As a corollary of Proposition 211, we obtain the following:

Corollary 218. *Let $\alpha < 2$. Unless* $\mathsf{P} = \mathsf{NP}$*, there exists no polynomial time algorithm which computes an α-approximately optimal coalition structure for T-NFGs, T-Matching games, and T-ISG$_{\geq 0}$.*

In some cases, OPTCS may be expected to be intractable because the cooperative game is defined on a combinatorial optimization domain which itself is intractable. We observe that even if computing the value of coalitions is intractable, solving OPTCS may be easy:

Observation 219. Given a maximum independent set game on graph $G = ([n], E)$, finding the value of the coalition $v([n])$ is NP-hard, but the optimal coalition structure is the one consisting of all singleton sets.

9.6 Summary of Results

We presented a general positive algorithmic result for coalition structure generation, namely that an optimal coalition structure can be computed in polynomial time if the player types are known and the number of player types is bounded by a constant. This is a useful result, as for many large multiagent systems it is a valid assumption that there are a lot of agents but the agents can be divided into a bounded number of strategic classes. For example, skill games are well motivated for coordinated rescue operation settings [Bachrach and Rosenschein, 2008, Bachrach et al., 2010b]. In these settings, there may be a large number of rescuers but they can be divided into a constant number of types such as firemen, policemen and medics.

We have also undertaken a detailed study of the complexity of computing an optimal coalition structure for a number of well-known cooperative games in that are rele-

Game class	Complexity of OPTCS
Coalition value oracle (valid type-partition & const. #types)	P (Theorem 203)
WVG (const number of weight values)	P (Corollary 206)
(T-)WTCSG (constant #skills or constant #types)	P (Corollary 208)
WCSG (constant #tasks, bounded tree-width skill graph)	P [Bachrach et al., 2010b]
SCG	P (Proposition 215)
EPCG and VPCG	P (Proposition 216)
NFG and matching game	P (Proposition 217)
$\text{ISG}_{\geq 0}$	P (Observation 214)
Independent Set Game	P (Observation 219)
ISG	Strongly NP-hard (Proposition 213)
(N, W^m)	NP-hard to approx. within a const. factor (Prop. 212)
WVG	Strongly NP-hard (Proposition 209); NP-hard to approximate within factor < 2 (Prop. 211)
T-Matching; T-NFG; T-ISG	NP-hard to approximate within factor < 2 (Cor. 218)
CSG	NP-hard [Bachrach et al., 2010b]

Table 9.1: Summary of complexity results for OPTCS

vant to AI, multiagent systems and operations research. The results are summarized in Table 9.1.

Chapter 10

Shapley Meets Shapley*

The unweighted version of matching games was introduced in the previous chapter (see Definition 192), where we concluded that its optimal coalition structure consists of the single set of all players in the game. For this reason, it makes sense to study solution concepts for matching games that divide the value $v([n])$ among the players. We introduce in this chapter a more general variant of these games, *weighted matching games*, and we will focus on the problem of computing one of the most well-known single-point solution concepts for this class of games: the Shapley value (Definition 31). The title of this chapter follows from the fact that both the the Shapley value and matching game concepts can be traced to the famous game theorist Lloyd S. Shapley.

Before reading this chapter, we advise the reader to make himself or herself familiar with Sections 1.2 and 1.3.2 of Chapter 1.

General weighted matching games are defined as follows.

Definition 220 (Matching Game)**.** A *(weighted) matching game* is a cooperative game (n, v) for which there is an undirected weighted graph $(G = ([n], E), w)$ (where $w : E \to \mathbb{R}_{\geq 0}$) such that for any $P \subseteq [n]$, $v(P)$ is the weight of a maximum weight matching of the subgraph of G induced b P. For a given weighted graph (G, w), we will denote by $MG(G, w)$ the matching game corresponding to graph G. For an unweighted graph $G = ([n], E)$, we denote by $MG(G)$ the matching game $MG(G, w)$ where $w(e) = 1$ for all $e \in E$.

Matching games constitute a fundamental class of cooperative games that help understand and model auctions and assignments.

*The contents of this chapter are based on Aziz and De Keijzer [2014, 2011c].

In economics and computer science, one of the most fundamental problems is the allocation of profits based on contributions of the nodes in a network. The problem has assumed even more importance as networks have become ubiquitous. The study we undertake in this chapter is related to this problem, as one way to address this problem is by defining a game over the network (a matching game in our case) and using its Shapley value as an allocation scheme for such profits.

After establishing some general insights, we will show that the Shapley value of matching games can be computed in polynomial time for some special cases: graphs with maximum degree two, and graphs that have a small modular decomposition into cliques or cocliques (complete k-partite graphs are a notable special case of this). The latter result extends to some other well-known classes of graph-based cooperative games.

We then show that computing the Shapley value of unweighted matching games is $\#\mathsf{P}$-complete in general. Finally, a fully polynomial-time randomized approximation scheme (FPRAS) is presented for general weighted matching games. This FPRAS can be considered the best positive result conceivable, in view of the $\#\mathsf{P}$-completeness result.

10.1 Background

Lloyd S. Shapley is one of the most influential game theorists in history. Among his numerous contributions, two of them are the following: (i) formulating the *assignment game* as a rich and versatile class of cooperative games [Shapley and Shubik, 1972], and (ii) proposing the *Shapley value* as a highly desirable solution concept for cooperative games [Shapley, 1953]. Both contributions have had far-reaching impact and were part of Shapley's Nobel Prize winning achievements.

The assignment game is a cooperative game based on bipartite graphs, and models the interaction between buyers and sellers. It is the *transferable utility* version of the well-known stable marriage setting and is a fundamental model that is used for modeling exchange markets and auctions [Roth and Sotomayor, 1990]. Assignment games were later generalized to *matching games* (see [Deng et al., 1999, Kern and Paulusma, 2003]).

Whereas the matching game is one of the most natural and important cooperative games, the Shapley value has been termed "the most important normative payoff division scheme" in cooperative game theory [Winter, 2002]. It is based on the idea that the payoff of an agent should be proportional to his marginal contributions to the payoff for the set of all players. For an excellent overview of the concept, we refer the reader to [Chapter 5, Moulin, 2003].

As network analysis becomes an increasingly important research area, centrality indices of graphs have received interest (see e.g., [Brandes and Erlebach, 2005]). The idea is to get a ranking of vertices according to their ability to connect with other vertices. Recently, a Shapley-value based game-theoretic approach has been used to gauge the centrality or connectivity of vertices by representing different valuation functions with a graph [see e.g., Michalak et al., 2013]. The motivation is that the Shapley value of a vertex captures various synergies which standard centrality measures do not. In this vein, the Shapley value is of importance outside game theory as well, as it constitutes an interesting method of gauging centrality/connectivity of the vertices. In particular it quantifies in a principled manner the ability of a vertex to match with other vertices to increase the value of coalitions.

The complexity of computing the Shapley value of important classes of cooperative games has been the topic of detailed studies. The papers Deng and Papadimitriou [1994] and Ieong and Shoham [2005] present algorithms to compute the Shapley value of *graph games* and *marginal contribution nets* respectively. On the other hand, computing the Shapley value is known to be intractable for a number of cooperative games (see e.g., [Elkind et al., 2009, Aziz et al., 2009]). Among the classes of cooperative games, matching games are one of the most well-studied. The core of matching games is characterized in Deng et al. [1999], where it is also shown that various computational problems regarding the core and the least core of matching games can be solved in polynomial time. For matching games, there has been considerable algorithmic research on the *nucleolus*: an alternative single-valued solution concept (see e.g., [Solymosi and Raghavan, 1994, Kern and Paulusma, 2003]).

10.2 Contributions and Outline

We address a gap in the compuational cooperative game theory literature, by studying the algorithmic aspects and computational complexity of the Shapley value for matching games for the first time. This gap is surprising on two fronts: First, computational aspects of Shapley values have been extensively studied for a number of cooperative games (see e.g., [Deng and Papadimitriou, 1994, Ieong and Shoham, 2005, Elkind et al., 2009]). Secondly, matching games are a well-established class of cooperative games, and the structure and computational complexity of computing important solution concepts such as the core, least core, and nucleolus have been examined in-depth for matching games (see e.g., [Alkan and Gale, 1990, Solymosi and Raghavan, 1994, Kern and Paulusma, 2003, Deng and Fang, 2008, Biró et al., 2011, 2013]).

We establish first some general insights (Section 10.3) and we consider two particular special cases for which the exact Shapley value can be computed in polynomial time (Section 10.4): graphs with a constant size decomposition into clique and coclique modules (these include e.g., complete k-partite graphs, for k constant), and graphs with maximum degree two. The non-trivial algorithm required for graphs of maximum degree two illustrates that exact computation of the Shapley value quickly becomes rather complex, even for very simple graph classes. We then move on to the central results of this chapter, which concern the general problem: we prove in Section 10.5 that the computational complexity of computing the Shapley value of matching games is $\#\mathsf{P}$-complete even if the graph is unweighted. The proof relies on Berge's Lemma and the fact that a certain matrix related to the Pascal triangle has a non-zero determinant. We subsequently present in Section 10.6 an *FPRAS* (i.e., a *fully polynomial time randomized approximation scheme*) for computing the Shapley value of (weighted) matching games. In view of our $\#\mathsf{P}$-completeness result, the FPRAS is the best possible result we can hope for.

Lastly, we discuss in Section 10.7 some prospectives for future research.

10.3 Preliminaries

We introduce first some essential basic notions related to graphs and matchings. While some readers may regard most of these notions as rudimentary, we define these anyway so that there will arise no confusion about the terminology used.

Matching basics. Given an undirected graph $G = ([n], E)$ (with vertex set $[n]$ and edge set E), a *matching* of G is a subset M of E such that $e \cap e' = \varnothing$ when $e, e' \in M$, $e \neq e'$. When discussing a particular matching M, we refer to the edges of a matching M as *matched edges*, and those outside M as *unmatched edges*. A *matched graph* is a pair (G, M) where G is a graph and M is a matching of G. A *maximum matching* of G is a matching of maximum cardinality among the set of all matchings of G.

We call a vertex i *exposed* or *unmatched* in (G, M) when i is not in any edge of M. Otherwise, we call i *matched*. An *alternating path* P in (G, M) is a path in G where the edges of P alternate between edges in M and edges in $E \backslash M$. An *augmenting path* P (with respect to a matching M) is an alternating path in G of which the endpoints are both exposed vertices. An augmenting path thus has odd length, starts with an unmatched edge, and ends with an unmatched edge. The following lemma is fundamental to matching theory:

Lemma 221 (Berge's lemma)**.** *Let $G = (V, E)$ be a graph. A matching M of G maximum if and only if there is no augmenting path in G with respect to M.*

Suppose we have a matching M for a graph G that is not a maximum matching. Then by the above lemma, there is an augmenting path P. It can be seen that removing from M the matched edges of P and adding to M the unmatched edges of P, gives us a bigger matching (i.e., a matching with one additional edge). We refer to this as the operation of *augmenting M along P*. Likewise, it is possible to augment a matching along an even-length alternating path with one exposed vertex and one matched vertex as endpoints. Augmenting along such a path does not increase the cardinality of the matching.

Observe that if P is an alternating path that is not augmenting, then it still possible to augment the matching along P iff one of the endpoints of P is an exposed vertex. Edmonds' blossom algorithm [Edmonds, 1965] is a polynomial time algorithm for finding a maximum weight matching in a graph.

Let M_1 and M_2 be two distinct maximum matchings for an unweighted graph $G = ([n], E)$. Then M_2 can be obtained from M_1 by a sequence of augmentations along mutually disjoint even-length alternating paths and even-length alternating cycles. A rough sketch of a proof for this is as follows: We investigate the symmetric difference D of M_1 and M_2, and conclude that D must be a collection of disjoint even-length paths and even length cycles of which the edges alternate between edges in M_1 and edges in M_2. A cycle in D must be an alternating cycle in M_1, and a path in D must be an alternating path in M_1. After augmenting M_1 along such a cycle or path, we obtain a matching M_3 such that the symmetric difference between M_3 and M_2 is D minus the cycle or path that we augmented. So by augmenting along all paths and cycles in D, we obtain M_2.

Pivotal players and marginal contribution. If, for an unweighted matching game (n, v), a player $i \in [n]$, and a coalition $P \subseteq [n]\backslash\{i\}$, it holds that $v(P \cup \{i\}) = v(P)+1$, then we say that player i is *pivotal* (for coalition P, in game (n, v)). Similarly, if $\sigma : [n] \to [n]$ is a permutation on $[n]$, and i is pivotal for set of players $p(i, \sigma) = \{i' \mid \sigma^{-1}(i') < \sigma^{-1}(i)\}$ (i.e., the players occurring before i in σ), then we say that σ is pivotal for i.

For the general case of weighted matching games, when P is a coalition not containing player i, we refer to the value $v(P \cup \{i\}) - v(P)$ as *the marginal contribution of i to P*. When σ is a permutation on $[n]$, we refer to the value $v(p(i, \sigma) \cup \{i\}) - v(p(i, \sigma))$ as *the marginal contribution of i to σ*.

For an unweighted matching game, the raw Shapley value of a player is thus equal to the number of pivotal permutations.

The computational problem. We are interested in the following computational problem.

SHAPLEY
Instance: A weighted graph $(G = ([n], E), w)$ and a specified player $i \in [n]$.
Question: Compute $\varphi_i(MG(G))$.

General Insights

In this section, we gain some general insights about the Shapley value of matching games. First, if the graph is not connected, then the problem of computing the Shapley value of the graph reduces to computing the Shapley value of the respective connected components.

Lemma 222 (Shapley value in connected components). *Let $(G = ([n], E), w)$ be a weighted graph with k connected components for some $k \in \mathbb{N}_{\geq 1}$, and let the respective vertex sets of these connected components be $P_1, \dots, P_k$. Let v be the characteristic function of the matching game $MG(G)$ on that graph, and let $c : [n] \to [k]$ be the function that maps a vertex i to the number j such that $j \in P_k$. Then, for every vertex i it holds that $\varphi_i(v) = \varphi_i(v_{c(i)})$, where v_j denotes the characteristic function of the matching game on the subgraph induced by P_j, for $j \in [k]$.*

Proof. We prove this for $k = 2$. For $k > 2$, the claim then holds by straightforward induction.

Therefore, let P_1 and P_2 be the vertex sets of the two connected components of G, and let (n, v'_{P_1}) and (n, v'_{P_2}) denote the matching game obtained by removing from the graph all edges among vertices in respectively P_1 and P_2. Note that v is the sum of v'_{P_1} and v'_{P_2}.

By the additivity property of the Shapley value, it thus holds for every player i that $\varphi_i(v) = \varphi_i(v'_{P_1}) + \varphi_i(v'_{P_2})$. It is therefore sufficient to show that $\varphi_i(v'_{P_1}) = \varphi_i(v_{P_2})$ for all $i \in P_1$ and that $\varphi_i(v'_{P_2}) = \varphi_i(v_{P_1})$ for all $i \in P_2$. This follows inductively if we prove that for any graph $H = ([m], F)$, for some $m \in \mathbb{N}$, it holds that $\varphi_i(v_{MG(H)}) = \varphi_i(v_{MG(([m+1],F))})$ for all $i \in [m]$, which follows from the following derivation:

$$
\begin{aligned}
&\varphi_i(v_{MG(H)}) \\
&= \tfrac{1}{m!} \sum_{P : P \subseteq [m] \setminus \{i\}} |P|!(m - |P| - 1)!(v(P \cup \{i\}) - v(P)) \\
&= \tfrac{1}{(m+1)!} \sum_{P : P \subseteq [m] \setminus \{i\}} (|P| + 1 + m - |P|)|P|!(m - |P| - 1)!(v(P \cup \{i\}) - v(P))
\end{aligned}
$$

$$
\begin{aligned}
&= \tfrac{1}{(m+1)!} \sum_{P:P\subseteq[m]\setminus\{i\}} (|P|+1)!(m-|P|-1)!(v(P\cup\{i\})-v(P)) \\
&\quad + \tfrac{1}{(m+1)!} \sum_{P:P\subseteq[m]\setminus\{i\}} |P|!(m-|P|)!(v(P\cup\{i\})-v(P)) \\
&= \tfrac{1}{(m+1)!} \sum_{P:P\subseteq[m+1]\setminus\{i\},m+1\in P} |P|!(m+1-|P|-1)!(v(P\cup\{i\})-v(P)) \\
&\quad + \tfrac{1}{(m+1)!} \sum_{P:P\subseteq[m+1]\setminus\{i\},j\notin P} |P|!(m+1-|P|-1)!(v(P\cup\{i\})-v(P)) \\
&= \tfrac{1}{(m+1)!} \sum_{P:P\subseteq[m+1]\setminus\{i\}} |P|!(m+1-|P|-1)!(v(P\cup\{i\})-v(P)) \\
&= \varphi_i(v_{MG([m+1],F)}).
\end{aligned}
$$

□

It is rather straightforward to see that a vertex has a Shapley value zero if and only if it is not connected to any other vertex.

Observation 223. A player in a matching game has a non-zero Shapley value if and only if there is an edge (with non-zero weight) in the graph that contains the player. It can thus be decided in linear time whether a player in a matching game has a Shapley value of zero.

Next, we present another lemma concerning the Shapley value of unweighted matching games.

Lemma 224. *Let (n, v) be an unweighted matching game. If for each $s \in [n-1]$, the number of coalitions of size s for which player i is pivotal in (n, v) can be computed in time $f(n)$ for some function $f : \mathbb{N} \to \mathbb{R}_{\geq 0}$, then the Shapley value of i can be computed in time $nf(n)$.*

Proof. Let η_i^s be the number of coalitions of size $s \in [n-1]$ for which a vertex i is pivotal in (n, v).

$$
\begin{aligned}
\varphi_i(n,v) &= \frac{1}{n!} \sum_{P:P\subseteq[n]\setminus\{i\}} |P|!(n-|P|-1)!(v(P\cup\{i\})-v(P)) \\
&= \frac{1}{n!} \sum_{s\in[n-1]} \sum_{\substack{P:P\subseteq[n]\setminus\{i\} \\ |P|=s}} s!(n-s-1)!(v(P\cup\{i\})-v(P))
\end{aligned}
$$

$$= \frac{1}{n!} \sum_{s \in [n-1]} s!(n-s-1)! \sum_{\substack{P: P \subseteq [n] \setminus \{i\}) \\ |P|=s}} (v(P \cup \{i\}) - v(P))$$

$$= \frac{1}{n!} \sum_{s \in [n-1]} s!(n-s-1)! \eta_i^s.$$

Therefore, the problem of computing the Shapley value reduces to computing η_i^s for all $s \in [n-1]$. $\square$

10.4 Exact Algorithms for Restricted Graph Classes

Some classes of matching games for which computing the Shapley value is trivial are symmetric graphs (e.g. cliques and cycles), and graphs with a constant number of vertices. We proceed to prove this for two additional special cases: weighted graphs that admit constant size (co)clique modular decompositions, and unweighted graphs with degree at most two.

10.4.1 Graphs with a Constant Number of Clique or Coclique Modules

An important concept in the context of undirected graphs is that of a *module*. A subset of vertices $P \subseteq [n]$ is a module iff all members of P have the same set of neighbors in $[n] \setminus P$. We can extend this notion to weighted graphs by requiring that all members of P are connected to the same set of neighbors, by edges of the same weight. A *modular decomposition* is a partition of the vertex set into modules.

A *clique module* (resp. *coclique module*) of a weighted graph is a module of which the vertices are pairwise connected by edges of the same weight (resp. pairwise disconnected). Note that every graph has a trivial modular decomposition into cliques (and cocliques): the partition of $[n]$ into singletons.

We prove that if a weighted graph G has a size k modular decomposition consisting of only cliques or only cocliques, for some fixed $k \in \mathbb{N}_{\geq 1}$, then the Shapley value of $MG(G)$ can be found in polynomial time. In fact, we will show that this holds for the more general class of *subgraph-based* games: We call a cooperative game (n, v) *subgraph-based* if there exists a weighted graph $(G = ([n], E), w)$ such that for $P_1, P_2 \subseteq [n]$, it holds that $v(P_1) = v(P_2)$ if the weighted subgraphs of (G, w) induced by P_1 and P_2 are isomorphic.

Theorem 225. *Let $k \in \mathbb{N}_{\geq 0}$. Consider a subgraph-based cooperative game (n, v). Then, the Shapley value of (n, v) can be computed in polynomial time if the following three conditions hold:*

- *the weighted graph $(G = ([n], E), w)$ associated to (n, v) is given, or can be computed in time polynomial in the size of the representation of (n, v).*
- *there exists a modular decomposition $\gamma(G, w)$ into k cocliques or k cliques, and G is unweighted in the latter case (i.e., $w(e) = 1$ for all $e \in E$).*
- *$v(P)$ can be computed in polynomial time for all $P \subseteq [n]$.*

Proof. For (G, w), first note that one can find in polynomial time a minimum cardinality modular decomposition into cocliques: simply check for each pair of vertices whether they are disconnected and connected to identical sets of vertices through edges with identical weights. If so, then they can be put in the same module. Similarly, a minimum cardinality modular decomposition into cliques can be found in polynomial time in case the graph is unweighted, by finding a minimum cardinality modular decomposition into cocliques in the complement of G (i.e., the graph that contains only those edges not in E).

Next, we make use of the definition of a *player type*: Definition 201 given in Chapter 9. We first show that all players in the same module of $\gamma(G, w)$ are of the same player type. Let i, i' be two players in the same module M in $\gamma(G, w)$. Then, for every coalition $P \subseteq [n]\backslash\{i, i'\}$, the weighted subgraphs of (G, w) induced by $P \cup \{i\}$ and $P \cup \{i'\}$ are isomorphic (because the weighted subgraph of (G, w) induced by M is a clique or coclique), so $v(P \cup \{i\}) = v(P \cup \{i'\})$. Therefore, the players can be divided into a constant number k of player types.

Ueda et al. [2011] showed that any cooperative game (n, v) in which the value of a given coalition can be computed in polynomial time, and there is known partition of the players into sets of the same player type of size at most k, then the Shapley value of (n, v) can be computed in polynomial time via dynamic programming. The number of player types in our game is k. Therefore the result of Ueda et al. [2011] can be applied, and this proves our claim. □

For matching games, the function v can be evaluated using any polynomial time maximum weight matching algorithm. Therefore, the above result implies that computing the Shapley value can be done in polynomial time for classes of graphs where we can find efficiently a size k modular decomposition into cliques or cocliques. This includes the class of complete k-partite graphs and any strong product[1] of an arbitrary

[1]The *strong product* of two graphs $G_1 = (V, E_1)$ and $G_2 = (U, E_2)$ is defined as the graph $(V \times U, E')$, where $E' = \{\{(i_V, i_U), (j_V, j_U)\} \subseteq V \times U \mid i_U = j_U \wedge \{i_V, j_V\} \in E_1 \vee \{i_U, j_U\} \in E_2\}$.

size clique (or coclique) with a graph of k vertices.

Corollary 226. *For matching games based on complete k-partite graphs, where k is a constant, the Shapley value can be computed in polynomial time.*

Theorem 225 also applies to cooperative games such as (s,t)-vertex connectivity games and min-cost spanning tree games [Deng and Fang, 2008, Deng et al., 1999], as these are subgraph-based games.

10.4.2 Graphs of Degree at Most Two

We first examine *linear graphs* (or: “paths”), i.e., unweighted connected graphs in which two vertices have out-degree one and the remaining vertices have out-degree two.

Lemma 227. *The Shapley value of a player in a matching game (n,v) on an unweighted linear graph can be computed in $O(n^4)$ time.*

Proof. Assume without loss of generality that the vertex set is $[n]$ and the edge set is $\{\{j,j+1\} \mid j \in [n-1]\}$, and that $i \in [n]$ is the player of which we want to compute the Shapley value. Fix any $s \in [n-1]$, and let η_i^s be the number of coalitions of size s for which i is pivotal. We compute η_i^s by subdividing in separate cases and taking the sum of them:

- The number $\eta_i^{s,\text{left}} = |\{P \cup \{i+1\} \mid P \subseteq [n] \backslash \{i, i-1, i+1\}, i \text{ is pivotal for } P\}|$. Intuitively: the number of coalitions P where i is pivotal such that adding i to P extends the left of a line segment.
- The number $\eta_i^{s,\text{right}} = |\{P \cup \{i-1\} \mid P \subseteq [n] \backslash \{i, i-1, i+1\}, i \text{ is pivotal for } P\}|$.
- The number $\eta_i^{s,\text{connect}} = |\{P \cup \{i-1, i+1\} \mid P \subseteq [n] \backslash \{i, i-1, i+1\}, i$ is pivotal for $P\}|$. Intuitively: the number of coalitions P where i is pivotal, such that i connects two line segments.
- $\eta_i^{s,\text{isolated}} = |\{P \mid P \subseteq [n] \backslash \{i, i-1, i+1\}, i \text{ is pivotal for } P\}|$.

It is immediate that $\eta_i^{s,\text{isolated}} = 0$, since adding i to a coalition P not containing $i+1$ nor $i-1$ results in a coalition forming a subgraph in which i is an isolated vertex. For the remaining three values, $\eta_i^{s,\text{left}}$, $\eta_i^{s,\text{right}}$, and $\eta_i^{s,\text{connect}}$, we show below how to compute them efficiently.

- For $\eta_i^{s,\text{left}}$, observe that adding a vertex to the left of a (non-empty) line segment L increases the cardinality of a maximum matching if and only if L has an even number of edges (and thus an odd number of vertices). Therefore, define $\eta_i^{s,\text{left}}(k)$ to be the number of coalitions P of size s for which i is pivotal such that P contains the line segment $\{i+1,\ldots,i+k+1\}$, and does not contain $\{i-1,i+k+2\}$. The number $\eta_i^{s,\text{left}}(k)$ is easy to determine:

$$\eta_i^{s,\text{left}}(k) = \begin{cases} 0 & \text{if } k \text{ is odd,} \\ \binom{|[n]\setminus\{i-1,\ldots,i+k+2\}|}{s-|\{i-1,\ldots,i+k+1\}\cap[n]|} & \text{otherwise.} \end{cases}$$

  We can then express $\eta_i^{s,\text{left}}$ as $\sum_{k\in[\max\{n-i-1,s-1\}]} \eta_i^{s,\text{left}}(k)$. There is only a linear number of terms in this sum, and all of them can be computed in linear time.

- $\eta_i^{s,\text{right}}$ is computed in an analogous fashion.

- For $\eta_i^{s,\text{connect}}$, observe that adding a vertex i to a coalition such that i connects two line segments L_1 and L_2, increases the cardinality of a maximum matching if and only if L_1 and L_2 do not both have an odd number of edges (or equivalently: not both have an even number of vertices). Therefore, define $\eta_i^{s,\text{connect}}(k_1,k_2)$ to be the number of coalitions P of size s for which i is pivotal such that P contains the line segments $\{i-k_1-1,\ldots,i-1\}$ and $\{i+1,\ldots,i+k_2+1\}$, and does not contain $\{i-k_1-2,i+k_2+2\}$. The number $\eta_i^{s,\text{connect}}(k_1,k_2)$ is easy to determine:

$$\eta_i^{s,\text{connect}}(k_1,k_2) = \begin{cases} 0 & \text{if } k_1 \text{ and } k_2 \text{ are both odd,} \\ \binom{|[n]\setminus(\{i-k-2,\ldots,i+k+2\}|}{s-|\{i-k-1,\ldots,i+k+1\}\cap[n]|} & \text{otherwise.} \end{cases}$$

  We can then express $\eta_i^{s,\text{connect}}$ as

$$\sum_{k_1\in[\max\{i-2,s-1\}]} \sum_{k_2\in[\max\{n-i-1,s-k_1-2\}]} \eta_i^{s,\text{left}}(k_1,k_2).$$

  The number of terms in this sum is quadratic, and all of these terms can be computed in linear time. We can thus compute $\eta_i^{s,\text{connect}}$ in $O(n^3)$ time.

The claim now follows from Lemma 224. $\square$

Theorem 228. *The Shapley value of a matching game on a unweighted graph with maximum degree* 2 *can be computed in polynomial time.*

Proof. An unweighted graph with degree at most two is a disjoint union of cycles and linear graphs. From Lemma 222, we can compute the Shapley value of the connected components separately. From Lemma 227, we know that the Shapley value of linear graphs can be computed in polynomial time. Due to the anonymity property of the Shapley value, the Shapley value of a cycle is uniform. □

The above proof for linear graphs demonstrates that computation of the Shapley value of a matching game already becomes involved for even the simplest of graph structures. We would be interested in seeing an extension of this result that enables us to exactly compute the Shapley value in *any* non-trivial class of graphs that contains a vertex of degree at least three.

10.5 Computational Complexity of the General Problem

In this section, we examine the computational complexity of the general problem of computing the Shapley value for matching games. As we mentioned in Section 10.3, SHAPLEY is equivalent to the problem of counting the number of pivotal permutations for a player in an unweighted matching game, and is therefore a counting problem. It is moreover easy to see that this counting problem is a member of the complexity class $\#\mathsf{P}$.[2]

For certain cooperative games such as weighted voting games [Elkind et al., 2009], intractability of computing the Shapley value can be established by proving that even checking whether a player gets non-zero Shapley value is NP-complete. Proposition 223 tells us that this is not the case for matching games. Before we proceed, we establish some notation. Let $G = ([n], E)$ be a graph. Let $\alpha_k(G)$ be the number of vertex sets $P \subseteq [n]$ such that $|P| = k$ and the subgraph $G(P)$ of G induced by P admits a perfect matching. Then $\overline{\alpha_k}(G) = \binom{n}{k} - \alpha_k(G)$ is the number of subsets $P \subseteq [n]$ of size k such that $G(P)$ does not admit a perfect matching. In order to characterize the complexity of SHAPLEY, we first define the following problem.

#MATCHABLESUBGRAPHS$_k$
Instance: Undirected and unweighted graph $G = ([n], E)$ and a number $k \in \mathbb{N}_{\geq 0}$.
Question: Compute $\alpha_k(G)$.

[2]Informally: $\#\mathsf{P}$ is the class of computational problems that correspond to counting the number of accepting paths on a non-deterministic Turing machine. We refer the reader to any introductory text on complexity theory.

Lemma 229. #MATCHABLESUBGRAPHS$_k$ *is* $\#\mathsf{P}$*-complete.*

Proof. In Colbourn et al. [1995] it is proved that the following problem is $\#\mathsf{P}$-complete: Given an undirected and unweighted bipartite graph $G = (S \cup I, E)$, compute the number of subsets $B \subseteq S$, such that $G(B \cup I)$ admits a perfect matching.[3] The problem is equivalent to #MatchableSubgraphs$_{2|I|}$. □

Theorem 230. *Computing the Shapley value of a matching game on an unweighted graph is* $\#\mathsf{P}$*-complete.*

Proof. We give a polynomial time Turing reduction from #MATCHABLESUBGRAPHS$_k$ to SHAPLEY. We show that if there exists a polynomial time algorithm for SHAPLEY, then we can solve #MATCHABLESUBGRAPHS$_k$ for a given graph G in polynomial time, by solving SHAPLEY on a set of graphs that we construct from G. For each of these graphs, we show that a linear equation holds that relates the Shapley value of a vertex of G to the values α_k and $\overline{\alpha_k}$. The coefficient matrix of the resulting system of linear equations will then turn out to be invertible, hence it can be solved in polynomial time via Gaussian elimination in order to compute the values α_k and $\overline{\alpha_k}$. We remind the reader that the symbol κ is used to denote the raw Shapley value, as defined in Section 10.3.

Let $G = ([n], E)$ be the given graph, and let G_0 be the graph in which a new vertex y_0 is added to G that is connected to all vertices in $[n]$. For $i > 0$, let G_i be G_0 with i additional vertices $y_1, y_2, \ldots, y_i$ and i additional edges $\{\{y_j, y_{j-1}\} \mid j \in [i]\}$.

The first part of the proof consists of showing that the following set of equations hold:

$$\kappa_{y_i}(MG(G_i)) = \begin{cases} C(i) + \sum_{k=0}^{n}(k+i)!(n-k)!\overline{\alpha_k}(G) & \text{if } i \text{ is even,} \quad (10.1) \\ C(i) + \sum_{k=0}^{n}(k+i)!(n-k)!\alpha_k(G) & \text{if } i \text{ is odd,} \quad (10.2) \end{cases}$$

where

$$C(i) = \sum_{k \in [\lfloor i/2 \rfloor]} \sum_{j=0}^{n+1-2k} (j+2k-1)!(n+i-j-2k+1)!\binom{n+i-2k}{j}.$$

Define a *type* 1 *pivotal coalition for* y_i *in* $MG(G_i)$ as a pivotal coalition for i in $MG(G_i)$ that *does not* contain all players $y_0, \ldots, y_{i-1}$. Define a *type* 2 *pivotal coalition for* y_i *in* $MG(G_i)$ as a pivotal coalition for y_i in $MG(G_i)$ that *does* contain

[3] The proof of Colbourn resolved “an exceptionally difficult problem” [Colbourn et al., 1995]. Interestingly, the corresponding decision problem of checking whether there exists a subgraph of size k that does not admit a perfect matching, appears to be open.

all players $y_0, \dots, y_{i-1}$. Denote by $H_i^{\text{type 1}}(s)$ (resp. $H_i^{\text{type 2}}(s)$) the set of type 1 (resp. type 2) pivotal coalitions for i in $MG(G_i)$ that are of size s. From (1.9), Definition 31, it follows that

$$\kappa_{y_i}(MG(G_i)) = \sum_{s \in [n+i]} s!(n+i-s)!|H_i^{\text{type 1}}(s)| + \sum_{s \in [n+i]} s!(n+i-s)!|H_i^{\text{type 2}}(s)|. \tag{10.3}$$

First we characterize the coalitions in $H_i^{\text{type 2}}(s)$.

Lemma 231. *If i is even, a coalition P of $MG(G_i)$ is in $H_i^{\textit{type 2}}(s)$ if and only if $G(P \cap [n])$ is not perfectly matchable (and $\{y_0, \dots, y_{i-1}\} \subseteq P, |P| = s$). If i is odd, a coalition P of $MG(G_i)$ is in $H_i^{\textit{type 2}}(s)$ if and only if $G(P \cap [n])$ is perfectly matchable (and $\{y_0, \dots, y_{i-1}\} \subseteq P, |P| = s$).*

Proof. Case of even i. ($\Rightarrow$) Let M be a maximum matching for $G_i(P)$. P is pivotal, so M is not a perfect matching. We can assume though, that all vertices $\{y_0, \dots, y_{i-1}\}$ are matched to each other in the matched graph $(G_i(P), M)$, because $G_i(\{y_0, \dots, y_{i-1}\})$ is a linear graph with an even number of vertices, and is thus perfectly matchable. It follows that the exposed nodes of $(G_i(P), M)$ are all in $[n]$, and therefore the matching M restricted to $[n]$ is a maximum matching for $G(P \setminus \{y_0, \dots, y_{i-1}\}) = G(P \cap [n])$ that is non-perfect.

($\Leftarrow$) Let M be a maximum (non-perfect) matching for $G(P \cap [n])$ and let y be an exposed vertex of $(G(P \cap [n]), M)$. Then $M' = M \cup \{\{y_{i'}, y_{i'+1}\} \mid i' \text{ even } \wedge i' < i\}$ is a maximum matching for $G_i(P)$, by Berge's lemma (Lemma 221), as it is clear that there is no augmenting path in $(G_i(P), M')$. Moreover, observe that in $(G_i(P), M')$ there is an even-length alternating path from y to y_{i-1}. Therefore, there is in (G_i, M') an augmenting path from y to y_i, and it follows again by Berge's lemma that P is pivotal.

Case of odd i. ($\Rightarrow$) Let M' be a maximum matching for $G_i(P)$. Coalition P is pivotal, so in $(G_i(P), M')$ there is an even-length alternating path $\mathcal{P}$ from an exposed node y to y_{i-1}. Obtain the matching M by augmenting M' along $\mathcal{P}$. Matching M is then a maximum matching for $G_i(P)$ in which y_{i-1} is exposed. $G_i(\{y_0, \dots, y_{i-1}\})$ is a linear graph and M is maximum, so it follows that y_{i-1} is the only exposed node in $(G_i(P), M)$ among $\{y_0, \dots y_{i-1}\}$. Therefore $P \cap [n]$ must be matched to each other in $(G(P), M)$ (for otherwise, in $(G_i(P), M)$ there would be an augmenting path from y_{i-1} to an exposed node of $P \cap [n]$, contradicting the fact that M is a maximum matching for $G_i(P)$). It follows that $G(P \cap [n])$ is perfectly matchable.

($\Leftarrow$) Let M be a maximum perfect matching for $G(P \cap [n])$. Let M' be a maximum matching for $G_i(\{y_0, \dots, y_{i-1}\})$ in which y_{i-1} is the only exposed node. Then $M \cup M'$ is a matching for $G_i(P)$ in which y_{i-1} is the only exposed node. $M \cup M'$ is clearly a maximum matching, and in $(G_i, M \cup M')$ the edge $\{y_{i-1}, y_i\}$ is exposed. So P is pivotal. $\square$

From the above lemma, it follows that the coalitions in $H_i^{\text{type 2}}(s)$ are precisely the coalitions of the form $P \cup \{y_0, \ldots, y_{i-1}\}$, where $P \subset [n]$ is such that for even i, $G(P)$ is not perfectly matchable, and for odd i, $G(P)$ is perfectly matchable. Therefore $|H_i^{\text{type 2}}(s)| = \overline{\alpha_{s-i}}(G)$ for even i and $|H_i^{\text{type 2}}(s)| = \alpha_{s-i}(G)$ for odd i, and this implies:

$$\sum_{s \in [n+i]} s!(n+i-s)!|H_i^{\text{type 2}}(s)| = \begin{cases} \sum_{k=0}^{n}(k+i)!(n-k)!\overline{\alpha_k}(G) & \text{if } i \text{ is even,} \\ \sum_{k=0}^{n}(k+i)!(n-k)!\alpha_k(G) & \text{if } i \text{ is odd.} \end{cases}$$

In words: the second summation of (10.3) equals the summation of (10.1) when i is even, and the summation of (10.2) when i is odd. Therefore, it suffices to prove that the first summation of (10.3) equals $C(i)$.

For this sake, define $H_i^{\text{type 1}}(s,k)$ for $k \in [\lfloor i/2 \rfloor]$ as $\{P \in H_i^{\text{type 1}}(s) \mid y_{i-2k} \notin P \wedge \{y_{i-1}, \ldots, y_{i-2k+1}\} \subseteq P\}$. Observe that $\{H_i^{\text{type 1}}(s,1), \ldots, H_i^{\text{type 1}}(s, i/2)\}$ is a partition of $H_i^{\text{type 1}}(s)$. For a given k and s, note that the set $H_i^{\text{type 1}}(s,k)$ consists of all coalitions of the form $P \cup \{y_{i-1}, \ldots, y_{i-2k+1}\}$, where $P \subseteq [n] \cup \{y_0, \ldots, y_{i-2k-1}\}$, $|P| = s - 2k + 1$. Hence, $|H_i^{\text{type 1}}(s,k)| = \binom{n+i-2k}{s-2k+1}$ (defining $\binom{a}{b} = 0$ whenever $b < 0$ or $b > a$). Therefore:

$$\sum_{s \in [n+i]} s!(n+i-s)!|H_i^{\text{type 1}}(s)| = \sum_{k \in [\lfloor i/2 \rfloor]} \sum_{s=2k-1}^{n+i-1} s!(n+i-s)!\binom{n+i-2k}{s-2k+1}$$
$$= \sum_{k \in [\lfloor i/2 \rfloor]} \sum_{j=0}^{n+i-2k} (j+2k-1)!(n+i-j-2k+1)!\binom{n+i-2k}{j}.$$

This shows that (10.1) and (10.2) hold.

The second part of the proof consists of showing that all $\alpha_k(G), k \in [n]$ can be computed from $\kappa_{y_i}(MG(G_i))$ in polynomial time, using (10.1) and (10.2), for $i \in [n] \cup \{0\}$. This is sufficient to complete the proof, because the graphs $G_0, \ldots, G_n$ can clearly be constructed from G in polynomial time, hence a polynomial time algorithm that computes α_k from $\kappa_{y_i}(MG(G_i)), i \in [n]$ yields a polynomial time Turing reduction.

Let $\beta_i(G) = \alpha_i(G)$ for even i and let $\beta_i(G) = \overline{\alpha_i}(G)$ for odd i. We can represent (10.1) and (10.2) for $i \in [n] \cup \{0\}$ as the following system of equations:

$$\begin{pmatrix} 0!n! & 1!(n-1)! & \cdots & n!0! \\ 1!n! & & \cdots & (n+1)!0! \\ \vdots & \vdots & \ddots & \vdots \\ n!n! & & \cdots & (2n)!0! \end{pmatrix} \times \begin{pmatrix} \beta_0(G) \\ \beta_1(G) \\ \vdots \\ \beta_n(G) \end{pmatrix} = \begin{pmatrix} \kappa_{y_0}(MG(G_0)) - C(0) \\ \kappa_{y_1}(MG(G_1)) - C(1) \\ \vdots \\ \kappa_{y_n}(MG(G_n)) - C(n) \end{pmatrix} \tag{10.4}$$

Denote by A the $(n+1)\times(n+1)$ matrix in the above equation. Recall that a scalar multiplication of a column by a constant c multiplies the determinant by c. Therefore, A is non-singular if and only if non-singularity also holds for the $(n+1) \times (n+1)$ matrix B, defined by $B_{ij} = (i+j)!$. B is a matrix that is related to Pascal's triangle, and it is known that its determinant is equal to $\prod_{i=0}^{n} i!^2 \neq 0$ [Bacher, 2002, Aziz et al., 2009]. It follows that A is nonsingular, so our system of equations (10.4) is linearly independent and has a unique solution. Note that all entries in the system can be computed in polynomial time (assuming that the Shapley value of a matching game is polynomial time computable): The constants $C(i)$ consist of polynomially many terms, and all factorials and binomial coefficients that occur in (10.4) are taken over numbers of magnitude polynomial in n.

Therefore, we can use Gaussian elimination to solve (10.4) in $O(n^3)$ time. It follows that for all $i \in [n]$, $\beta_i(G)$ can be computed in polynomial time, and hence $\alpha_i(G)$ can be computed in polynomial time. Therefore, if there exists an algorithm that solves SHAPLEY in polynomial time, then it can be used to solve #MATCHABLESUBGRAPHS$_k$ in polynomial time. □

10.6 An Approximation Algorithm

In this section, we show that although computing exactly the Shapley value of matching games is a hard problem, approximating it is much easier.

Let Σ be a finite alphabet in which we agree to describe our problem instances and solutions. A *fully polynomial time randomized approximation scheme (FPRAS)* for a function $f : \Sigma^* \to \mathbb{Q}$ is an algorithm that takes input $x \in \Sigma^*$ and a parameter $\epsilon \in \mathbb{Q}_{>0}$, and returns with probability at least $\frac{3}{4}$ a number in between $f(x)/(1+\epsilon)$ and $(1+\epsilon)f(x)$. Moreover, an FPRAS is required to run in time polynomial in the size of x and $1/\epsilon$. The probability of $\frac{3}{4}$ is chosen arbitrarily: by a standard amplification technique, it can be replaced by an arbitrary number $\delta \in [0,1]$. The resulting algorithm would then run in time polynomial in $n, 1/\epsilon$, and $\log(1/\delta)$.

We will now formulate an algorithm that approximates the raw Shapley value of a player in a weighted matching game, and show that it is an FPRAS. Note that we cannot utilize approximation results in [Liben-Nowell et al., 2011] and [Bachrach et al., 2010a] since matching games are neither convex nor simple. Our FPRAS is based on Monte Carlo sampling, and works as follows: Let $((G = ([n], E), w), i, \epsilon)$ be the input, where (G, w) is the weighted graph representing matching game $MG(G, w)$, $i \in [n]$ is a player in $MG(G)$, and ϵ is the precision parameter. For notational convenience, we write κ_i as a shorthand for $\kappa_i(MG(G))$. The algorithm first determines whether $\kappa_i = 0$ (using Observation 223). If so, then it outputs 0 and terminates. If not, then it samples $\lceil 4n^2(n-1)^2/\epsilon^2 \rceil$ permutations of the player set uniformly at random. Denote this multiset of sampled permutations by $\mathcal{P}$. The algorithm then outputs the average marginal contribution of player i over the permutations in $\mathcal{P}$, and terminates. Note that this average marginal contribution is efficiently computable: it is given by $1/\lceil 4n^2(n-1)^2/\epsilon^2 \rceil$ times the sum of the marginal contributions of player i to each of the sampled permutations. Determining these marginal contributions can be done in polynomial time, using any maximum weight matching algorithm. Denote our sampling algorithm by MATCHINGGAME-SAMPLER.

MATCHINGGAME-SAMPLER resembles the algorithms given in [Mann and Shapley, 1960, Liben-Nowell et al., 2011]: the differences are that the algorithm takes a different number of samples, and that it determines whether the Shapley value of player i is 0 prior to running the sampling procedure. Moreover, its proof of correctness requires different insights.[4]

Theorem 232. MATCHINGGAME-SAMPLER *is an FPRAS for computing the raw Shapley value in a weighted matching game.*

Proof. Let $((G = ([n], E), w), i, \epsilon)$ be the input. Denote by $\bar{\kappa}_i$ the output of the algorithm. If $\kappa_i = 0$, then MATCHINGGAME-SAMPLER is guaranteed to output the right solution, so assume that $\kappa_i > 0$. Let $w_i^{\max}$ be the maximum weight among the edges attached to i, and let $e_i^{\max} \in E$ be an edge that is attached to i such that $w(e_i^{\max}) = w_i^{\max}$. Let X be a random variable that takes the value of $n!$ times the marginal contribution of player i in a uniformly randomly sampled permutation of the players. Note that $\mathbf{E}[X] = \kappa_i$. Note that the marginal contribution of a player in any permutation is at most $w_i^{\max}$, so X is at most $w_i^{\max} n!$.

Let j be the neighbor of i connected by $e_i^{\max}$. Observe that any permutation in which j is positioned first, and i is positioned second, is a permutation for i in which the marginal contribution of i is $w_i^{\max}$. There are $(n-2)!$ such permutations, so the raw Shapley value κ_i of i is at least $w_i^{\max}(n-2)!$. For the variance of X we derive that $\mathbf{Var}[X] = \mathbf{E}[X^2] - \mathbf{E}[X]^2 \leq \mathbf{E}[X^2] \leq (w_i^{\max})^2 n!^2 \leq n^2(n-1)^2\kappa_i^2$.

[4] To be precise, this applies only to [Liben-Nowell et al., 2011]. For the sampling algorithm in Mann and Shapley [1960], no proof or approximation-quality analysis of any kind is given.

Observe that $\bar{\kappa}_i$ is a random variable that is equal to $\frac{\sum_{j \in [\lceil 4n^2(n-1)^2/\epsilon^2 \rceil]} X_j}{\lceil 4n^2(n-1)^2/\epsilon^2 \rceil}$, where X_j are independent random variables with the same distribution as X. From this we obtain that $\mathbf{E}[\bar{\kappa}_i] = \mathbf{E}[X] = \kappa_i$. The desired approximation guarantee then follows from Chebyshev's inequality,[5] and completes the proof:

$$\begin{aligned} \mathbf{Pr}[|\bar{\kappa}_i - \kappa_i| \geq \epsilon\kappa_i] &\leq \frac{\mathbf{Var}[\bar{\kappa}_i]}{\epsilon^2\kappa_i^2} = \frac{\mathbf{Var}\left[\frac{1}{\lceil 4n^2(n-1)^2/\epsilon^2 \rceil} \sum_{j \in [\lceil 4n^2(n-1)^2/\epsilon^2 \rceil]} X_j\right]}{\epsilon^2\kappa_i^2} \\ &= \frac{\left(\frac{\mathbf{Var}[X]}{\lceil 4n^2(n-1)^2/\epsilon^2 \rceil}\right)}{\epsilon^2\kappa_i^2} \leq \frac{n^2(n-1)^2\kappa_i^2}{(4n^2(n-1)^2/\epsilon^2) \cdot \epsilon^2\kappa_i^2} \leq \frac{1}{4}. \end{aligned}$$

□

Corollary 233. *The algorithm that runs* MATCHINGGAME-SAMPLER *and returns its output scaled down by $1/n!$, is an FPRAS for computing the Shapley value of a weighted matching game.*

Observe that MATCHINGGAME-SAMPLER is an FPRAS in the strong sense that its running time does not depend on the weights of the edges. Due to the #P-completeness result stated in Theorem 229, this FPRAS is the best one can hope for, and provides us (based on our best judgment) with a complete answer to the precise complexity of this problem.

10.7 Conclusions

In this chapter, we examined the structure, algorithms, and computational complexity for the problem of computing the Shapley value in a matching game. There are many special cases of the problem that have not been treated in this chapter, but nonetheless are potentially worthwhile to analyze: trees, bipartite graphs, connected regular graphs, and series-parallel graphs. Among these, bipartite graphs are especially interesting, since they model two-sided markets. One may pursue the same questions for *fractional matching games* in which the value of a coalition is the maximum size of a fractional matching [Chen et al., 2012]. Moreover, our study motivates exploring the connections with some objects in matching theory. The matching polytope is one of the most well-studied objects in polyhedral combinatorics [Plummer, 1992]. It will be interesting to identify any relation between the matching polytope of a graph and

[5] Here, one could also choose to apply Hoeffding's inequality instead of Chebyshev's inequality, but this will not result in an asymptotically better bound.

the Shapley value of the corresponding matching game. Secondly, network flows are fundamentally connected to matchings for the case of bipartite graphs. An interesting research direction is to explore the connection of network flow games [Kalai and Zemel, 1982] with matching games and check whether computing the Shapley values for these game classes reduce to each other.

Bibliography

I. Abraham, M. Babaioff, S. Dughmi, and T. Roughgarden. Combinatorial auctions with restricted complements. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 3–16. ACM, 2012. 167, 169, 171, 197

H. Ackermann, H. Röglin, and B. Vöcking. Pure Nash equilibria in player-specific and weighted congestion games. In P. Spirakis, M. Mavronicolas, and S. C. Kontogiannis, editors, *Proceedings of the 2nd Workshop on Internet and Network Economics*, volume 4286 of *Lecture Notes in Computer Science*, pages 50–61. Springer, 2006. 19, 89, 115

S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact price of anarchy for polynomial congestion games. In W. Thomas and P. Weil, editors, *Proceedings of the 24th Symposium on Theoretical Aspects of Computer Science*, volume 4393 of *Lecture Notes in Computer Science*, pages 218–229. Springer, 2006. 19, 72

J. Alcalde-Unzu and E. Molis. Exchange of indivisible goods and indifferences: The top trading absorbing sets mechanisms. *Games and Economic Behavior*, 73(1):1–16, 2011. 28, 209, 210, 212, 217, 219, 220

A. Alkan and D. Gale. The core of the matching game. *Games and Economic Behavior*, 2(3):203–212, 1990. 247

A. Anagnostopoulos, L. Becchetti, B. de Keijzer, and G. Schäfer. Inefficiency of games with social context. In Berthold Vöcking, editor, *Proceedings of the 6th International Symposium on Algorithmic Game Theory*, volume 8146 of *Lecture Notes in Computer Science*, pages 219–230. Springer, 2013. 117

N. Andelman, M. Feldman, and Y. Mansour. Strong price of anarchy. In *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms*, pages 189–198. ACM, 2007. 70, 71

E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proceedings of the 45th IEEE Annual Symposium on Foundations of Computer Science*, pages 295–304. IEEE Computer Society, 2004. 14, 19

E. Anshelevich, A. Dasgupta, J. M. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008. 19, 166

A. Archer and É. Tardos. Frugal path mechanisms. In *Proceedings of the 13th ACM-SIAM Symposium on Discrete Algorithms*, pages 991–999. ACM, 2002. 148

R. J. Aumann. Acceptable points in games of perfect information. *Pacific Journal of Mathematics*, 10:381–417, 1960. 67

L. Ausubel and P. Cramton. Demand reduction and inefficiency in multi-unit auctions. Technical report, University of Maryland, 2002. 33, 34, 57

B. Awerbuch, Y. Azar, E. F. Grove, M.-Y. Kao, P. Krishnan, and J. S. Vitter. Load balancing in the L_p norm. In *Proceedings of the 36th IEEE Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, 1995. 165

B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on the Theory of Computing*, pages 57–66. ACM, 2005. 19, 71, 115

B. Awerbuch, Y. Azar, Y. Richter, and D. Tsur. Tradeoffs in worst-case equilibria. *Theoretical Computer Science*, 361(2–3):200–209, 2006. 70, 71

H. Aziz. Complexity of comparison of influence of players in simple games. In *Proceedings of the 2nd International Workshop on Computational Social Choice*, pages 61–72, 2008. 238

H. Aziz and B. de Keijzer. Complexity of coalition structure generation. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, pages 191–198. International Foundation for Autonomous Agents and Multiagent Systems, 2011a. 227

H. Aziz and B. de Keijzer. Complexity of coalition structure generation. *CoRR*, abs/1101.1007, 2011b. 227

H. Aziz and B. de Keijzer. Shapley meets shapley. *CoRR*, abs/1307.0332, 2011c. 245

H. Aziz and B. de Keijzer. Housing markets with indifferences: A tale of two mechanisms. In Jörg Hoffmann and Bart Selman, editors, *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, pages 1249–1255. AAAI Press, 2012. 207

H. Aziz and B. de Keijzer. Shapley meets shapley. In N. Portier and E. Mayr, editors, *Proceedings of the 31st International Symposium on Theoretical Aspects of Computer Science*, volume 25 of *LIPIcs*. Schloss Dagtuhl, 2014. 245

H. Aziz, O. Lachish, M. Paterson, and R. Savani. Wiretapping a hidden network. In S. Leonardi, editor, *Proceedings of the 5th International Workshop on Internet and Network Economics*, volume 5929 of *Lecture Notes in Computer Science*, pages 438–446. Springer, 2009. 229, 231, 247, 260

H. Aziz, F. Brandt, and P. Harrenstein. Monotone cooperative games and their threshold versions. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1017–1024. International Foundation for Autonomous Agents and Multiagent Systems, 2010. 229, 230, 243

M. Babaioff, R. Kleinberg, and C. H. Papadimitriou. Congestion games with malicious players. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 103–112. ACM, 2007. 19, 88

M. Babaioff, R. Kleinberg, and C. H. Papadimitriou. Congestion games with malicious players. *Games and Economic Behavior*, 67(1):22–35, 2009. 19

R. Bacher. Determinants of matrices related to the pascal triangle. *Journal de théorie des nombres de Bordeaux*, 14:19–41, 2002. 260

Y. Bachrach and J. S. Rosenschein. Coalitional skill games. In *Proceedings of the 7th International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1023–1030. International Foundation for Autonomous Agents and Multiagent Systems, 2008. 228, 229, 232, 238, 243

Y. Bachrach and J. S. Rosenschein. Power in threshold network flow games. *Autonomous Agents and Multi-Agent Systems*, 18(1):106–132, 2009. 228, 229, 230, 232

Y. Bachrach, E. Markakis, E. Resnick, A. D. Procaccia, J. S. Rosenschein, and A. Saberi. Approximating power indices: theoretical and empirical analysis. *Autonomous Agents and Multi-Agent Systems*, 20:105–122, 2010a. 261

Y. Bachrach, R. Meir, K. Jung, and P. Kohli. Coalitional structure generation in skill games. In M. Fox and D. Poole, editors, *Proceedings of the 24th AAAI Conference on Artificial Intelligence*. AAAI Press, 2010b. 229, 238, 243, 244

R. Banner and A. Orda. Bottleneck routing games in communication networks. *IEEE Journal on Selected Areas in Communication*, 25(6):1173–1179, 2007. 71

T. Bergstrom. Systems of benevolent utility functions. *Journal of Public Economic Theory*, 1(1):71–100, 1999. 87

K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms*, pages 700–709. ACM, 2011. 35, 37, 47, 66

V. Bilò, A. Celi, M. Flammini, and V. Gallotti. Social context congestion games. In A. Kosowski and M. Yamashita, editors, *Proceedings of the 18th International Colloquium Structural Information and Communication Complexity*, volume 6796 of *Lecture Notes in Computer Science*, pages 282–293. Springer, 2011a. 19, 120

V. Bilò, A. Fanelli, M. Flammini, and L. Moscardelli. Graphical congestion games. *Algorithmica*, 61(2):274–297, 2011b. 19, 120

K. Binmore and J. Swierzbinski. Treasury auctions: Uniform or discriminatory? *Review of Economic Design*, 5(4):387–410, 2000. 34

P. Biró, W. Kern, and D. Paulusma. Computing solutions for matching games. *International Journal of Game Theory*, 41(1):75–90, 2011. 247

P. Biró, M. Bornhoff, P. A. Golovach, W. Kern, and D. Paulusma. Solutions for the stable roommates problem with payments. *Theoretical Computer Science*, 2013. 247

A. Blum, M. T. Hajiaghayi, K. Ligett, and A. Roth. Regret minimization and the price of total anarchy. In *Proceedings of the 40th Annual ACM Symposium on the Theory of Computing*, pages 373–382. ACM, 2008. 88, 89, 94, 124

L. Blumrosen and S. Dobzinski. Welfare maximization in congestion games. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 52–61. ACM, 2006. 19, 27, 165, 166, 168, 194, 196

U. Brandes and T. Erlebach, editors. *Network Analysis*, volume 3418 of *Lecture Notes in Computer Science*. Springer, 2005. 247

F. Brandt and G. Weiß. Antisocial agents and Vickrey auctions. In J.-J. Ch. Meyer and M. Tambe, editors, *Proceedings of the 8th Workshop on Agent Theories, Architectures and Languages*, volume 3222 of *Lecture Notes in Computer Science*, pages 335–347. Springer, 2002. 144

F. Brandt, T. Sandholm, and Y. Shoham. Spiteful bidding in sealed-bid auctions. In *Proceedings of the 20th International Joint Conference on Artifical intelligence*, pages 1207–1214. AAAI Press, 2007. 88, 120, 144, 145, 147

M. Bresky. Pure equilibrium strategies in multi-unit auctions with private value bidders. Technical Report 376, CERGE Economics Institute, 2008. 34

I. Brocas. Endogenous entry in auctions with negative externalities. *Theory and Decision*, 54(2):125–149, 2003. 145

R. Buehler, Z. Goldman, D. Liben-Nowell, Y. Pei, J. Quadri, A. Sharp, S. Taggart, T. Wexler, and K. Woods. The price of civil society. In N. Chen, E. Elkind, and E. Koutsoupias, editors, *Proceedings of the 7th International Workshop on Internet and Network Economics*, volume 7090 of *Lecture Notes in Computer Science*, pages 375–382. Springer, 2011. 120

C. Busch and R. Kannan. Stretch in bottleneck games. In J. Gudmundsson, J. Mestre, and T. Viglas, editors, *Proceedings of the 18th Annual International Computing and Combinatorics Conference*, volume 7434 of *Lecture Notes in Computer Science*, pages 592–603. Springer, 2012. 71

C. Busch and M. Magdon-Ismail. Atomic routing games on maximum congestion. *Theoretical Computer Science*, 410(36):3337–3347, 2009. 71

I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, and E. Papaioannou. The impact of altruism on the efficiency of atomic congestion games. In M. Wirsing, M. Hofmann, and A. Rauschmayer, editors, *Proceedings of the 5th Symposium on Trustworthy Global Computing*, volume 6084, pages 179–189. Springer, 2010. 19, 86, 88, 89, 90, 91, 99, 103, 104, 106, 114, 120

I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 81–90. ACM, 2011. 121

D. Chakrabarty, A. Mehta, V. Nagarajan, and V. V. Vazirani. Fairness and optimality in congestion games. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, pages 52–57. ACM, 2005. 19, 165, 166

N. Chen, P. Lu, and H. Zhang. Computing the nucleolus of matching, cover and clique games. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*. AAAI Press, 2012. 262

P.-A. Chen. *The Effects of Altruism and Spite on Games*. PhD thesis, University of Southern California, 2011. 144

P.-A. Chen and D. Kempe. Altruism, selfishness, and spite in traffic routing. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 140–149. ACM, 2008. 19, 86, 88, 89, 109, 120, 145

P.-A. Chen, M. David, and D. Kempe. Better vaccination strategies for better people. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, pages 179–188. ACM, 2010. 86, 88, 89, 115, 145

P.-A. Chen, B. de Keijzer, D. Kempe, and G. Schäfer. The robust price of anarchy of altruistic games. In N. Chen, E. Elkind, and E. Koutsoupias, editors, *Proceedings of the 7th Workshop on Internet and Network Economics*, volume 7090 of *Lecture Notes in Computer Science*, pages 383–390. Springer, 2011a. 85, 120

P.-A. Chen, B. de Keijzer, D. Kempe, and G. Schäfer. The robust price of anarchy of altruistic games. *CoRR*, abs/1112.3680, 2011b. 85, 120

G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th ACM Symposium on the Theory of Computing*, pages 67–73. ACM, 2005a. 19, 71, 72, 99, 115

G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In G. S. Brodal and S. Leonardi, editors, *Proceedings of the 13th Annual European Symposium on Algorithms*, volume 3669 of *Lecture Notes in Computer Science*, pages 59–70. Springer, 2005b. 19

G. Christodoulou, A. Kovács, and M. Schapira. Bayesian combinatorial auctions. In L. Aceto, I. Damgård, L. A. Goldberg, M. M. Halldórsson, A. Ingólfsdóttir, and I. Walukiewicz, editors, *Proceedings of the 35t International Colloquium on Automata, Languages and Programming*, volume 5125 of *Lecture Notes in Computer Science*, pages 820–832. Springer, 2008. 35, 37, 47, 66

C. J. Colbourn, J. S. Provan, and D. Vertigan. The complexity of computing the Tutte polynomial on transversal matroids. *Combinatorica*, 15(1):1–10, 1995. 257

R. Cole, J. R. Correa, V. Gkatzelis, V. Mirrokni, and N. Olver. Inner product spaces for minsum coordination mechanisms. In *Proceedings of the 43rd ACM Symposium on the Theory of Computing*, pages 539–548. ACM, 2011. 132, 133

V. Conitzer and T. Sandholm. Expressive negotiation over donations to charities. In *Proceedings of the 5th ACM conference on Electronic Commerce*, pages 51–60. ACM, 2004. 166

V. Conitzer and T. Sandholm. Expressive negotiation in settings with externalities. In *Proceedings of the 20th AAAI Conference on Artificial Intelligence - Volume 1*, pages 255–260. AAAI Press, 2005. 166

V. Conitzer and T. Sandholm. Complexity of constructing solutions in the core based on synergies among coalitions. *Artificial Intelligence*, 170:607–619, 2006. 237

J. R. Correa and M. Queyranne. Efficiency of equilibria in restricted uniform machine scheduling with total weighted completion time as social cost. *Naval Research Logistics*, 59(5):384–395, 2012. 132

A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. *ACM Transactions on Algorithms*, 3(1), 2007. 70, 71

X. Deng and Q. Fang. Algorithmic cooperative game theory. In A. Chinchuluun, P. M. Pardalos, A. Migdalas, and L. Pitsoulis, editors, *Pareto Optimality, Game Theory And Equilibria*, volume 17 of *Springer Optimization and Its Applications*, pages 159–185. Springer, 2008. 229, 231, 247, 254

X. Deng and C. H. Papadimitriou. On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, 12(2):257–266, 1994. 229, 231, 241, 247

X. Deng, T. Ibaraki, and H. Nagamochi. Algorithmic aspects of the core of combinatorial optimization games. *Mathematics of Operations Research*, 24(3):751–766, 1999. 246, 247, 254

S. Dobzinski and S. Dughmi. On the power of randomization in algorithmic mechanism design. In *Proceedings of the 50th IEEE Annual Symposium on Foundations of Computer Science*, pages 505–514. IEEE Computer Society, 2009. 204

eBid. `http://www.uk.ebid.net`, 2013. 34

J. Edmonds. Paths, trees and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965. 249

J. Elias, F. Martignon, K. Avrachenkov, and G. Neglia. Socially-aware network design games. In *Proceedings of the 29th IEEE Conference on Computer Communications*, pages 41–45. IEEE Computer Society, 2010. 88, 89, 120

E. Elkind, L. A. Goldberg, P. W. Goldberg, and M. J. Wooldridge. Computational complexity of weighted threshold games. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, pages 718–723. AAAI Press, 2007. 228, 229, 230

E. Elkind, L. A. Goldberg, P. W. Goldberg, and M. Wooldridge. On the computational complexity of weighted voting games. *Annals of Mathematics and Artificial Intelligence*, 56(2):109–131, 2009. 247, 256

R. Engelbrecht-Wiggans and C. M. Kahn. Multi-unit auctions with uniform prices. *Economic Theory*, 12(2):227–258, 1998. 34

E. Fehr and K. M. Schmidt. The economics of fairness, reciprocity and altruism: Experimental evidence and new theories. In S.-C. Kolm and J. M. Ythier, editors, *Handbook on the Economics of Giving, Reciprocity and Altruism*, volume 1, chapter 8, pages 615–691. Elsevier, 2006. 120

J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63(1):21–41, 2001. 165

M. Feldman, H. Fu, N. Gravin, and B. Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the 45th ACM Symposium on the Theory of Computing*, pages 201–210. ACM, 2013. 35, 37, 46, 49, 52, 59, 66

A. Fiat, H. Kaplan, M. Levy, and S. Olonetsky. Strong price of anarchy for machine load balancing. In L. Arge, C. Cachin, and T. Jurdzinski, editors, *Proceedings of the 34th International Colloquium on Automata, Languages and Programming*, volume 4596 of *Lecture Notes in Computer Science*, pages 583–594. Springer, 2007. 71

A. Fiat, A.R. Karlin, E. Koutsoupias, and A. Vidali. Approaching utopia: Strong truthfulness and externality-resistant mechanisms. *CoRR*, abs/1303.1646, 2012. 145

D. Fotakis. Stackelberg strategies for atomic congestion games. In L. Arge, M. Hoffmann, and E. Welzl, editors, *Proceedings of the 15th European Symposium on Algorithms*, volume 4698 of *Lecture Notes in Computer Science*, pages 299–310. Springer, 2007. 19

D. Fotakis, V. Gkatzelis, A. C. Kaporis, and P.G. Spirakis. The impact of social ignorance on weighted congestion games. In S. Leonardi, editor, *Proceedings of the*

5th International Workshop on Internet and Network Economics, volume 5929 of *Lecture Notes in Computer Science*, pages 316–327. Springer, 2009. 19, 120

D. Fotakis, P. Krysta, and O. Telelis. Externalities among advertisers in sponsored search. In M. Mavronicolas and M. Papadoupoulou, editors, *Proceedings of the 2nd Symposium on Algorithmic Game Theory*, volume 5814 of *Lecture Notes in Computer Science*, pages 105–116. Springer, 2011. 145

M. Friedman. *A Program for Monetary Stability*. Fordham University Press, 1960. 34

H. Fu, R. Kleinberg, and R. Lavi. Conditional equilibrium outcomes via ascending price processes with applications to combinatorial auctions with item bidding. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, page 586. ACM, 2012. 35

M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. The price of anarchy for restricted parallel links. *Parallel Processing Letters*, 16(1):117–132, 2006. 70, 71, 78

M. L. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979. 195

J. Geanakoplos, D. Pearce, and E. Stacchetti. Psychological games and sequential rationality. *Games and Economic Behavior*, 1(1):60–79, 1989. 89

H. Gintis, S. Bowles, R. Boyd, and E. Fehr. *Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life*. MIT Press, 2005. 89

T. Harks, M. Klimm, and R.H. Möhring. Strong Nash equilibria in games with the lexicographical improvement property. In S. Leonardi, editor, *Proceedings of the 5th Workshop on Internet and Network Economics*, volume 5929 of *Lecturne Notes in Computer Science*, pages 463–470. Springer, 2009. 67, 69, 71

T. Harks, M. Hoefer, M. Klimm, and A. Skopalik. Computing pure Nash and strong equilibria in bottleneck congestion games. In M. de Berg and U. Meyer, editors, *Proceedings of the 18th Annual European Symposium on Algorithms: Part II*, volume 6347 of *Lecture Notes in Computer Science*, pages 29–38. Springer, 2010. 71

A. Hassidim, H. Kaplan, Y. Mansour, and N. Nisan. Non-price equilibria in markets of discrete goods. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 295–296. ACM, 2011. 35, 37, 44

J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182:105–142, 1999. 172

A. Hayrapetyan, É. Tardos, and T. Wexler. The effect of collusion in congestion games. In *Proceedings of the 38th Annual ACM Symposium on the Theory of Computing*, pages 89–98. ACM, 2006. 19, 72, 88

M. Hoefer and A. Skopalik. Stability and convergence in selfish scheduling with altruistic agents. In S. Leonardi, editor, *Proceedings of the 5th International Workshop on Internet and Network Economics*, volume 5929 of *Lecture Notes in Computer Science*, pages 616–622. Springer, 2009a. 120

M. Hoefer and A. Skopalik. Altruism in atomic congestion games. In A. Fiat and P. Sanders, editors, *Proceedings of the 17th European Symposium on Algorithms*, volume 5757 of *Lecture Notes in Computer Science*, pages 179–189. Springer, 2009b. 19, 86, 89, 99, 120

M. Hoefer and A. Skopalik. Social context in potential games. In P. W. Goldberg and M. Guo, editors, *Proceedings of the 8th International Conference on Internet and Network Economics*, volume 7695, pages 364–377. Springer, 2012. 120

R. Hoeksma and M. Uetz. The price of anarchy for minsum related machine scheduling. In R. Solis-Oba and G. Persiano, editors, *Proceedings of the 9th International Workshop on Approximation and Online Algorithms*, volume 7164 of *Lecture Notes in Computer Science*, pages 261–273. Springer, 2012. 132, 133

S. Ieong and Y. Shoham. Marginal contribution nets: A compact representation scheme for coalitional games. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, pages 193–202. ACM, 2005. 247

N. Immorlica, E. Li, V. S. Mirrokni, and A. S. Schulz. Coordination mechanisms for selfish scheduling. *Theoretical Computer Science*, 410(17):1589–1598, 2009. 71

P. Jaramillo and V. Manjunath. The difference indifference makes in strategy-proof allocation of objects. *Journal of Economic Theory*, 147(5):1913–1946, 2011. 28, 209, 210, 212, 217

P. Jehiel, B. Moldovanu, and E. Stacchetti. Multidimensional mechanism design for auctions with externalities. *Journal of Economic Theory*, 85(2):258–294, 1999. 145

E. Kalai and D. Schmeidler. An admissible set occurring in various bargaining situations. *Journal of Economic Theory*, 14:402–411, 1977. 214

E. Kalai and E. Zemel. Generalized network problems yielding totally balanced games. *Operations Research*, 30:998–1008, 1982. 263

V. Kann, S. Khanna, J. Lagergren, and A. Panconesi. On the hardness of approximating max k-cut and its dual. *Chicago Journal of Theoretical Computer Science*, 2, 1997. 166

G. Karakostas and A. Viglas. Equilibria for networks with malicious users. *Mathematical Programming*, 110:591–613, 2007. 88

A. Karlin, D. Kempe, and T. Tamir. Beyond VCG: Frugality of truthful mechanisms. In *Proceedings of the 46th IEEE Annual Symposium on Foundations of Computer Science*, pages 615–624. IEEE Computer Society, 2005. 148

B. de Keijzer and Guido Schäfer. Finding social optima in congestion games with positive externalities. In L. Epstein and P. Ferragina, editors, *Proceedings of the 20th Annual European Symposium on Algorithms*, volume 7501 of *Lecture Notes in Computer Science*, pages 395–406. Springer, 2012. 163

B. de Keijzer, G. Schäfer, and O. A. Telelis. On the inefficiency of equilibria in linear bottleneck congestion games. In S. C. Kontogiannis, E. Koutsoupias, and P. G. Spirakis, editors, *Proceedings of the 3rd International Symposium on Algorithmic Game Theory*, volume 6386, pages 335–346. Springer, 2010. 67

B. de Keijzer, E. Markakis, G. Schäfer, and O. Telelis. On the inefficiency of standard multi-unit auctions. *CoRR*, abs/1303.1646, 2013a. 31

B. de Keijzer, E. Markakis, G. Schäfer, and O. Telelis. On the inefficiency of standard multi-unit auctions. In *Proceedings of the 21st European Symposium on Algorithms*, Lecture Notes in Computer Science. Springer, 2013b. 31

H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack Problems*. Springer, 2004. 229

D. Kempe and M. Mahdian. A cascade model for externalities in sponsored search. In C. H. Papadimitriou and S. Zhang, editors, *Proceedings of the 4th Workshop on Internet and Network Economics*, volume 5385 of *Lecture Notes in Computer Science*, pages 585–596. Springer, 2008. 145

W. Kern and D. Paulusma. Matching games: the least core and the nucleolus. *Mathematics of Operations Research*, 28(2):294–308, 2003. 229, 231, 246, 247

A. W. J. Kolen and J. K. Lenstra. Combinatorics in operations research. In R. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics*, chapter 35, pages 1875–1910. Elsevier, 1995. 166

E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In C. Meinel and S. Tison, editors, *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science*, volume 1770 of *Lecture Notes in Computer Science*, pages 404–413. Springer, 1999. 14, 70, 71

E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009. 14

E. Koutsoupias, M. Mavronicolas, and P. G. Spirakis. Approximate equilibria and ball fusion. *Theory of Computing Systems*, 36(6):683–693, 2003. 70, 71

V. Krishna. *Auction Theory*. Academic Press, 2002. 34, 57, 65, 143

M. Langberg, Y. Rabani, and C. Swamy. Approximation algorithms for graph homomorphism problems. In J. Díaz, K. Jansen, J. D. P. Rolim, and U. Zwick, editors, *Proceedings of the 9th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, volume 4110 of *Lecture Notes in Computer Science*, pages 176–187. Springer, 2006. 199

R. Lavi and C. Swamy. Truthful and near-optimal mechanism design via linear programming. In *Proceedings of the 46th IEEE Annual Symposium on Foundations of Computer Science*, pages 595–604. IEEE Computer Society, 2005. 205

J. Ledyard. Public goods: A survey of experimental resesarch. In J. Kagel and A. Roth, editors, *Handbook of Experimental Economics*, chapter 2, pages 111–194. Princeton University Press, 1997. 86, 88, 89

B. Lehmann, D. J. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006. 31

D. K. Levine. Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, 1(3):593–622, 1998. 88, 89

L. Liang and Q. Qi. Cooperative or vindictive: Bidding strategies in sponsored search auctions. In X. Deng and F. C. Graham, editors, *Proceedings of the 3rd Workshop on Internet and Network Economics*, volume 4858 of *Lecture Notes in Computer Science*, pages 167–178. Springer, 2007. 161

D. Liben-Nowell, A. Sharp, T. Wexle, and K. Woods. Computing Shapley value in supermodular coalitional games. In J. Gudmundsson, J. Mestre, and T. Viglas, editors, *Proceedings of the 18th International Computing and Combinatorics Conference*, volume 7434 of *Lecture Notes in Computer Science*, pages 568–579. Springer, 2011. 261

B. Lucier and R. Paes-Leme. Gsp auctions with correlated types. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 71–80. ACM, 2011. 17

T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A new model for selfish routing. *Theoretical Computer Science*, 406:187–206, October 2008. 91, 106

J. Ma. Strategy-proofness and the strict core in a market with indivisibilities. *International Journal of Game Theory*, 23(1):75–83, 1994. 216

E. Maasland and S. Onderstal. Auctions with financial externalities. *Economic Theory*, 32(3):551–574, 2007. 144

F. Maffioli and G. Galbiati. Approximability of hard combinatorial optimization problems: an introduction. *Annals of Operations Research*, 96(1):221–236, 11 2000. 240

I. Mann and L. S. Shapley. Values of large games, iv: Evaluating the electoral college by montecarlo techniques. Technical Report RM-2651, RAND Corporation, 1960. 261

E. Markakis and O. Telelis. Uniform price auctions: Equilibria and efficiency. In M. Serna, editor, *Proceedings of the 5th Symposium on Algorithmic Game Theory*, volume 7615 of *Lecture Notes in Computer Science*, pages 227–238. Springer, 2012. 33, 34, 35, 36, 37, 40, 41, 46, 47, 65, 66

T. Matsui and Y. Matsui. A survey of algorithms for calculating power indices of weighted majority games. *Journal of the Operations Research Society of Japan*, 43 (1):71–86, 2000. 237, 239

C. A. Meyers and A. S. Schulz. The complexity of welfare maximization in congestion games. *Networks*, pages 252–260, 2012. 19, 166

T. Michalak, J. Sroka, T. Rahwan, M. Wooldridge, P. Mcburney, and N. Jennings. A distributed algorithm for anytime coalition structure generation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1007–1014. International Foundation for Autonomous Agents and Multiagent Systems, 2010. 228

T. P. Michalak, K. V. Aadithya, P. L. Szczepanski, B. Ravindran, and N. R. Jennings. Efficient computation of the shapley value for game-theoretic network centrality. *Journal of Artificial Intelligence Research*, 46:607–650, 2013. 247

I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, 1996. 19, 89, 115, 165

P. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004. 34, 35, 42, 46, 65

P. R. Milgrom and R. J. Weber. Distributional strategies for games with incomplete information. *Mathematics of Operations Research*, 10(4):619–632, 1985. 12

D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14: 124–143, 1996. 19

J. Morgan, K. Steiglitz, and G. Reis. The spite motive and equilibrium behavior in auctions. *Contributions to Economic Analysis and Policy*, 2(1):1102–1127, 2003. 144, 147

T. Moscibroda, S. Schmid, and R. Wattenhofer. When selfish meets evil: Byzantine players in a virus inoculation game. In *Proceedings of the 25th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, pages 35–44. ACM, 2006. 88, 145

H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003. 246

U. Nadav and T. Roughgarden. The limits of smoothness: A primal-dual framework for price of anarchy bounds. In A. Saberi, editor, *Proceedings of the 6th Workshop on Internet and Network Economics*, volume 6484 of *Lecture Notes in Computer Science*, pages 319–326. Springer, 2010. 17

J. F. Nash. Equilibrium points in n-person games. In *Proceedings of the National Academy of Sciences*, volume 36, pages 48–49. Mack Printing Company, 1950. 8

N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007. 95, 97, 165, 204

C. Noussair. Equilibria in a multi-object uniform price sealed bid auction with multi-unit demands. *Economic Theory*, 5(2):337–351, 1995. 34

A. Ockenfels, D. H. Reiley, and A. Sadrieh. Online auctions. In T. Hendershott, editor, *Economics and Information Systems*, volume 1 of *Handbooks in Information Systems*, chapter 12, pages 571–628. Elsevier, 2006. 34

U.S. Department of Treasury. Uniform-price auctions: Update of the treasury experience. Available at `http://www.treasury.gov/press-center/press-releases/Documents/upas.pdf`, 1998. 34

N. Ohta, V. Conitzer, R. Ichimura, Y. Sakurai, A. Iwasaki, and M. Yokoo. Coalition structure generation utilizing compact characteristic function representations. In *15th International Conference on the Principles and Practice of Constraint Programming*, pages 623–638, 2009. 229

M. D. Plummer. Matching theory - a sampler: from Dénes König to the present. *Discrete Mathematics*, 100:177–219, 1992. 262

T. Quint and J. Wako. On houseswapping, the strict core, segmentation, and linear programming. *Mathematics of Operations Research*, 29(4):861–877, 2004. 210

M. Rabin. Incorporating fairness into game theory and economics. *American Economic Review*, 83(5):1281–1302, 1993. 89

T. Rahwan, S. Ramchurn, N. Jennings, and A. Giovannucci. An anytime algorithm for optimal coalition structure generation. *Journal of Artificial Intelligence Research*, 34:521–567, 2009. 228

P. J. Reny. On the existence of pure and mixed strategy Nash equilibria in discontinuous games. *Econometrica*, 67:1029–1056, 1999. 34

R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973. 19, 166

J. Roskind and R. E. Tarjan. A note on finding minimum-cost edge-disjoint spanning trees. *Mathematics of Operations Research*, 10(4):701–708, 1985. 242

A. Roth. The price of malice in linear congestion games. In C. H. Papadimitriou and S. Zhang, editors, *Proceedings of the 4th Workshop on Internet and Network Economics*, volume 5385 of *Lecture Notes in Computer Science*, pages 118–125. Springer, 2008. 19, 88

A. Roth and M. A. O. Sotomayor. *Two-Sided Matching: A Study in Game Theoretic Modelling and Analysis*. Cambridge University Press, 1990. 246

T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005. 165

T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st Annual ACM Symposium on the Theory of Computing*, pages 513–522. ACM, 2009. 15, 16, 85, 86, 92, 93, 118, 124

T. Roughgarden. The price of anarchy in games of incomplete information. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 862–879. ACM, 2012. 17, 35, 46, 137

T. Roughgarden and F. Schoppmann. Local smoothness and the price of anarchy in atomic splittable congestion games. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms*, pages 255–267. ACM, 2011. 17, 19, 115

T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49 (2):236–259, 2002. 19

S. Sahni and T. Gonzalez. P-complete approximation problems. *Journal of the ACM*, 23(3):555–565, 1976. 166

T. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohmé. Coalition structure generation with worst case guarantees. *Artificial Intelligence*, 111(1–2):209–238, 1999. 228

T. Service and J. Adams. Approximate coalition structure generation. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*. AAAI Press, 2010. 228

L. S. Shapley. A value for n-person games. *Annals of Mathematical Studies*, 28:307–317, 1953. 246

L. S. Shapley and H. Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23–37, 1974. 209, 210

L. S. Shapley and M. Shubik. The assignment game i: The core. *International Journal of Game Theory*, 1:111–130, 1972. 246

T. Shrot, Y. Aumann, and S. Kraus. On agent types in coalition formation problems. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems*, pages 757–764. International Foundation for Autonomous Agents and Multiagent Systems, 2010. 234, 237

T. Solymosi and T. E. S. Raghavan. An algorithm for finding the nucleolus of assignment games. *International Journal of Game Theory*, 23:119–143, 1994. 247

T. Sönmez and M. U. Ünver. Matching, allocation, and exchange of discrete resources. In J. Benhabib, M. O. Jackson, and A. Bisin, editors, *Handbook of Social Economics*, volume 1 of *Handbooks in Economics*, chapter 17, pages 781–852. Elsevier, 2011. 209, 212

V. Syrgkanis. Bayesian games and the smoothness framework. *CoRR*, abs/1203.5155, 2012. 17, 35, 46

V. Syrgkanis and É. Tardos. Composable and efficient mechanisms. In *Proceedings of the 45th ACM Symposium on the Theory of Computing*, pages 211–220. ACM, 2013. 17, 33, 34, 35, 36, 37, 46, 49, 59, 60, 61, 62, 63, 65, 66

R. Tarjan. Depth-first search and linear graph algorithms. *SIAM Journal on Computing*, 1(2):146–160, 1972. 214

A. D. Taylor and W. S. Zwicker. *Simple Games: Desirability Relations, Trading, Pseudoweightings*. Princeton University Press, 1999. 230, 233

S. Ueda, M. Kitaki, A. Iwasaki, and M. Yokoo. Concise characteristic function representations in coalitional games based on agent types. In T. Walsh, editor, *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, pages 393–399. AAAI Press, 2011. 253

I. Vetsikas and N. Jennings. Outperforming the competition in multi-unit sealed bid auctions. In *Proceedings of the 6th International Conference on Autonomous Agents and Multiagent Systems*, pages 702–709. International Foundation for Autonomous Agents and Multiagent Systems, 2007. 144, 145

A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Proceedings of the 43rd IEEE Annual Symposium on Foundations of Computer Science*, pages 416–425. IEEE Computer Society, 2002. 86, 90

W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961. 34

E. Winter. The shapley value. In R. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 3, chapter 53, pages 2025–2054. Elsevier, 2002. 23, 246

A. C.-C. Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proceedings of the 18th IEEE Annual Symposium on Foundations of Computer Science*, pages 222–227. IEEE Computer Society, 1977. 236

D. Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3(1):103–128, 2007. 172

List of Symbols

The following list gives a short description of the symbols used throughout this thesis.

Symbol	**Description**
$\mathbb{R}$	The set of real numbers.
$\mathbb{Q}$	The set of rational numbers.
$\mathbb{N}$	The set of natural numbers (including 0).
$\mathbb{R}_{\geq a}$,$\mathbb{Q}_{\geq a}$,$\mathbb{N}_{\geq a}$	The sets of real, rational, and natural numbers greater than or equal to a.
$\mathbb{R}_{>a}$,$\mathbb{Q}_{>a}$,$\mathbb{N}_{>a}$	The sets of real, rational, and natural numbers greater than a.
e	The base of the natural logarithm.
Γ	A strategic game or cooperative game.
G	A graph.
E	The edge set of a graph.
V	The vertex set of a graph (usually identified with $[n]$), or the set of type vectors of an incomplete information game.
i	A player in a game or a vertex of a graph.
j	A facility of a congestion game, or a machine in a minsum scheduling game, or a slot in a generalized second price auction, or a house in a housing market.
n	The number of players of a game or the number of vertices of a graph.
m	The number of facilities of a congestion game, or the number of machines in a scheduling game, or the number of edges of a graph.
u	The vector of utility functions of a utility maximization game.
c	The vector of cost functions of a cost minimization game.

u_i	The utility function of player i in a utility maximization game.
c_i	The cost function of player i in a cost minimization game.
Σ	The set of strategy profiles of a strategic game, i.e., the Cartesian product $\Sigma_1 \times \cdots \times \Sigma_m$.
Σ_i	The strategies of player i in a strategic game.
PE_Γ	The pure equilibria of game Γ.
NE_Γ	The Nash equilibria of game Γ.
ClE_Γ	The correlated equilibria of game Γ.
$PBNE_\Gamma$	The pure Bayes-Nash equilibria of incomplete information game Γ.
$MBNE_\Gamma$	The mixed Bayes-Nash equilibria of incomplete information game Γ.
U	A social welfare function.
C	A social cost function.
V_i	The set of type vectors of player i in an incomplete information game.
π	The type distribution of an incomplete information game.
π_i	The type distribution of player i in an incomplete information game.
P	A set of players in a game.
$P_j(s)$	The set of players (in a congestion game) that choose facility j under strategy profile s.
d	The vector of delay functions of a congestion game.
d_j	The delay function of facility j in a congestion game.
Φ	A potential function for a game.
v	A vertex of a graph or the characteristic function of a cooperative game.
v_i	The valuation function of player i in an auction game.
σ	A probability distribution (i.e., a probability mass function or probability density function), typically on the strategy profiles of a game; or a permutation on the set of players of a game.
$\Delta(X)$	The set of probability distributions on the set X.
s	A strategy profile of a game.
s_i	The strategy chosen by player i.
s_{-i}	The vector s excluding its i'th element.
s_{-P}	The vector s excluding the indices in P.
(s'_i, s_{-i})	The vector obtained from vector s by replacing its i'th element by s'_i.

(s'_P, s_{-P})	The vector obtained from vector s by replacing the elements on indices P by the elements of the $\lvert P\rvert$-dimensional vector s'_P.
α	An altruism vector or social context.
$\check{\alpha}$	The lowest altruism level of altruism vector α.
$\hat{\alpha}$	The highest altruism level of altruism vector α.
φ	The Shapley value of a cooperative game or the golden ratio.
φ_i	The Shapley value of player i in a cooperative game.
κ_i	The raw Shapley value of player i in a cooperative game.
W	The Lambert W function.
W_{-1}	The lower branch of the Lambert W function.
$\preceq$	A relation, or a partial order, or a tie-breaking rule, or a vector of player preferences of a housing market.
$p_{i,j}$	The processing time of job i on machine j in a scheduling game.
a_j	The speed of machine j in a scheduling game with related machines.
w_i	The length of job i in a scheduling game with related machines, or the weight of player i in a weighted voting game.
β	The vector of click-through rates of a generalized second price auction.
β_j	The click-through rate of slot j in a generalized second price auction.
$r(s)$	The ranking under strategy profile s in a generalized second price auction, or the runner-up in a second price generalized procurement auction.
μ_i	The penalty multiplier of player i in a generalized procurement auction.
$w(s)$	The winner under s in a generalized procurement auction.
ϵ-PE_Γ	The pure ϵ-equilibria of game Γ.
LE	The set of limit equilibria of a generalized procurement auction.
$\mathcal{P}$	A coalition structure of a cooperative game.
q	The quota of a weighted voting game.
$[q; w_1, \ldots, w_n]$	A weighted voting game with quota q and player weights $w_1, \ldots, w_n$.
H	A housing market.
$\preceq_i$	The preferences of player i in a housing market.
ω	The initial endowment of a housing market.
$MG(G)$	The matching game associated to graph G.
$MG(G, w)$	The matching game associated to weighted graph (G, w).
(G, M)	A matched graph, where G is a graph and M is a matching of G.

Summary

This Ph.D. thesis discusses research in algorithmic game theory, carried out by the author and various collarborators. Algorithmic game theory is an area of research that lies in the intersection of economics, mathematics and computer science. The field encompasses the combination of game theory (the study of strategic decision making and models of conflict) with theoretical computer science. The central object of study in the thesis is a *game*, which is an abstraction of a situation in which multiple autonomous entities (called *players*) participate by making autonomous decisions.

The topics discussed are roughly centered around the notions of cooperation, externalities, and in general "other-regarding behavior". The study of these topics is motivated by the observation that the usual assumption that players act perfect rational, and in isolation (classically assumed in strategy game theory) is often inaccurate: in many realistic scenarios, players are embedded into some sort of social context and/or are able to cooperate with eachother.

In this thesis we investigate models that include (or allow for the inclusion of) the aspects of cooperation and externalities in games. We will be interested in the impact of this inclusion to various established notions and results in algorithmic game theory; most prominently the *price of anarchy*, which measures the worst-case quality of an *equilibrium* (i.e., roughly a "stable outcome") of a game. Another prominent theme is the study of the consequences to algorithmic problems, that the presence of externalities induces. Some related topics are studied as well: various problems in cooperative game theory in the final chapters of the thesis, and the price of anarchy of multi-unit auctions in the second chapter of the thesis.

The most important results and insights obtained in this thesis that are related are summarized as follows:

- Lower bounds and upper bounds on the price of anarchy of multi-unit auctions. The upper bounds are shown to be optimal with respect to the currently known techniques.

- An asymptotic characterization of the price of anarchy of *linear bottleneck congestion game* when cooperation among the players can occur. Linear bottleneck congestion games form a class of games that model e.g. load balancing scenarios in which multiple autonomous participants have to schedule their jobs on a set of machines.

- The price of anarchy of various games may deteriorate when players behave more altruistically, and it is possible to derive good upper bounds on the price of anarchy of some important classes of games, by extending a known technique.

- The same ideas may be used for a more general scenario in which players have particular, player-specific, attidudes towards other players. For many such *social contexts*, some broad classes of games can be shown to have a low price of anarchy.

- It is possible to characterize the set of equilibria and the price of anarchy of procurement auctions in a precise way, in the presence of players that behave spitefully.

- There exist efficient algorithms that find a good-quality profile of strategies for congestion games in which players have positive externalities for sets of other players, from which they may derive additional utility. Congestion games form a class of games of central importance, and model all sorts of situations in which there are resources of which the performance depends on the set of entities making use of it.

- For the setting of housing markets, there exists a broad and simple family of mechanisms that compute allocations of houses to players, and satisfy a large set of desirable properties. Moreover, there exists a mechanism that has been previously proposed and falls within this class, but unfortunately does not compute an allocation in an satisfying amount of time.

- Computing the best structure for a set of players to cooperate is difficult in general, but easy for a number of special cases. Most importantly, there exists an efficient algorithm for this problem in case the number of types of the players is bounded and it is known which player has which type.

- The hardness of computing the Shapley value in a matching game can be completely characterized. The Shapley value provides an important method for dividing a value among a set of cooperating players. Matching games form a class of cooperative games that help understand and model auctions and assignments.

Samenvatting

Nederlandse vertaling titel: ***Externaliteiten en Samenwerking in Algoritmische Speltheorie.***
Dit proefschrift behandelt onderzoek in algoritmische speltheorie, uitgevoerd door de auteur en verschillende andere mede-onderzoekers. Algoritmische speltheorie is een onderzoeksgebied dat in de intersectie van economie, wiskunde, en informatica ligt. Het gebied omvat de combinatie van speltheorie (de studie van conflictmodellen en het maken van strategische beslissingen) en theoretische informatica. Het object dat centraal staat in speltheorie is het *spel*, hetgeen een abstractie is van een situatie waarin meerdere autonome entiteiten (genaamd *spelers*) participeren door het maken van autonome beslissingen.

De onderwerpen die in het proefschrift besproken worden draaien om de concepten van samenwerking, externaliteiten, en algemener gezegd: "met-anderen-rekeninghoudend gedrag". De motivatie achter de studie van deze onderwerpen is de observatie dat de gebruikelijke aanname (in klassieke speltheorie) dat spelers zich perfect rationeel gedragen, en in isolatie opereren, vaak inaccuraat is: in veel realistische scenarios bevinden de spelers zich in een sociale context en zijn de spelers in staat om met elkaar samen te werken.

We onderzoeken speltheoretische modellen waarin de concepten van samenwerking en externaliteiten geïncorporeerd zijn, of een dergelijke incorporatie toestaan. We zijn vervolgens geinteresseerd in de impact die dit heeft op verscheidene gevestigde begrippen en resultaten binnen de algoritmische speltheorie. Een van de meer prominente noties die we bestuderen is de *prijs van anarchie*: een manier om de slechtst mogelijke kwaliteit van een *evenwicht* (i.e., grof weg een "stabiele uitkomst") van een spel te meten. Een ander prominent thema wordt gevormd door de gevolgen van externaliteiten op algoritmische problemen. Er worden bovendien enkele gerelateerde onderwerpen bestudeerd: verschillende problemen in coöperatieve speltheorie staan centraal in de laatste hoofdstukken van het proefschrift, en de prijs van anarchie in *multi-unit* veilingen in het tweede hoofdstuk van het proefschrift.

De belangrijkste resultaten en inzichten die verkregen worden in dit proefschrift zijn hieronder geschetst.

- Bovengrenzen en ondergrenzen op de prijs van anarchie van multi-unit veilingen. We bewijzen dat de bovengrenzen optimaal zijn met respect tot de huidig bekende technieken.

- Een asymptotische karakterisatie van de prijs van anarchie van *lineaire bottleneck congestiespellen* wanneer samenwerking tussen de spelers mogelijk is. Lineaire bottleneck congestiespellen vormen een klasse spellen die onder andere taakverdelingsscenarios modelleren, waarin meerdere autonome deelnemers een aantal taken moeten plannen op een verzameling machines.

- The prijs van anarchie van verscheidene spellen kan slechter worden wanneer spelers zich altruistischer gedragen, en het is mogelijk om scherpe bovengrenzen af te leiden op de prijs van anarchie, voor enkele belangrijke spelklasses, door middel van een extensie op een bekende en veelgebruikte techniek.

- Dezelfde ideeën kunnen gebruikt worden voor een algemener scenario waarin spelers een bepaalde speler-specifieke houding hebben naar andere spelers. Voor vele van zulke *sociale contexten* kunnen we aantonen dat enkele grote spelklasses een lage prijs van anarchie hebben.

- Het is mogelijk om de verzameling van evenwichten en de prijs van anarchie van inkoopveilingen op precieze wijze te karakteriseren, in situaties waar de spelers in de veiling een afgunstige houding tegen elkaar hebben.

- Er bestaan efficiënte algoritmes die strategieprofielen van hoge kwaliteit vinden in congestiespellen waarin spelers positieve externaliteiten hebben voor sets van andere spelers. Congestiespellen vormen een spelklasse die allerlei situaties modelleren waarbij er sprake is van een verzameling hulpbronnen waar de spelers gebruik van kunnen maken. De kwaliteit van de prestatie van deze hulpbronnen hangt af van de verzameling spelers die er gebruik van maken.

- Voor huizenmarkten bestaat er een grote en eenvoudige familie van mechanismes die allocaties genereren van huizen aan spelers, met een grote verzameling belangrijke goede eigenschappen. Bovendien bestaat er een eerder geintroduceerd mechanisme voor huizenmarkten die binnen deze klasse valt, waarvan we kunnen bewijzen dat hij helaas geen goede tijdscomplexiteit heeft.

- Het berekenen van de beste structuur waarin een groep van spelers met elkaar kan samenwerken is moeilijk in het algemeen, maar makkelijk voor een aantal speciale gevallen. Een van de belangrijkere van zulke gevallen is wanneer het aantal spelertypes begrensd is, en het bekend is welke speler van welk type is.
- De moeilijkheid van het berekenen van de Shapley-waarde in een *matchingspel* kan volledig gekarakteriseerd worden. De Shapley-waarde biedt een belangrijke methode voor het verdelen van een waarde onder een set van samenwerkende spelers. Matchingspellen vormen een klasse coöperatieve spellen die inzicht bieden in verscheidene veilingen en scenarios die betrekking hebben op het toewijzen van taken en andere objecten.

BART DE KEIJZER

"COOPERATION AND EXTERNALITIES IN ALGORITHMIC GAME THEORY"

BART DE KEIJZER - "COOPERATION AND EXTERNALITIES IN ALGORITHMIC GAME THEORY"