

Designing Mechanisms for Good Equilibria

vorgelegt von
Dipl.-Math. Philipp Friedrich Freiherr von Falkenhausen
Hamburg

Von der Fakultät II – Mathematik und Naturwissenschaften
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Naturwissenschaften
– Dr. rer. nat. –

genehmigte Dissertation

Promotionsausschuss

Vorsitzender: Prof. Dr. Jochen Blath
Berichter: Dr. Tobias Harks (Assistant Professor)
Prof. Dr. Rolf H. Möhring
Prof. Dr. Britta Peis

Tag der wissenschaftlichen Aussprache: 27. Mai 2014

Berlin 2014
D 83

Preface

Over the three years that went into writing this thesis, I have received outstanding support. I would like to thank everyone who has had a part in this.

Tobias Harks introduced me to algorithmic game theory as my diploma thesis advisor. Our fruitful collaboration motivated me for further research and I am thoroughly glad to still have him as my advisor. I am grateful for his support and guidance and for the chance to work with him. Visits in Maastricht and Berlin brought days of intense mathematical discussions and convivial evenings that I have always looked forward to.

The Group for Combinatorial Optimization and Graph Algorithms has been a great home for doing research. I am thankful for Rolf H. Möhring's generous leadership and the group's open nature. A special mention goes to Max Klimm and Kai-Simon Goetzmann, with whom I shared an office and many jokes over the last years. Max has also been a great sparring partner for research questions and helped proofreading parts of this thesis. My thanks also go to Britta Peis for taking the third assessment of this thesis.

Exposure to the international research community has been a tremendous experience for me. I would like to thank Richard Cole and Tim Roughgarden for inviting me to do research with them.

Finally, I thank my parents for their continuous encouragement and undoubting support and the freedom they gave me throughout the years.

Berlin, March 2014

Philipp von Falkenhausen

This research was supported by the Deutsche Forschungsgemeinschaft within the research training group 'Methods for Discrete Structures' (GRK 1408).

Contents

1	Introduction	5
1.1	Extending the Price of Anarchy	6
1.2	Cost Sharing for Capacitated Facility Location Games	7
1.3	Quantitative Comparative Statics for a Multimarket Paradox	8
1.4	Conversion of Regret-Minimizing Algorithms in Atomic Splittable Routing	8
2	Preliminaries – Equilibria and Their Efficiency	11
2.1	Nash Equilibrium in Strategic Games	11
2.2	Frameworks for the Analysis of Equilibria	16
3	Cost Sharing for Capacitated Facility Location Games	19
3.1	Contribution and Outline	20
3.2	Model and Problem Statement	22
3.3	Characterizing Pure Nash Equilibria	24
3.4	An Optimal Protocol for the Price of Stability	27
3.5	An Optimal Protocol for the Price of Anarchy	31
3.6	Concave and Convex Costs	42
3.7	Computational Complexity of Cost Shares	44
3.8	Non-Matroid Strategy Spaces	51
3.9	Conclusion	53
4	Quantitative Comparative Statics for a Multimarket Paradox	55
4.1	Comparative Statics	57
4.2	Contribution and Outline	58
4.3	Multimarket Cournot Oligopoly Model	59
4.4	Maximum Profit Loss of an Individual Firm	61
4.5	Effect of Price Shocks on Aggregates	69
4.6	Conclusion	74
5	Convergence of Regret-Minimizing Algorithms in Atomic Splittable Routing	75
5.1	Contribution and Outline	77
5.2	Atomic Splittable Routing and No-regret Play	78
5.3	Affine Latency: Learning Converges	79

5.4 General Latency	86
5.5 Strict ρ -approximate Nash Equilibrium	92
5.6 Conclusion	95
Bibliography	95

Chapter 1

Introduction

“Game Theory is a bag of analytical tools designed to help us understand the phenomena that we observe when decision-makers interact.” Osborne and Rubinstein, 1994

Game theoretic thinking dates back at least as far as Cournot (1838), who found that a firm with a single competitor will make its actions dependent on the competitor’s actions and vice versa, until the actions are in equilibrium. The foundation of game theory as its own field of research is generally attributed to von Neumann (1928) and von Neumann and Morgenstern (1944), who formalized the observation that human interaction very closely resembles that of players in parlor games, and that the study of games therefore yields insights for the social sciences. Their definition of a strategic game consists of a set of players, a set of possible actions for each player, and a private objective function for each player that depends on the action of the player herself as well as the actions of the other players. *Dynamic* analysis of such games has since studied the behavior of the players as they react to one another, while *static* analysis assumes the existence of *equilibrium* states where the players cease to change their actions.

In the two decades that have passed since the sentence at the top of this page was written, a new direction of research has been fueled by the application of tools from the design and analysis of algorithms to game theory. Such research has been termed algorithmic game theory and has brought a range of new insights, like an understanding of the complexity of finding equilibria.

A recent focal point of algorithmic game theory has been the quantification of equilibrium efficiency. Early findings in economics, such as the example of suboptimal selfish routing by Pigou (1920), had already shown inefficient equilibria in real life situations. Especially the Braess Paradox (Braess 1968) from civil engineering, where an increase in network capacity leads to a decreased equilibrium flow, highlights how important it is to consider the efficiency of the resulting equilibria when one has the power to design a system, for example as a policy maker. Algorithmic game theory research has revisited such problems, developing methodology for quantifying and improving equilibrium efficiency.

1.1 Extending the Price of Anarchy

Analogous to the ‘approximation factor’ of an approximation algorithm and the ‘competitive ratio’ of an online algorithm, the *price of anarchy* was introduced by Koutsoupias and Papadimitriou (1999) as the worst ratio of a pure or mixed Nash equilibrium to a socially optimal allocation in a given strategic game. By now, price of anarchy research has successfully established wide knowledge about the efficiency of equilibria in standard problems including routing (e.g. Roughgarden and Tardos 2002, Roughgarden 2002, Correa et al. 2004), load balancing (e.g. Czumaj et al. 2010, Czumaj and Vöcking 2007), facility location (e.g. Vetta 2002), and recently also auctions (e.g. Paes Leme et al. 2012). We refer to the book by Nisan et al. (2007, Chapters 17-21) for an overview.

While significant effort went into finding price of anarchy bounds for standard settings, the price of anarchy has in many places also served as an inspiration to develop similar worst-case measures for the efficiency of a game’s outcome. By now, there is a large body of research that uses variations of the original approach, and we give an overview over the most important approaches. An obvious variation to the price of anarchy definition is to regard the worst correlated or coarse correlated equilibrium instead of the worst pure or mixed equilibrium. Results of this kind often use the ‘smoothness’ framework introduced by Roughgarden (2009), which guarantees the same efficiency bound for all of the aforementioned equilibrium classes. The benefit of using correlated and coarse correlated equilibria is that these require fewer assumptions and thus are less prone to certain criticisms of pure and mixed equilibria that we discuss in Section 2.1. Also, there are settings where pure and mixed equilibria have no meaningful interpretation while correlated and coarse correlated equilibria do.

In addition to the efficiency of the static equilibria, game dynamics have been studied from a price of anarchy perspective, for example using the *price of total anarchy* approach introduced by Blum et al. (2008) for games where the players use no-regret learning strategies. Another form of dynamics are best-response dynamics that have mostly been studied in games with potential functions (e.g. Mirrokni and Vetta 2004).

A further measure for the efficiency of equilibria considered in this work is the *price of stability* (Correa et al. 2004, Anshelevich et al. 2008). It uses the best equilibrium instead of worst equilibrium like the price of anarchy. This more optimistic approach is justified in the presence of a central coordinator that is able to propose the most efficient equilibrium to the players, who then have no reason to choose a deviating strategy. Combining the price of stability with the price of anarchy also gives insight into the interval that contains all equilibria of a game.

All of the above approaches assume a given type of game and set out to determine the efficiency of its equilibria. However, the price of anarchy is also used as a design metric when, say, a policy maker or the operator of a marketplace designs the game to be played. In this spirit, Christodoulou et al. (2009) designed machine scheduling policies with the goal of a low price of anarchy and Chen et al. (2010) considered the worst-case price of anarchy their objective function when designing cost sharing protocols for network games.

This thesis considers the efficiency of equilibria in three settings: first from a design perspective aiming for optimal equilibria, then from the point of view of a policy maker looking to understand the maximum effect of a parameter change on the game’s equilibrium, and finally in

a routing game where the efficiency of distributed learning dynamics is studied.

1.2 Cost Sharing for Capacitated Facility Location Games

Historically, congestion games are the natural habitat of the price of anarchy – the measure was introduced in the context of atomic routing on parallel links, a well-studied subclass of congestion games. Since then, price of anarchy bounds for many other congestion game variants have been found.

The most common mindset of price of anarchy results in the area of congestion games is the interpretation of cost as a thing that the players inherently experience and the game designer has no control over, as is the case for example for actual congestion delays in road traffic. Different lines of research (outside of congestion games) rather assume a design perspective and ask the question how the interaction of players can be formed in order to achieve a low price of anarchy. An example is the research on coordination mechanisms by Christodoulou et al. (2009) and follow-up papers (Caragiannis 2013, Cole et al. 2011, Immorlica et al. 2009), where players choose machines for a job and the design goal is a mechanism for processing the jobs such that the resulting game has a low price of anarchy.

We bring such a design perspective to congestion games, assuming the resource cost to be monetary so that it can be assigned to the players in arbitrary ways. As the players' behavior heavily depends on how they incur cost, the way that cost is assigned to the players determines a game's equilibrium outcomes. In Chapter 3, we investigate efficient ways of allocating monetary cost of a resource to its users, formalized as cost sharing protocols.

Contributions of Chapter 3

We study the design of cost sharing protocols for weighted congestion games with multi-resource strategies. Each resource has monetary cost – given as a non-decreasing but otherwise arbitrary function of the weights of its users – and additionally for each player there is a player-specific delay when connecting to a given resource. We allow any matroid base set as strategy set of a player, including for example spanning trees in a graph. Our design goal is to devise cost sharing protocols so as to minimize the resulting price of anarchy and price of stability. We investigate *basic* protocols that guarantee the existence of at least one pure Nash equilibrium, and *separable* protocols that additionally require that the resulting cost shares only depend on the set of players on a resource. We find basic and separable protocols which guarantee that the price of stability grows at most logarithmically and price of anarchy at most linearly in the number of players. For the special case of symmetric games without delays we give a basic protocol that even has a worst-case price of anarchy logarithmic in the number of players. All protocols are complemented by matching lower bound instances. The foundation of our results is a structural insight that we use to characterize the strategy profiles that are suitable to be a Nash equilibrium under any cost sharing protocol.

The above protocols are determined by an algorithm that, for a given instance, provides the player's cost shares. However, the algorithm does not run in polynomial time and in fact, we find that is not only NP-complete to find cost shares optimal with respect to the price of anarchy or

price of stability, but even $c \log(n)$ approximations are strongly NP-hard for any $0 < c < 1$. For a more restricted setting, we give a polynomial time approximation algorithm and show that the cost shares computed by this algorithm have price of anarchy of at most n and price of stability of at most \mathcal{H}_n , matching the worst-case bounds.

1.3 Quantitative Comparative Statics for a Multimarket Paradox

Comparative statics is a line of research in economics that studies how a game's equilibria change in reaction to the change of exogenous parameters. Comparative statics is focused on qualitative results, usually the monotonicity of the change of an objective function for a given marginal parameter change, and in many cases no quantitative results are known. Our analysis in Chapter 4 is inspired by a well known paradox in multimarket Cournot competition, where a positive price shock on a monopoly market may actually reduce the monopolist's profit, as pointed out by [Bulow et al. \(1985\)](#). Observing the price shock as a parameter change, the authors employ comparative statics to attain a qualitative result, leaving open the extent to which the monopolist's profit may suffer.

In some cases, known price of anarchy results, for example those for Cournot competition by [Johari and Tsitsiklis \(2005\)](#), can imply a quantitative bound for the effect of such parameter changes on equilibrium welfare. However, the price of anarchy – a comparison of an equilibrium with a non-equilibrium – is not designed for comparing two equilibria and the bounds need not be tight.

Contributions of Chapter 4

In Chapter 4, we propose a quantitative approach to comparative statics that allows to bound the maximum effect of an exogenous parameter change on a system's equilibrium. We use our approach to quantify for the first time the worst case profit reduction for multimarket Cournot competition exposed to arbitrary positive price shocks. For markets with affine price functions and firms with convex cost technologies, we show that the relative profit loss of any firm is at most 25%, no matter how many firms compete in the oligopoly. We further investigate the impact of positive price shocks on total profit of all firms as well as on social surplus. We again find tight bounds for these measures, showing that total profit and social surplus decrease by at most 25% and 16.7%, respectively.

1.4 Conversion of Regret-Minimizing Algorithms in Atomic Splittable Routing

There are many arguments why pure and even mixed Nash equilibria may be a too restrictive subset of the strategy space to capture actual behavior. Key among these criticisms is the question how players arrive at an equilibrium, especially given the computational intractability of finding equilibria and the games' distributed nature. This has driven efforts to extend known price of anarchy bounds to broader equilibrium concepts, in particular correlated and coarse correlated equilibria. These appear as the convergence limit of distributed learning processes, in the latter case of so called no-regret learning.

While price of anarchy bounds for pure and mixed Nash equilibria do not in general extend to coarse correlated equilibria, they have been shown to do so in certain classes of games with additional structure. More recently, a technique known as smoothness (Roughgarden 2009) has allowed such results for several types of congestion games. Using an extension of the smoothness framework, Roughgarden and Schoppmann (2011) were able to show that price of anarchy bounds for atomic splittable routing games with polynomial cost functions coincide for pure, mixed and correlated Nash equilibria. The question whether these bounds also apply to coarse correlated equilibria is the topic of Chapter 5.

Contributions of Chapter 5

For atomic splittable routing games with linear cost functions, we prove that no-regret learners converge to Nash equilibria in the following sense: at almost all time steps, the cost of the flow is close to that of the best responses of players given the flow of other players. The rate of convergence depends polynomially on the players' regret bounds, the network size, and the maximum slope of a latency function. A consequence of our analysis is that coarse correlated equilibria are essentially unique in atomic splittable routing games with affine cost functions.

Moreover, we prove that networks with non-linear cost functions exhibit qualitatively different behavior – for example, even in symmetric games there are coarse correlated equilibria with expected cost worse than that of every Nash equilibrium. In fact, we show for games with cubic cost that coarse correlated equilibria are more expensive than pure Nash equilibria in worst-case instances, implying that price of anarchy bounds for pure equilibria do not hold for coarse correlated equilibria. We provide analogs of our convergence results in such networks that are close to the best possible.

Chapter 2

Preliminaries – Equilibria and Their Efficiency

This chapter introduces the major concepts used throughout the thesis. It aims to do so in a way that gives a comprehensive overview over existing methodologies. We start out with the definition of a strategic game, which is at all times our model of interaction, then introduce relevant equilibrium concepts and finally cover the price of anarchy and price of stability as measures for the efficiency of equilibria.

2.1 Nash Equilibrium in Strategic Games

The strategic game is our basic model for the interaction of a group of selfish players. We denote the set of players by $N = \{1, \dots, n\}$ and for each player $i \in N$ the set of available strategies by X_i . We call a collection of strategies $x_i \in X_i$, one for each player $i \in N$, a strategy profile $x = (x_i)_{i \in N}$ and denote by $X = X_1 \times \dots \times X_n$ the joint strategy space. Each player has a cost function $c_i : X \rightarrow \mathbb{R}$ that returns her cost in a given strategy profile, that is, her cost is dependent on both her own strategy as well as the strategy of all other players.

Definition 2.1 (strategic game)

A strategic game is a tuple $G = (N, X, c)$, where N is a finite and nonempty set of players, $X = \times_{i \in N} X_i$ is a strategy space, and $c : X \rightarrow \mathbb{R}^n$ is a joint cost function, such that the cost of player $i \in N$ in a strategy profile $x \in X$ is $c_i(x)$.

We assume that each player strives to minimize her cost, a situation known as a *minimization game*. Equivalently, in a maximization game, each player has a utility function that she maximizes.

Example 2.2 (Congestion Games). Rosenthal (1973) introduced *congestion games* as ‘a class of games possessing pure-strategy Nash equilibria’. In a congestion game, there is a set of players N and a ground set of resources R . Each player’s strategy set is a subset $X_i \subseteq 2^R$, for example a collection of paths of some underlying network. Each resource is associated with a congestion function that maps a given number of users to a non-negative congestion cost, for example the travel time on a road given the number of cars using that road. Then, player i ’s cost in a strategy profile x is the sum of the congestion costs of the resources in x_i .

The most common solution concepts in game theory are the pure Nash equilibrium and its variants. A pure Nash equilibrium is a strategy profile in which no player can unilaterally decrease her cost, i.e., switching one's own strategy is not profitable unless other players switch their strategy as well. For a formal definition, we denote for a game (N, X, c) by (z_i, x_{-i}) the strategy profile where player i plays some strategy $z_i \in X_i$ and all other players play according to a joint strategy profile $x \in X$.

Definition 2.3 (pure Nash equilibrium)

Given a strategic game (N, X, c) , a strategy profile $x \in X$ is a pure Nash equilibrium if for each player $i \in N$,

$$c_i(x) \leq c_i(y_i, x_{-i}),$$

for all alternative strategies $y_i \in X_i$.

An interpretation of this definition is to say that x_i is player i 's best response to the actions x_{-i} of the other players. We formalize this to define player i 's best response function

$$BR_i(z_{-i}) := \{z_i \in X_i : c_i(z_i, z_{-i}) \leq c_i(z'_i, z_{-i}) \text{ for all } z'_i \in X_i\}.$$

Then, a profile x is a pure Nash equilibrium if for every player $x_i \in BR_i(x_{-i})$.

The pure Nash equilibrium is arguably the most straightforward solution concept for analyzing the interaction of individuals and in particular captures the circular relation between the players' actions that inspired von Neumann (1928). While it leads to powerful results, it relies on strong assumptions and there are concerns regarding its applicability. The most notable drawback is that many games do not possess a pure Nash equilibrium, as demonstrated in the ‘Rock Paper Scissors’ game below.

Example 2.4 (Rock Paper Scissors). There are two players, $N = \{1, 2\}$, that both simultaneously draw a shape from the set {‘rock’, ‘paper’, ‘scissor’}. ‘rock’ beats ‘scissor’, which in turn beats ‘paper’ while ‘paper’ beats ‘rock’, meaning that for example player 1 wins if she chooses ‘paper’ while player 2 chooses ‘rock’. The looser has to pay the winner \$1 and no payment is made in a draw. Clearly, this game has no pure Nash equilibrium, as any strategy profile offers one of the players a profitable deviation.

Numerous variants of the pure Nash equilibrium concept have been proposed to address situations where different assumptions are necessary. We introduce some of the most common generalizations.

The most notable variant of the pure Nash equilibrium is the *mixed* Nash equilibrium. It introduces randomization, that is, the players specify a probability distribution called *mixed strategy* that they draw their strategy from instead of fixing a single ‘pure’ strategy. In other words, a mixed strategy of a player $i \in N$ is a probability distribution $\sigma_i : X_i \rightarrow [0, 1]$.

Definition 2.5 (mixed Nash equilibrium)

Given a strategic game (N, X, c) and a set of independent probability distributions $\sigma_i : X_i \rightarrow [0, 1]$, one for each player $i \in N$, the product distribution $\sigma := \sigma_1 \times \dots \times \sigma_k$ is a mixed Nash equilibrium if no player can improve her expected cost, that is,

$$\mathbb{E}_{x \sim \sigma}[c_i(x)] \leq \mathbb{E}_{x_{-i} \sim \sigma_{-i}}[c_i(z_i, x_{-i})]$$

for all pure strategies $z_i \in X_i$.

Note that by linearity of the expected value, comparing the expected cost in x to the expected cost in an alternative pure strategy $z_i \in X_i$ is equivalent to comparing to all mixed strategies. Clearly, a pure Nash equilibrium can be seen as a mixed Nash equilibrium where each distribution σ_i has a single atomic weight.

Mixed strategies allow an equilibrium in the above ‘Rock Paper Scissors’ game when both players play each shape with equal probability. In fact, [Nash \(1950\)](#) showed that every finite game¹ has an equilibrium in mixed strategies. However, whether the mixed Nash equilibrium is applicable for a given type of game or even in general has been subject to debate, see [Osborne and Rubinstein \(1994, § 3.2\)](#) for an overview of the discussion. Most notably, there is no clear interpretation for a mixed strategy in single-shot games. Also, even in repeated games it is unclear how the players arrive at a mixed Nash equilibrium: [Chen et al. \(2009\)](#) and [Daskalakis et al. \(2009\)](#) have shown that finding a mixed Nash equilibrium is intractable (formally, PPAD-complete) even in two-player games, and, given that it is hard to computationally find the equilibrium, there is significant doubt that one can assume uncoordinated players to find equilibrium strategies.

Another drawback of mixed equilibria is that it is necessary to assume the players to be *risk-neutral* to use the expected value to express the players’ objective. However, experiments have rebutted this assumption ([Allais 1953](#)) and without this assumption, mixed strategy Nash equilibria need not exist ([Fiat and Papadimitriou 2010](#)).

2.1.1 Repeated Games and Learning Equilibria

Studying a given type of equilibrium is much more plausible if it comes with an explanation of how players end up playing their equilibrium strategy. The coarse correlated equilibria studied in Chapter 5 are the convergence limit of a distributed learning process, and so are correlated equilibria, which we also introduce for completeness of this overview of equilibrium concepts. Correlated equilibria are a generalization of mixed equilibria to joint (non-product) distributions over strategy profiles.

Definition 2.6 (correlated equilibrium)

Given a strategic game (N, X, c) , a probability distribution $\sigma : X \rightarrow [0, 1]$ is a correlated equilibrium if for every player $i \in N$,

$$\mathbb{E}_{x \sim \sigma}[c_i(x)|x_i] \leq \mathbb{E}_{x \sim \sigma}[c_i(f(x_i), x_{-i})|x_i] \quad (2.1)$$

for all functions $f : X_i \rightarrow X_i$.

A common interpretation is that of a mediator drawing a joint strategy profile x from a known distribution σ , and informing each player privately about her draw x_i . The distribution σ is a correlated equilibrium if following the mediator’s recommendation has no more expected cost than deviating to a different strategy based on the recommendation. Mixed Nash equilibria are

¹ A strategic game (N, X, c) is a finite game if each player has a finite number of strategies, i.e., when X is finite.

precisely the correlated equilibria that are product distributions. Correlated equilibria have not only been studied in context of the interpretation with the mediator, but also because they are relatively tractable: a correlated equilibrium can be described by a small set of linear equations and can thus be computed in polynomial time (Gilboa and Zemel 1989). Additionally, such equilibria can be ‘learned’ in the sense described below through learning algorithms with vanishing internal regret (cf. Blum and Mansour 2007).

Chapter 5 deals with the efficiency of coarse correlated equilibria. Such equilibria occur if all players employ algorithms with vanishing external regret to choose their strategy in a repeated game.

Definition 2.7 (external regret and no-regret algorithm)

Let X_i be a convex set of strategies and c_i^1, c_i^2, \dots an infinite sequence of cost functions, where each $c_i^t : X_i \rightarrow \mathbb{R}$. Given an online algorithm that selects, at each time step t given the cost functions c_i^1, \dots, c_i^{t-1} , a solution $x_i^t \in X_i$, the *external regret* $\mathcal{R}(T)$ of the algorithm at a timestep T is the cost of the choices x_i^t , $t = 1, \dots, T$, compared to the best static solution in hindsight, that is,

$$\mathcal{R}(T) := \sum_{t=1}^T c_i^t(x_i^t) - \min_{y_i \in X_i} \sum_{t=1}^T c_i^t(y_i) \quad (2.2)$$

An algorithm is called *no-regret* algorithm if its average regret $\frac{1}{T}\mathcal{R}(T)$ approaches 0 as the time limit T grows, i.e., $\frac{1}{T}\mathcal{R}(T) = o(1)$.

Numerous no-regret algorithms are known, for example Zinkevich (2003) gives a simple gradient based algorithm with external regret at most $\mathcal{R}(T) = O(\sqrt{T})$ for closed, convex, bounded and non-empty strategy spaces X_i and differentiable cost functions c_i^t with bounded first derivative. For settings where c_i^t are strictly concave and differentiable, Hazan et al. (2007) give an algorithm with regret at most $\mathcal{R}(T) = O(\log(T))$.

Given a repeated strategic game (N, X, c) , that is, a setting where the players repeatedly choose a strategy $x_i^t \in X_i$ based on the history of previous strategy profiles x^1, \dots, x^{t-1} , a player i may employ a no-regret algorithm, using the cost functions $c_i^t(x_i^t) = c_i(x_i^t, x_{-i}^t)$. Note that the algorithm’s performance guarantee holds independent of the other players’ behavior, and consequently a player may employ a no-regret algorithm even if it is unclear whether the other player do so as well. When all players simultaneously choose their strategies with no-regret procedures, we speak of a no-regret sequence.

Definition 2.8 (no-regret sequence)

Given a strategic game (N, X, C) that is played repeatedly, we call a sequence of strategy profiles $(x^T)_{T \in \mathbb{N}}$ a no-regret sequence if for each player i there is a regret bound $\mathcal{R}_i(T)$ with

$$\sum_{t=1}^T c_i(x^t) - \min_{y_i \in X_i} \sum_{t=1}^T c_i(y_i, x_{-i}^t) \leq \mathcal{R}_i(T)$$

with $\mathcal{R}_i(T) = o(T)$.

The empirical distribution of such no-regret sequence converges, as $T \rightarrow \infty$, to what is known as a coarse correlated Nash equilibrium.

Definition 2.9 (coarse correlated equilibrium)

Give a strategic game (N, X, c) , a (possibly non-product) probability distribution $\sigma : X \rightarrow [0, 1]$ is a coarse correlated equilibrium if for every player $i \in N$,

$$\mathbb{E}_{x \sim \sigma}[c_i(x)] \leq \mathbb{E}_{x \sim \sigma}[c_i(z_i, x_{-i})] \quad (2.3)$$

for all $z_i \in X_i$.

Not only do no-regret sequences converge to coarse correlated equilibria, but conversely, sampling from a coarse correlated equilibrium gives a no-regret sequence for all players. In this sense, the two concepts are dual. Note that coarse correlated equilibria are a further generalization of correlated equilibria as (2.1) implies (2.3) when the player disregards the recommendation of the mediator when considering a deviation.

2.1.2 Approximate Equilibrium

The pure Nash equilibrium and the variants defined above guarantee that no player can profitably deviate. Depending on the application at hand, this may be too restrictive. There can be strategy profiles where players can profitably deviate, but with such negligible benefit that e.g. this benefit is outweighed by a cost inherent to changing the strategy. Approximate equilibria encompass precisely such strategy profiles.

Definition 2.10 (ρ -approximate pure Nash equilibrium)

Given a game (N, X, c) and an approximation factor $\rho \geq 0$, a strategy profile is a ρ -approximate pure Nash equilibrium if the combined improvement from unilateral deviations is at most ρ , i.e.,

$$\sum_{i \in N} c_i(x) \leq \rho + \sum_{i \in N} \min_{y_i \in X_i} c_i(y_i, x_{-i}).$$

Similar definitions are possible for ρ -approximate mixed, correlated and coarse correlated Nash equilibria.

Our definition of ρ -approximate Nash equilibrium differs somewhat from the definition commonly found in the literature, see for example Awerbuch et al. (2008) and Caragiannis et al. (2011). We use the term ‘strict ρ -approximate Nash equilibrium’ for strategy profiles that fit the more common definition.

Definition 2.11 (strict ρ -approximate pure Nash equilibrium)

Given a game (N, X, c) and an approximation factor $\rho \geq 0$, strategy profile is a strict ρ -approximate pure Nash equilibrium if for each player $i \in N$ the improvement from a unilateral deviation is at most ρ , i.e.,

$$c_i(x) \leq \rho + \min_{y_i \in X_i} c_i(y_i, x_{-i}).$$

The key difference here is that a strict ρ -approximate Nash equilibrium limits the individual improvement of a single player to ρ , while a ρ -approximate Nash equilibrium does not impose limits for individual players, only the aggregate improvement of all players combined is limited to ρ . Clearly, any strict ρ -approximate Nash equilibrium of a game with n players is also an

$(n\rho)$ -approximate Nash equilibrium, but not vice versa. Hence, with the appropriate scaling, any known result for strict ρ -approximate equilibria can be used for ρ -approximate equilibria. In Chapter 5, we use the wider definition because it allows us to show a stronger approximation factor ρ for the class of strategy profiles that we are interested in.

2.2 Frameworks for the Analysis of Equilibria

To analyze equilibria, one first needs a notion of *social cost*, i.e., some ordering or valuation function over a game’s strategy profiles. Most commonly used is utilitarian welfare, which in minimization games casts social cost as the sum of the costs of all players, $C(x) = \sum_{i \in N} c_i(x)$. For maximization games, as studied in Chapter 4, the equivalent is the sum of the players’ utilities.

We leverage frameworks developed in the past decades for the use in algorithmic game theory to measure the efficiency of equilibria. We consider the already mentioned *price of anarchy* and the *price of stability* as the two prevailing performance metrics used in the literature. Both compare the social cost of an equilibrium to a minimal cost strategy profile. The price of anarchy uses the most expensive equilibrium of a game, while the price of stability uses the cheapest equilibrium. To this end, we denote for a game G the set of pure Nash equilibria by $PNE(G)$.

Definition 2.12 (price of anarchy and price of stability)

The price of anarchy of a game G is the ratio of the social cost of the most expensive pure Nash equilibrium to the social cost of an optimal profile y , i.e.,

$$POA(G) = \sup_{x \in PNE(G)} \frac{C(x)}{C(y)}.$$

The price of anarchy of a class of games \mathcal{G} is the worst-case price of anarchy across all games in \mathcal{G} , i.e.,

$$POA(\mathcal{G}) = \sup_{G \in \mathcal{G}} POA(G).$$

The price of stability is a similar measure for the least expensive equilibrium of a game, i.e.,

$$POS(G) = \inf_{x \in PNE(G)} \frac{C(x)}{C(y)} \quad \text{and} \quad POS(\mathcal{G}) = \sup_{G \in \mathcal{G}} POS(G).$$

Comparing the cost of an equilibrium to an optimal strategy profile exposes the efficiency loss caused by the selfish or uncoordinated behavior of the players. Doing so with a worst-case perspective across a class of games is inspired by the worst-case analysis of the performance of approximation algorithms using an *approximation factor* (cf. Williamson and Shmoys 2011), where a solution obtained in polynomial time by an approximation algorithm is compared to an optimal solution. Another measure of similar style is the *competitive ratio* (cf. Borodin and El-Yaniv 1998) from research in online algorithms that compares the performance of an online algorithm with that of an offline algorithm.

Above, we have defined the price of anarchy and price of stability with regard to the pure Nash equilibrium, as this is the predominant solution concept in Chapters 3 and 4. In Chapter 5

we study the price of anarchy of coarse correlated equilibria, which is similar to the above definition when $PNE(G)$ is replaced with set of coarse correlated equilibria $CCE(G)$. We also use the price of total anarchy, which measures the cost of the most expensive no-regret sequence compared to an optimal profile and converges to the price of anarchy of coarse correlated equilibria as the regret goes to zero.

Chapter 3

Cost Sharing for Capacitated Facility Location Games

The existence of pure Nash equilibria and their price of anarchy and price of stability in congestion games has been studied extensively (Ackermann et al. 2009, Even-Dar et al. 2007, Fotakis et al. 2009, Gairing et al. 2011, Harks and Klimm 2012, Ieong et al. 2005, Milchtaich 1996). However, with applications like road traffic and telecommunication networks in mind, the classic congestion game treats costs as unavoidable physical delays. More recently, research has been undertaken into a model with similar structure, but where cost is considered monetary. This allows for a different range of applications and at the same time calls for new approaches.

Consider a setting where players jointly use resources that have load-dependent monetary costs and player-specific delays. The monetary cost of each resource must be shared by its users while the player-specific delays are unavoidable physical quantities. This setting arises in facility location models, where users share the monetary cost of opened facilities and additionally experience delays measured by the distance to the closest open facility. As players might value the delay and monetary cost differently, these delays depend on both the player and the facility the player is assigned to. Another example appears in distributed network design, where players jointly install capacities on a subgraph satisfying user-specific connectivity requirements. Besides the monetary cost for installing enough capacity, each player experiences player-specific delays on the resources used.

In distributed systems, players selfishly select resources for their demands based on the cost shares they have to pay and delays they experience. Hence, the way the monetary cost of a resource is shared among its users determines the equilibrium states of the strategic game induced. In this chapter, we investigate efficient ways of allocating monetary cost to a resource's users, formalized as cost sharing protocols. Our analysis focuses on providing cost sharing protocols that guarantee an optimal price of anarchy or price of stability.

Chen et al. (2010) have provided an axiomatization of cost sharing protocols in the context of network design games. We use their axioms as listed below to establish classes of cost sharing protocols.

- (1) *Budget-balance*: The cost of each resource is exactly covered by the cost shares collected from the players using the resource.

- (2) *Stability*: There is at least one pure strategy Nash equilibrium in each game induced by the cost sharing protocol.
- (3) *Separability*: When assigning the cost shares on a given resource, the protocol has no information about the load on other resources.

We call a cost sharing protocol *basic* if it satisfies (1)-(2) and *separable* if it satisfies (1)-(3). While condition (1) is straightforward, the stability condition (2) requires the existence of at least one pure Nash equilibrium, see [Osborne and Rubinstein \(1994, § 3.2\)](#) for several drawbacks of using mixed Nash equilibria instead. Condition (3) is for instance crucial for practical applications in which cost sharing protocols must be distributed because each resource has only local information about its own usage.

The only previously known result for cost sharing protocols related to congestion games is by [Kollia and Roughgarden \(2011\)](#), who considered weighted congestion games and proposed a cost sharing protocol based on the Shapley value for which they are able to prove existence of pure Nash equilibria and corresponding bounds on the price of anarchy and price of stability. They focus on polynomial cost per unit functions with non-negative coefficients, but it is not known if this protocol is optimal. There is further work on cost sharing in congestion games that does not assume a protocol design perspective but instead studies the performance of a given procedure. See [Albers \(2009\)](#), [Anshelevich et al. \(2008\)](#), [Epstein et al. \(2009\)](#), [Hoefer \(2011\)](#), [Rozenfeld and Tennenholz \(2006\)](#) regarding average cost sharing assuming non-increasing marginal cost functions modeling economies of scale or buy-at-bulk. There is also a large body of papers studying cost sharing protocols for continuous and convex strategy spaces assuming convex cost sharing functions, cf. [Chen and Zhang \(2012\)](#), [Harks and Miller \(2011\)](#), [Johari and Tsitsiklis \(2006, 2009\)](#), [Moulin \(2008, 2010\)](#).

Cost sharing approaches to facility location problems and network design problems were analyzed in [Könemann et al. \(2008\)](#) and [Pál and Tardos \(2003\)](#). In these works it is only required that total cost shares cover (approximately) total cost as the players pay for the service of being connected. In contrast to their work, we require the stricter notion that the cost of every individual resource is paid for by the players using it. Our model also includes the case of the congested facility location problem considered by [Desrochers et al. \(1995\)](#), where players jointly pay for the congestion related costs and the opening costs of the facilities. Instead of a game-theoretic problem formulation, they consider a centralized approach using mixed-integer programming techniques for minimizing total cost.

As a final pointer to further references, note that cost sharing is a central topic in the area of cooperative game theory, cf. [Archer et al. \(2004\)](#), [Bogomolnaia et al. \(2010\)](#), [Moulin and Shenker \(2001\)](#) or the survey of [Moulin \(2002\)](#).

3.1 Contribution and Outline

We study the design of cost sharing protocols for games in which every player wants to use resources that together form a basis of a player-specific matroid defined on the set of resources. We demonstrate that the aforementioned class of facility location games can be represented by

Table 3.1: Summary of the results. Unless indicated, these are tight worst-case bounds for both basic and separable cost sharing protocols.

	Player-spec. Matroids		Symmetric Matroids without Delays	
	price of stability	price of anarchy	price of stability	price of anarchy
general cost	\mathcal{H}_n	n	\mathcal{H}_n	\mathcal{H}_n (basic), n (separable)
concave cost	1	n	1	$\leq n$
convex cost	1	1	1	1

player-specific matroids. As we assume general non-decreasing cost functions on the resources, our model also includes the case of facility location with hard capacities. Our setting has further applications in network design games, where each player wants to allocate a spanning tree in a graph.

Price of Anarchy and Price of Stability. The main objective of this chapter is to design basic and separable cost sharing protocols so as to minimize the efficiency loss of the equilibria induced.

The first main result presented in Section 3.3 is a characterization of Nash equilibria in games with general cost functions and strategy spaces defined by player-specific matroids, showing that only a subclass of strategy profiles (called *decharged*) are candidates for being a pure Nash equilibrium. This allows to give lower bounds by constructing instances where all discharged profiles are expensive. Our characterization of Nash equilibria using the notion of discharged profiles strictly generalizes a characterization introduced in [von Falkenhausen and Harks \(2011\)](#). While [von Falkenhausen and Harks \(2011\)](#) assumed symmetric singleton strategy spaces, i.e., that a strategy consists only of a single resource and every player has access to every resource, we allow in this thesis player-specific matroid strategy spaces and, additionally, individual delay for each player and resource.

As a second contribution, we give in Section 3.4 an algorithm that constructs discharged profiles, allowing to establish a protocol with price of stability matching our lower bounds. In Section 3.5, we further prove that this protocol is also optimal with respect to the price of anarchy. The protocol we use for these positive results also fulfills the stricter separability requirement from [Chen et al. \(2010\)](#), i.e., when assigning the cost shares on a given resource, the protocol has no information about the load on other resources. For basic cost sharing protocols that do not need to fulfill the separability requirement, a stronger price of anarchy bound is possible when there are no delays and all players share a common strategy space. To this end, we introduce a second protocol, again based on the discharged strategy profiles. When the class of cost functions is restricted to be either concave or convex, we show in Section 3.6 a drastic improvement of the price of stability and price of anarchy. The results are summarized in Table 3.1.

Computational Complexity. In Section 3.7, we study the computational complexity of computing optimal cost shares that minimize the cost of the best/worst induced Nash equilibrium.

We prove that both problems are strongly NP-hard and there are no polynomial time $c \log(n)$ approximations for any $0 < c < 1$, unless $P = NP$. These hardness results even hold for instances with unweighted players, zero delays, singleton strategies and unit fixed costs. In light of the hardness even for this restricted class of problems, we study approximation algorithms for the case of unweighted players, zero delays and singleton strategies, still assuming general non-decreasing costs. This setting includes several interesting classes of problems such as scheduling applications, where each player is associated with a job of unit weight. The job can be processed on a job-specific set of machines, and the monetary cost on a resource (for instance energy costs as in Yao et al. 1995) is a non-decreasing function of its total load. Another application arises in capacitated facility location with delays in $\{0, \infty\}$. We devise a polynomial time algorithm computing cost shares with an approximation guarantee of \mathcal{H}_n and n for the problem of minimizing the cost of the best/worst Nash equilibrium, respectively.

Bibliographic Information. The results presented in this chapter are joint work with Tobias Harks. Parts of the results contained in this chapter appeared in *Mathematics of Operations Research*, see von Falkenhausen and Harks (2013a), and an abstract summarizing parts of the results appeared in *Proceedings of the 12th Cologne-Twente Workshop on Graphs and Combinatorial Optimization*, see von Falkenhausen and Harks (2013b).

3.2 Model and Problem Statement

In a *facility location model*, there is a set N of players that choose facilities (also called resources) from a set $R = \{r_1, \dots, r_m\}$. For each player $i \in N$, there is a set X_i of feasible resource choices, where a strategy $x_i \in X_i$ is a set of resources, i.e., $x_i \subseteq R$. Given strategies x_i for all players $i \in N$, we denote by $x = (x_i)_{i \in N}$ the joint strategy profile and correspondingly $X = \times_{i \in N} X_i$.

A vector $d = (d_i)_{i \in N}$ specifies the player's weights, i.e., the load each player imposes on a resource she uses. In a strategy profile x these loads sum up on each resource $r \in R$ to the load $\ell_r(x) = \sum_{i \in N_r(x)} d_i$, where $N_r(x) = \{i : r \in x_i\}$ is the set of players using r . The resources' cost functions are given by a vector $c = (c_r)_{r \in R}$, they are non-decreasing in the load. Each player i incurs a player-specific delay $t_{i,r} \geq 0$ when using resource r . Altogether, a *facility location model* is represented by a tuple $I = (N, R, X, d, c, t)$.

Unless noted otherwise, we assume that a player's strategy set X_i is the base set of some matroid \mathcal{M}_i . For a comprehensive introduction to matroid theory we refer to Schrijver (2003). We only state the properties necessary for our purposes.

Definition 3.1 (base set of a matroid)

A set system X_i with ground set R , i.e., where $x_i \subseteq R$ for all $x_i \in X_i$, is the base set of a matroid if it fulfills two conditions. First, X_i needs to be an *anti-chain*: for any $x_i, y_i \in X_i$, $x_i \subseteq y_i$ implies $x_i = y_i$. Second, it has the *basis exchange property*: given a set from X_i , one can alter it element by element towards another set from X_i such that all intermediate sets are also contained in X_i . In other words, for any sets $x_i, y_i \in X_i$ and any $r \in x_i$ there is some $r' \in y_i \setminus x_i$, $r \neq r'$, such that $x_i \setminus \{r\} \cup \{r'\} \in X_i$.

As we make frequent use of the exchange property, we denote a strategy profile x_i where resource r is exchanged by resource r' by $x_i + r' - r$.

Note that a facility location model is not a strategic game: a resource's cost function c_r is not equivalent to the cost that a player using this resource incurs. Instead, the cost $c_r(\ell_r(x))$ in a given strategy profile is distributed among the users $N_r(x)$ by a cost sharing protocol as defined below.

To measure the social cost of a profile x , we use utilitarian welfare, i.e., the sum of the individual players' costs and delays, and we denote $C(x) = \sum_{r \in R} c_r(\ell_r(x)) + \sum_{i \in N: r \in x_i} t_{i,r}$. Abusing notation, we often refer to the cost on resource r by $c_r(x)$ instead of $c_r(\ell_r(x))$. We now give two examples of such models.

Example 3.2 (capacitated facility location games). In a capacitated facility location game, motivated by the classic FACILITY LOCATION problem, every player i chooses exactly one resource, that is $|x_i| = 1$ for all $x_i \in X_i$ and $i \in N$ and, hence, \mathcal{M}_i corresponds to a uniform matroid of rank one. Hard capacities on the facilities can be modeled by cost functions that sharply increase at the corresponding capacity.

Example 3.3 (MST games). We are given an undirected graph $G = (V, E)$ with non-negative and non-decreasing edge cost functions $c_e, e \in E$. In a minimum spanning tree (MST) game, every player i is associated with demand of size $d_i > 0$ and a subgraph G_i of G . A strategy for player i is then to route its demand along a spanning tree for G_i . Formally, we set $R = E$ and the sets $x_i, i \in N$, are the spanning trees of G_i , hence \mathcal{M}_i corresponds to the graphical matroid.

We study how different ways of sharing the costs of a resource affect the resulting pure Nash equilibria of the induced game. To model this, we introduce a formalized notion of cost sharing protocols.

Definition 3.4 (cost sharing protocol)

A cost sharing protocol Ξ assigns to any facility location model (N, R, X, d, c, t) cost share functions $\xi_{i,r} : X \rightarrow \mathbb{R}$ for all $i \in N$ and $r \in R$ and thus *induces* a strategic game (N, X, ξ) where each player i 's private cost function is $\xi_i(x) = \sum_{r \in x_i} (\xi_{i,r}(x) + t_{i,r})$. Thus, a cost sharing protocol maps facility location models to strategic games.

We assume that every player strives to minimize her private cost, and, following our definition from Chapter 2, a strategy profile x is a Nash equilibrium if

$$\sum_{r \in x_i} (\xi_{i,r}(x) + t_{i,r}) \leq \sum_{r \in z_i} (\xi_{i,r}(z_i, x_{-i}) + t_{i,r}) \quad \text{for all strategies } z_i \in X_i \text{ and players } i \in N.$$

We quantify the efficiency of Nash equilibria using the price of anarchy and the price of stability. Extending Definition 2.12, we define for a cost sharing protocol Ξ by $POS(\Xi)$ and $POA(\Xi)$ the worst case price of stability and price of anarchy across games induced by protocol Ξ .

The main goal of this chapter is to design cost sharing protocols that minimize the price of anarchy and price of stability, respectively. Of course, the attainable objective values crucially depend on the design space that we permit. The following properties have been first proposed by [Chen et al. \(2010\)](#) in the context of designing cost sharing protocols for network design games.

Definition 3.5 (properties of cost sharing protocols)

A cost sharing protocol Ξ is

- (1) *stable* if it induces only games that admit at least one pure Nash equilibrium.
- (2) *basic* if it is stable and additionally *budget balanced*, i.e., if it assigns for all facility location models (N, R, X, d, c, t) cost share functions $\xi_{i,r}$ such that for all $r \in R$ and $x \in X$

$$c_r(x) = \sum_{i \in N_r(x)} \xi_{i,r}(x),$$

with $\xi_{i,r}(x) = 0$ for all $i \notin N_r(x)$. This property requires $c_r(0) = 0$ for unused resources, which we assume throughout the chapter.

- (3) *separable* if it is basic and if it induces only games for which in any two profiles $x, x' \in X$ for every resource $r \in R$,

$$N_r(x) = N_r(x') \Rightarrow \xi_{i,r}(x) = \xi_{i,r}(x') \text{ for all } i \in N_r(x).$$

Informally, separability means that in a profile x the values $\xi_{i,r}(x)$, $i \in N$, depend only on the set $N_r(x)$ of players sharing resource r and disregard all other information contained in x . Still, separable protocols can assign cost share functions that are specifically tailored to the given facility location model, for example based on an optimal profile. We denote by \mathcal{B}_n and \mathcal{S}_n the set of basic and separable protocols for facility location games with n players, respectively. We obtain the following optimization problems that we address in this chapter:

$$\min_{\Xi \in \mathcal{B}_n} POA(\Xi), \min_{\Xi \in \mathcal{B}_n} POS(\Xi), \min_{\Xi \in \mathcal{S}_n} POA(\Xi), \min_{\Xi \in \mathcal{S}_n} POS(\Xi).$$

3.3 Characterizing Pure Nash Equilibria

One of the main concepts of this chapter is that of *decharged* strategy profiles. Intuitively, a discharged profile has the minimum structure that is necessary for a Nash equilibrium. That is, a Nash equilibrium is necessarily discharged and for each discharged strategy profile there is a cost sharing protocol such that the profile is an equilibrium in the induced game. This allows to reduce the problem of quantifying the price of anarchy and price of stability to quantifying the possible cost of discharged strategy profiles, reducing the complexity of the design of sharing protocols. Moreover, the concept gives rise to protocols that make such profiles Nash equilibria.

The underlying idea of discharged strategy profiles is to compare the cost of a resource with the cost of the cheapest alternatives for each of the resource's users, as this is the maximum each user may pay in a Nash equilibrium. To this end, for a given facility location model I and strategy profile $z \in X$, denote the cost of the cheapest alternative of player i to resource $r \in z_i$ by

$$\Delta_i^r(z) = \min_{\substack{s \in R \\ z_i + s - r \in X_i}} (c_s(z_i + s - r, z_{-i}) + t_{i,s}).$$

Note that $\Delta_i^r(z)$ is defined independent of any cost sharing protocol, that is, we use the actual cost of the resource s , which is the maximum cost share a protocol can assign player i when she switches to s .

Definition 3.6 (decharged strategy profile)

A profile z of a facility location model I is discharged if it fulfills the following two properties:

- (1) Each player i is willing to pay the delay $t_{i,r}$ for each chosen resource $r \in z_i$:

$$t_{i,r} \leq \Delta_i^r(z). \quad (\text{D1})$$

- (2) The users $N_r(z)$ of each resource r , after having paid their delays, are willing to share the cost of the resource:

$$c_r(z) \leq \sum_{i \in N_r(z)} (\Delta_i^r(z) - t_{i,r}). \quad (\text{D2})$$

Note that (D1) implies that each summand $\Delta_i^r(z) - t_{i,r}$ in (D2) is non-negative.

Profiles that are not discharged are called *charged*, similarly resources are called discharged or charged depending on whether both (D1) and (D2) are fulfilled for all players using the resource.

Lemma 3.7. *Any pure Nash equilibrium x in a facility location game with a separable cost sharing protocol is discharged.*

Proof. Let x be a Nash equilibrium under a separable protocol ξ . We first show that (D1) holds for all players $i \in N$ and all resources $r \in x_i$. Let $s = \operatorname{argmin}_{s \in R: x_i + s - r \in X_i} c_s(x_i + s - r, x_{-i}) + t_{i,s}$ be the cheapest alternative for i to r and let $z_i = x_i + s - r$. Then the Nash inequality for player i can be restated as

$$\begin{aligned} t_{i,r} &\leq t_{i,r} + \sum_{r^* \in z_i} (\xi_{i,r^*}(z_i, x_{-i}) + t_{i,r^*}) - \sum_{r^* \in x_i} (\xi_{i,r^*}(x) + t_{i,r^*}) \\ &= \xi_{i,s}(z_i, x_{-i}) + t_{i,s} - (\xi_{i,r}(x) + t_{i,r}) + t_{i,r} \leq c_s(z_i, x_{-i}) + t_{i,s} = \Delta_i^r(x). \end{aligned} \quad (3.1)$$

For (3.1) observe that going from x_i to z_i only r and s are exchanged and since ξ is separable the cost shares for all other resources remain the same, i.e., $\xi_{i,r^*}(x) = \xi_{i,r^*}(z_i, x_{-i})$ for $r^* \in x_i \cap z_i$.

We now show that (D2) holds for all resources $r \in R$, again using the Nash inequality for player i ,

$$\begin{aligned} c_r(x) &= \sum_{i \in N_r(x)} \xi_{i,r}(x) \leq \sum_{i \in N_r(x)} \left(\min_{\substack{s \in R \\ x_i + s - r \in X_i}} (\xi_{i,s}(x_i + s - r, x_{-i}) + t_{i,s}) - t_{i,r} \right) \\ &\leq \sum_{i \in N_r(x)} \left(\min_{\substack{s \in R \\ x_i + s - r \in X_i}} (c_s(x_i + s - r, x_{-i}) + t_{i,s}) - t_{i,r} \right) = \sum_{i \in N_r(x)} (\Delta_i^r(x) - t_{i,r}) \end{aligned} \quad (3.2)$$

Here, (3.2) comes from budget-balance. \square

Not only are all pure Nash equilibria decharged, but we can also find cost shares for any decharged strategy profile such that it is a pure Nash equilibrium.

Definition 3.8 (x -enforcing cost shares)

Given a facility location model and and a decharged strategy profile x , cost shares ξ are x -enforcing if

$$\xi_{i,r}(x) = \frac{\Delta_i^r(x) - t_{i,r}}{\sum_{j \in N_r(x)} \Delta_j^r(x) - t_{j,r}} \cdot c_r(x) \quad \text{for } r \in x_i,$$

and

$$\xi_{i,r}(z_i, x_{-i}) = c_r(z_i, x_{-i}) \quad \text{for } r \in z_i \setminus x_i \text{ for any } z_i \in X_i.$$

In other words, in x the cost of each resource is shared among the users proportional to their cheapest alternatives minus the delay and a unilaterally deviating player pays the entire cost.

We show in Section 3.4 that there always exists a separable protocol with such cost shares.

Lemma 3.9. *Given a decharged profile x and x -enforcing cost shares ξ , x is a pure Nash equilibrium.*

Proof. To establish that x is a Nash equilibrium, we first show that no player can improve by unilaterally exchanging a single resource in her strategy.

$$\begin{aligned} \xi_{i,r}(x) + t_{i,r} &= \frac{\Delta_i^r(x) - t_{i,r}}{\sum_{j \in N_r(x)} \Delta_j^r(x) - t_{j,r}} \cdot c_r(x) + t_{i,r} \\ &\leq \Delta_i^r(x) - t_{i,r} + t_{i,r} = \Delta_i^r(x) \end{aligned} \tag{3.3}$$

where (3.3) holds because for decharged profiles we have (D2) and hence the denominator of the cost share is not smaller than $c_r(x)$.

Having shown that no player can improve her cost by exchanging a single resource in her strategy, we use the basis exchange property of matroids as introduced in Section 3.2 to conclude that no player can improve by changing her strategy to an arbitrary $z_i \in X_i$ and, consequently, that x is a Nash equilibrium. To this end, fix such a $z_i \in X_i$ and denote by $G(x_i \triangle z_i)$ the bipartite graph (V, E) with $V = (x_i \setminus z_i) \cup (z_i \setminus x_i)$ and $E = \{(r, s) : r \in x_i \setminus z_i, s \in z_i \setminus x_i, (x_i + s - r) \in X_i\}$.

Proposition 3.10 (Schrijver 2003). *There exists a perfect matching in the graph $G(x_i \triangle z_i)$.*

Consider such a matching and observe that, as shown above, no player can improve unilaterally by exchanging a single resource across a matching edge (r, s) with $r \in x_i, s \in z_i$,

$$\xi_{i,r}(x) + t_{i,r} \leq \xi_{i,s}(x_i + s - r, x_{-i}) + t_{i,s} = \xi_{i,s}(z_i, x_{-i}) + t_{i,s},$$

where in the last step we use that player i 's cost share on s is independent of the other resources. Summing this up across all matching edges yields the desired

$$\xi_i(x) + \sum_{r \in x_i} t_{i,r} = \sum_{r \in x_i} (\xi_{i,r}(x) + t_{i,r}) \leq \sum_{s \in z_i} (\xi_{i,s}(z_i, x_{-i}) + t_{i,s}) = \xi_i(z_i, x_{-i}) + \sum_{s \in z_i} t_{i,s}.$$

□

Theorem 3.11 (Characterization of Pure Nash Equilibria for Separable Protocols). *Given a strategy profile $x \in X$ of a facility location model, there is a separable protocol such that x is a pure Nash equilibrium in the induced game if and only if x is decharged.*

Proof. Follows from Lemma 3.7 and Lemma 3.9, the existence of the separable protocol is the content of Section 3.4. \square

3.4 An Optimal Protocol for the Price of Stability

In this section, we deal with the existence and cost of decharged strategy profiles, which by Theorem 3.11 correspond to the possible pure Nash equilibria of a facility location game. In von Falkenhausen and Harks (2011) we showed that even in the very limited setting of symmetric singleton games without delays, there are instances where the cheapest decharged profile costs \mathcal{H}_n times as much as an optimal profile. Additionally, for these games we showed that the characterization of Nash equilibria as decharged profiles also holds for basic protocols, resulting in a lower bound for both classes of protocols. This lower bound also applies to the more general setting addressed here.

Lemma 3.12. [von Falkenhausen and Harks 2011] *The price of stability for symmetric singleton games without delays induced by basic or separable protocols is at least \mathcal{H}_n , i.e.,*

$$\min_{\Xi \in \mathcal{B}_n} POS(\Xi) \geq \mathcal{H}_n, \quad \min_{\Xi \in \mathcal{S}_n} POS(\Xi) \geq \mathcal{H}_n.$$

Symmetric singleton games are facility location games with the restrictions that $X_i = X_j = R$, i.e., the strategies of the players are single resources, and $t_{i,r} = 0$ for all $i, j \in N$ and $r \in R$. Consequently, the price of stability for facility location games induced by basic or separable protocols is at least \mathcal{H}_n .

To match this lower bound, we present an algorithm that ‘decharges’ any strategy profile of a facility location model, increasing the cost in the process by at most a factor \mathcal{H}_n . Finally, we give a separable protocol that is x -enforcing for the output of the algorithm. Using this protocol and starting the algorithm at an optimal strategy profile of a facility location model gives an upper bound on the price of stability.

3.4.1 Finding Cheap Decharged Strategy Profiles Algorithmically

Algorithm 1 is the cornerstone of the upper bound on the price of stability presented in this section. The algorithm returns for every instance (N, R, X, d, c, t) and input profile y a decharged profile x that costs no more than $\mathcal{H}_n \cdot C(y)$, as shown in the upcoming lemmata.

The algorithm iteratively moves players away from charged resources. If there is a player that is not even willing to pay her delay (case called *Transp*, lines 5-6), this player is immediately moved away. Otherwise, if there are players that have been moved before (case called *Shuffle*, lines 7-8) these players are moved in a last-in-first-out order. Finally, if all players are willing to pay a non-negative cost share and none of them have been moved before (case called *Kickoff*, lines 9-11), the one who is willing to pay the least is moved. In each iteration, the selected

Algorithm 1 Find decharged strategy profile x **Input:** Facility location model (N, R, X, d, c, t) , profile y **Output:** Decharged profile x

```

1:  $k \leftarrow 1$       {step number}
2:  $x^1 \leftarrow y$     {starts with profile  $y$ }
3: while there are charged resources do
4:   select the most expensive charged resource  $r^k \leftarrow \operatorname{argmax}_{r \in R \text{ is charged}} c_r(x^k)$ 
5:   if (D1) not fulfilled by all players on  $r^k$  then    {case called Transp}
6:     select such a player  $i^k \in N_{r^k}(x^k)$  with  $t_{i^k, r^k} > \Delta_{i^k}^{r^k}(x^k)$ 
7:   else if some player on  $r^k$  was moved before then    {case called Shuffle}
8:     select player  $i^k \in N_{r^k}(x^k)$  that was moved last
9:   else    {case called Kickoff}
10:    select player who is willing to pay the least
11:     $t^k \leftarrow \operatorname{argmin}_{i \in N_{r^k}(x^k)} \Delta_{i^k}^{r^k}(x^k) - t_{i^k, r^k}$ 
12:  end if
13:  select cheapest (1,1)-exchange  $s^k \leftarrow \operatorname{argmin}_{s \in R} c_s(x_{i^k}^k + s - r^k, x_{-i^k}^k) + t_{i^k, s}$ 
14:  execute (1,1)-exchange  $x^{k+1} \leftarrow (x_{i^k}^k + s^k - r^k, x_{-i^k}^k)$ 
15:  iterate  $k \leftarrow k + 1$ 
16: end while
17: return  $x \leftarrow x^k$ 

```

player is moved to the best available (1,1)-exchange (line 14). To show that the algorithm works as desired, we prove two lemmas: First, we show that the algorithm terminates, then we deal with the cost of the returned profile x .

Lemma 3.13. *Algorithm 1 terminates.*

Proof. For the proof, we adhere to the interpretation of each player i scheduling $|y_i|$ jobs on the resources. We fix a player and follow one of her jobs over the course of the algorithm. The job can be moved at most once by a Kickoff, afterwards multiple times by Transp and Shuffle moves. We show that these moves strictly decrease the cost of the job, independent of what happens in between the moves. More precisely, if a job of player i is moved in iteration k of the algorithm to resource s^k and stays there until iteration l (i.e., $s^k = r^l$) when it is moved to resource s^l , we show that $c_{s^k}(x^{k+1}) + t_{i, s^k} > c_{s^l}(x^{l+1}) + t_{i, s^l}$. Since there are only finitely many values for the cost of the job, this proves that the algorithm terminates.

If the move in iteration l is a Transp move, then the cost of the delay t_{i, r^l} is greater than the cost of the cheapest (1,1)-exchange and we have

$$c_{s^k}(x^{k+1}) + t_{i, s^k} \geq t_{i, r^l} = t_{i, r^l} > \Delta_i^{r^l}(x^l) = c_{s^l}(x^{l+1}) + t_{i, s^l}. \quad (3.4)$$

If on the other hand the move in iteration l is a Shuffle move, the last-in-first-out scheme of

Shuffles ensures that $N_{r^l}(x^{k+1}) \supseteq N_{r^l}(x^l)$ and hence

$$\begin{aligned} c_{s^k}(x^{k+1}) &\geq c_{r^l}(x^l) \\ &> \sum_{j \in N_{r^l}(x^l)} (\Delta_j^{r^l}(x^l) - t_{j,r^l}) \end{aligned} \quad (3.5)$$

$$\geq \Delta_i^{r^l}(x^l) - t_{i,r^l} = c_{s^l}(x^{l+1}) + t_{i,s^l} - t_{i,r^l}. \quad (3.6)$$

When the algorithm does a Shuffle on resource r^l , all users of the resource fulfill (D1) (otherwise a Transp would be done), but the resource does not fulfill (D2) (otherwise it would be discharged and hence be disregarded by the algorithm). Not fulfilling (D2) leads to (3.5) and since every user fulfills (D1) all summands in (3.5) are non-negative, which leads to (3.6). \square

Lemma 3.14. *Profile x returned by Algorithm 1 has at most \mathcal{H}_n times the cost of the input profile y .*

Proof. Throughout the proof of this lemma, we regard the set $Q = \{q_1, \dots\}$ of all jobs instead of the players they belong to. We denote the player to which a job q belongs by $i(q)$ and the resource on which job q is scheduled in profile y by $y(q)$. For ease of exposition, we often use jobs q interchangeably with $i(q)$, e.g. $t_{q,r}$ as a shorthand for $t_{i(q),r}$. If k is the iteration in which q is first moved by the algorithm, we define $p(q) = |N_{y(q)}(x^k)|$ to be the number of players on $y(q)$ in x^k . This first move can either be a Transp or a Kickoff. For a Transp, we have already seen in (3.4) that $c_{s^k}(x^{k+1}) + t_{q,s^k} < t_{q,y(q)}$. If the first move is a Kickoff, then the set of jobs using the resource $y(q)$ in x^k is a subset of the jobs using it in y , i.e., $N_{y(q)}(x^k) \subseteq N_{y(q)}(y)$, and we have

$$\begin{aligned} c_{y(q)}(y) &\geq c_{y(q)}(x^k) \\ &> \sum_{j \in N_{y(q)}(x^k)} (\Delta_j^{y(q)}(x^k) - t_{j,y(q)}) \end{aligned} \quad (3.7)$$

$$\geq p(q) \cdot (\Delta_q^{y(q)}(x^k) - t_{q,y(q)}) = p(q) \cdot (c_{s^k}(x^{k+1}) + t_{q,s^k} - t_{q,y(q)}). \quad (3.8)$$

(3.7) stems like (3.5) from the fact that in a Kickoff, the resource $y(q)$ does not fulfill (D2). Since all summands in (3.7) are positive (like in (3.5)) and the summand corresponding to job q is the smallest of the $p(q)$ summands (see line 11 of the algorithm), we have inequality (3.8).

Altogether, the first time a job is moved by the algorithm, we have

$$c_{s^k}(x^{k+1}) + t_{q,s^k} \leq \frac{1}{p(q)} \cdot c_{y(q)}(y) + t_{q,y(q)},$$

regardless of whether this move is a Kickoff or a Transp. Further Transp and Shuffle moves only reduce this quantity, hence for every further iteration \bar{k} of the algorithm where job q is on resource \bar{r} and no job has been moved there after him, we have

$$c_{\bar{r}}(x^{\bar{k}}) + t_{q,\bar{r}} \leq \frac{1}{p(q)} \cdot c_{y(q)}(y) + t_{q,y(q)}. \quad (3.9)$$

Particularly, for resources where jobs have been moved by the algorithm, the cost in the final profile x returned by the algorithm is determined by the last job that was moved to the resource in the sense of (3.9). In the following, we estimate the cost of such resources by summing up (3.9) over all resources. For resources where no job was moved, the cost is no greater than in the original profile y . Thus,

$$\begin{aligned} C(x) &= \sum_{\substack{r \in R \\ \text{jobs moved to } r}} \left(c_r(x) + \sum_{q \in N_r(x)} t_{q,r} \right) + \sum_{\substack{r \in R \\ \text{no jobs moved to } r}} \left(c_r(x) + \sum_{q \in N_r(x)} t_{q,r} \right) \\ &\leq \sum_{\substack{r \in R \\ \text{jobs moved to } r}} \sum_{q \in N_r(x)} \left(\frac{1}{p(q)} \cdot c_{y(q)}(y) + t_{q,y(q)} \right) + \sum_{\substack{r \in R \\ \text{no jobs moved to } r}} \left(c_r(y) + \sum_{q \in N_r(x)} t_{q,y(q)} \right) \\ &\leq \sum_{r \in R} \sum_{\substack{q \in N_r(y) \\ q \text{ moved}}} \frac{1}{p(q)} \cdot c_r(y) + \sum_{\substack{r \in R \\ \text{jobs remained on } r}} c_r(y) + \sum_{\substack{\text{all jobs } q}} t_{q,y(q)} \end{aligned} \quad (3.10)$$

$$\leq \sum_{r \in R} \mathcal{H}_n \cdot c_r(y) + \sum_{\substack{\text{all jobs } q}} t_{q,y(q)} \leq \mathcal{H}_n \cdot C(y). \quad (3.11)$$

In (3.10) we change the order of summation: before we grouped the jobs by the resource they are on in x , in (3.10) they are grouped by the resource they are on in y . For each resource r the $\frac{1}{p(q)}$ fractions sum up to at most \mathcal{H}_n or, if jobs remained on r , to at most $\mathcal{H}_n - 1$ by definition of $p(q)$, hence (3.11). \square

Corollary 3.15. *Every facility location model has a decharged strategy profile x at cost $C(x) \leq \mathcal{H}_n \cdot C(y)$, where y is a cost-optimal strategy profile.*

Proof. Follows from Lemma 3.13 and Lemma 3.14 \square

We now use the algorithm together with our insights about x -enforcing cost shares to construct a protocol that matches our lower bound for the price of stability. For this, we assume without loss of generality that the players are indexed by non-increasing weights $d_1 \geq d_2 \geq \dots \geq d_n$.

Definition 3.16 (enforcing protocol)

For any facility location model (N, R, X, d, t, c) , the enforcing protocol runs Algorithm 1 with a cost-optimal profile y as input to obtain a decharged profile x with cost $C(x) \leq \mathcal{H}_n \cdot C(y)$. Then, define for any profile z and resource r the set of *foreign players* $N_r^1(z) = N_r(z) \setminus N_r(x)$ and assign the cost share functions

$$\xi_{i,r}(z) = \begin{cases} \frac{\Delta_i^r(x) - t_{i,r}}{\sum_{j \in N_r(x)} \Delta_j^r(x) - t_{j,r}} \cdot c_r(x), & \text{if } r \in z_i, N_r(z) = N_r(x) \text{ and } c_r(x) > 0, \\ c_r(z), & \text{if } N_r^1(z) \neq \emptyset \text{ and } i = \min N_r^1(z), \\ c_r(z), & \text{if } N_r^1(z) = \emptyset, N_r(z) \subset N_r(x) \text{ and } i = \min N_r(z), \\ 0, & \text{else.} \end{cases}$$

One can easily verify that the protocol is budget-balanced and separable. The cost shares are x -enforcing and hence the price of stability for the enforcing protocol is \mathcal{H}_n .

Theorem 3.17. *The price of stability for facility location games induced by basic and separable protocols is \mathcal{H}_n , i.e.,*

$$\min_{\Xi \in \mathcal{B}_n} POS(\Xi) = \mathcal{H}_n, \quad \min_{\Xi \in \mathcal{S}_n} POS(\Xi) = \mathcal{H}_n.$$

Proof. The lower bound is given by Lemma 3.12. One can easily verify that the enforcing protocol is budget-balanced and separable. Furthermore, it is x -enforcing for a decharged profile x returned by Algorithm 1 with $C(x) \leq \mathcal{H}_n \cdot C(y)$ and hence the price of stability for the enforcing protocol is \mathcal{H}_n . \square

3.5 An Optimal Protocol for the Price of Anarchy

We now turn to the problem of finding a cost sharing protocol that minimizes the price of anarchy of the games it induces. [Von Falkenhausen and Harks \(2011\)](#) have shown that the price of anarchy for basic and separable protocols in symmetric singleton games without delays is \mathcal{H}_n . Relaxing these assumptions quickly raises the price of anarchy: we construct lower bound instances showing a price of anarchy of at least n for asymmetric singleton games without delays and symmetric singleton games with delays. We then give a complementing an upper bound, using our separable enforcing protocol in asymmetric matroid games with delays.

At the end of this section we deal with the special case of symmetric matroid games without delays. For such games, there is a basic protocol that gives a price of anarchy of \mathcal{H}_n , while we show that for separable protocols the price of anarchy is n .

Lemma 3.18. *The price of anarchy for facility location games with basic or separable cost sharing protocols is at least n ,*

$$\min_{\Xi \in \mathcal{B}_n} POA(\Xi) \geq n, \quad \min_{\Xi \in \mathcal{S}_n} POA(\Xi) \geq n.$$

Proof. Consider the facility location model (N, R, X, d, c, t) with n players $i \in N$, resource set $R = \{r_0, r_1, \dots, r_n\}$, strategy spaces $X_i = \{\{r_0\}, \{r_i\}\}$, weights $d_i = 1$ for all $i \in N$, delays $t_{i,r} = 0$ for all $i \in N, r \in R$ and constant resource cost $c_r \equiv 1$ independent of the load for all $r \in R$. Here, the optimal profile $y = (\{r_0\}, \dots, \{r_0\})$ has cost $C(y) = 1$. However, the profile $x = (\{r_1\}, \dots, \{r_n\})$ with cost $C(x) = n$ is a Nash equilibrium under any protocol that satisfies budget-balance. \square

Remark 3.19. Note that the above bound can equivalently be constructed with a symmetric singleton model with delays as follows: Let $X_i = \{\{r_0\}, \dots, \{r_n\}\}$ for all players $i \in N$ be the symmetric strategy spaces and assign delays $t_{i,r_0} = 0, t_{i,r_i} = 0$ and $t_{i,r} = 2$ for all $r \in R \setminus \{r_0, r_i\}$. Then, as above, y is an optimal profile and x is a Nash equilibrium under any budget balanced protocol.

Lemma 3.20. *The price of anarchy for facility location games with basic or separable protocols is at most n*

$$\min_{\Xi \in \mathcal{B}_n} POA(\Xi) \leq n, \quad \min_{\Xi \in \mathcal{S}_n} POA(\Xi) \leq n.$$

Proof. Let (N, R, X, d, c, t) be a facility location model. Let $\xi_{i,r}$ for $i \in N, r \in R$ be the cost share functions assigned by the enforcing protocol and let x be the decharged profile returned by Algorithm 1 for the protocol with intermediate profiles x^1, x^2, \dots, x and some optimal profile y as input. We show $C(z) \leq n \cdot C(y)$ for any pure Nash equilibrium z . In a first step, we link the cost in z to the cost in profiles (x_i, z_{-i}) via the Nash property. The major challenge of the proof is then to estimate the cost of these profiles in relation to the cost of y . To this end we employ properties of x and the intermediate profiles found in the analysis of the algorithm in the previous section.

In the following, we use notation for players and jobs interchangeably, denoting jobs by the letter q , for example $q \in N_r(x)$ for the jobs on a resource r . For each job q , define $y(q)$ and $p(q)$ as in the analysis of the algorithm and denote additionally by $x(q)$ the resource job q is on in profile x and by x^q the algorithm's intermediate profile in which q is first on $x(q)$. From the analysis of the algorithm, we know

$$c_{x(q)}(x^q) + t_{q,x(q)} \leq \frac{1}{p(q)} c_{y(q)}(y) + t_{q,y(q)} \quad (3.12)$$

for all jobs q that were moved by the algorithm. For jobs q that were not moved by the algorithm, we set $x^q = y$.

To prove the lemma, we show $C(z) \leq n \cdot C(y)$ for any pure Nash equilibrium z . To this end, we fix such a profile z and link it to the profiles (x_i, z_{-i}) via the Nash property,

$$\begin{aligned} C(z) &= \sum_{i \in N} \left(\xi_i(z) + \sum_{r \in z_i} t_{i,r} \right) \leq \sum_{i \in N} \left(\xi_i(x_i, z_{-i}) + \sum_{r \in x_i} t_{i,r} \right) = \sum_{i \in N} \left(\sum_{\substack{r \in x_i \\ N_r^1(z) = \emptyset}} \xi_{i,r}(x_i, z_{-i}) + \sum_{r \in x_i} t_{i,r} \right) \\ &= \sum_{\substack{r \in R \\ N_r^1(z) = \emptyset}} \sum_{i \in N_r(x)} (\xi_{i,r}(x_i, z_{-i}) + t_{i,r}) + \sum_{\substack{r \in R \\ N_r^1(z) \neq \emptyset}} \sum_{i \in N_r(x)} t_{i,r} \\ &\leq \sum_{\substack{r \in R \\ N_r^1(z) = \emptyset}} \left(\sum_{q \in N_r(x) \cap N_r(y)} (c_r(y) + t_{q,r}) + 2 \cdot \sum_{q \in N_r(x) \setminus N_r(y)} \left(\frac{1}{p(q)} c_{y(q)}(y) + t_{q,y(q)} \right) \right) + \sum_{\substack{r \in R \\ N_r^1(z) \neq \emptyset}} \sum_{q \in N_r(x)} t_{q,r} \end{aligned} \quad (3.13)$$

Proving (3.13) is a major challenge of this proof and beforehand we give a brief intuition for this inequality: for jobs that are moved by the algorithm we have an at most logarithmic cost-increase going from profile y to profile z , represented by the second term, while for jobs not moved by the algorithm, the cost-increase can even be linear as represented by the first term. In our worst-case example in Lemma 3.18, this linear cost-increase dominates the logarithmic cost-increase: no jobs are moved by the algorithm.

To prove (3.13), we partition the resources without foreign players into two sets,

- $R_1 = \{r \in R : N_r^1(z) = \emptyset \text{ and } |N_r(x) \setminus N_r(z)| \leq 1\}$ - resources where at most one job is missing,
- $R_2 = \{r \in R : N_r^1(z) = \emptyset \text{ and } |N_r(x) \setminus N_r(z)| > 1\}$ - resources, where multiple jobs are missing.

For the resources in both sets, we find bounds corresponding to (3.13) in two separate claims. Afterwards we combine the two claims to prove the lemma.

Claim 3.21. *For $r \in R_1$,*

$$\sum_{i \in N_r(x)} (\xi_{i,r}(x_i, z_{-i}) + t_{i,r}) \leq \sum_{q \in N_r(x) \cap N_r(y)} (c_r(y) + t_{q,r}) + 2 \cdot \sum_{q \in N_r(x) \setminus N_r(y)} \left(\frac{1}{p(q)} c_{y(q)}(y) + t_{q,y(q)} \right).$$

Proof. Recall that for all $r \in R$

$$c_r(x) \leq \begin{cases} c_r(y) & \text{if } N_r(x) \setminus N_r(y) = \emptyset, \\ c_r(x^{q_r}) \leq \frac{1}{p(q_r)} c_{y(q_r)}(y) + t_{q_r,y(q_r)} - t_{q_r,r} & \text{if } N_r(x) \setminus N_r(y) \neq \emptyset, \end{cases} \quad (3.14)$$

$$(3.15)$$

where job q_r denotes the last job moved to r by the algorithm. The second inequality (3.15) follows from (3.12). To prove the claim, we have for $r \in R_1$,

$$\sum_{i \in N_r(x)} \xi_{i,r}(x_i, z_{-i}) + t_{i,r} = \mathbf{1}_{N_r(z) \neq \emptyset} \sum_{i \in N_r(z)} \xi_{i,r}(z) + \mathbf{1}_{N_r(x) \setminus N_r(z) = \{i^*\}} \xi_{i^*,r}(x_{i^*}, z_{-i^*}) + \sum_{i \in N_r(x)} t_{i,r} \quad (3.16)$$

$$\leq \mathbf{1}_{N_r(z) \neq \emptyset} c_r(x) + \mathbf{1}_{N_r(x) \setminus N_r(z) = \{i^*\}} c_r(x) + \sum_{i \in N_r(x)} t_{i,r} \quad (3.17)$$

$$\leq \sum_{q \in N_r(x) \cap N_r(y)} c_r(y) + 2 \cdot \sum_{q \in N_r(x) \setminus N_r(y)} \left(\frac{1}{p(q)} c_{y(q)}(y) + t_{q,y(q)} - t_{q,r} \right) + \sum_{q \in N_r(x)} t_{q,r} \quad (3.18)$$

$$\leq \sum_{q \in N_r(x) \cap N_r(y)} (c_r(y) + t_{q,r}) + 2 \cdot \sum_{q \in N_r(x) \setminus N_r(y)} \left(\frac{1}{p(q)} c_{y(q)}(y) + t_{q,y(q)} \right),$$

where $\mathbf{1}$ denotes the indicator function tied to the condition in subscript and, if applicable, i^* is the single player using resource r in x but not in z . Equation (3.16) is due to the nature of R_1 and for (3.17) we use that there are no foreign players on resources $r \in R_1$. In inequality (3.18), we estimate using both (3.14) and (3.15). For the case where both indicator functions are true, we multiply the term from (3.15) by 2. For the term from (3.14) this is not necessary because, if both indicator functions are true and $N_r(x) \setminus N_r(y) = \emptyset$, then there are multiple jobs $q \in N_r(x) \cap N_r(y)$ and, hence, $\sum_{q \in N_r(x) \cap N_r(y)} c_r(y) \geq 2 \cdot c_r(y)$. \square

Claim 3.22. *For $r \in R_2$,*

$$\sum_{i \in N_r(x)} \xi_{i,r}(x_i, z_{-i}) + t_{i,r} \leq \sum_{q \in N_r(x) \cap N_r(y)} (c_r(y) + t_{i,r}) + \sum_{q \in N_r(x) \setminus N_r(y)} \left(\frac{1}{p(q)} c_{y(q)}(y) + t_{q,y(q)} \right).$$

Proof. We denote the jobs on $r \in R_2$ in profile x by $q_1^r, \dots, q_{|N_r(x)|}^r$ such that they are indexed with non-increasing weights $d_{q_1^r} \geq \dots \geq d_{q_{|N_r(x)|}^r}$. Let $s(r) = \min\{i : q_i^r \in N_r(z)\}$. Since the jobs are indexed in the same order as their players, the protocol assigns for $i \leq |N_r(x)|$

$$\xi_{q_i^r, r}(x_{q_i^r}, z_{-q_i^r}) = \begin{cases} c_r(\ell_r(z) + d_{q_i^r}) & \text{if } i < s(r), \\ c_r(\ell_r(z)) & \text{if } i = s(r), \\ 0 & \text{if } i > s(r). \end{cases} \quad (3.19)$$

We now define an automorphism $\sigma_r : \{q_1^r, \dots, q_{|N_r(x)|}^r\} \rightarrow \{q_1^r, \dots, q_{|N_r(x)|}^r\}$ that maps the first $t(r) = |N_r(x) \setminus N_r(y)|$ jobs (by index) to $N_r(x) \setminus N_r(y)$, such that

- $\sigma_r(q_1^r)$ is the last job that was moved to r by the algorithm,
- $\sigma_r(q_2^r)$ is the second-last job that was moved to r by the algorithm,
- ...
- $\sigma_r(q_{t(r)}^r)$ is the first job that was moved to r by the algorithm.

The remaining jobs are mapped arbitrarily to $N_r(x) \cap N_r(y)$, keeping σ_r bijective. Then,

$$\ell_r(z) \leq \ell_r(x) - \sum_{j=1}^{s(r)-1} d_{q_j^r} \quad (3.20)$$

$$\leq \ell_r(x) - \sum_{j=1}^{s(r)-1} d_{\sigma_r(q_j^r)} \quad (3.21)$$

$$\leq \ell_r(x^{\sigma_r(q_{s(r)}^r)}), \quad (3.22)$$

where (3.20) holds because $N_r(z) \subset N_r(x)$ and $q_1^r, \dots, q_{s(r)-1}^r \notin N_r(z)$ by definition of $s(r)$. Inequality (3.21) holds because we indexed the jobs from big to small and hence the first $s(r) - 1$ jobs are the ‘biggest’ jobs on resource r . For (3.22), if $s(r) \leq t(r)$, that is, if $\sigma_r(q_{s(r)}^r)$ was moved to resource r , then in profile $x^{\sigma_r(q_{s(r)}^r)}$ none of the jobs $\sigma_r(q_{s(r)-1}^r), \dots, \sigma_r(q_1^r)$ moved to r after job $\sigma_r(q_{s(r)}^r)$ are on resource r , and consequently (3.22) follows. Otherwise, if $s(r) > t(r)$, that is, if $\sigma_r(q_{s(r)}^r) \in N_r(y)$, then $y = x^{\sigma_r(q_{s(r)}^r)}$ and in this profile none of the jobs $\sigma_r(q_{t(r)}^r), \dots, \sigma_r(q_1^r)$ that were moved to resource r are on resource r and hence (3.22) follows. We find likewise for $i < s(r)$,

$$\ell_r(z) + d_{q_i^r} \leq \ell_r(x) - \sum_{j=1}^{i-1} d_{q_j^r} \leq \ell_r(x) - \sum_{j=1}^{i-1} d_{\sigma_r(q_j^r)} \leq \ell_r(x^{\sigma_r(q_i^r)}), \quad (3.23)$$

where the above inequalities hold for similar reasons as (3.20), (3.21), and (3.22). We complete

the proof of the claim by

$$\sum_{i \in N_r(x)} (\xi_{i,r}(x_i, z_{-i}) + t_{i,r}) = c_r(\ell_r(z)) + \sum_{i=1}^{s(r)-1} c_r(\ell_r(z) + d_{q_i^r}) + \sum_{i \in N_r(x)} t_{i,r} \quad (3.24)$$

$$\leq \sum_{i=1}^{s(r)} c_r(\ell_r(x^{\sigma_r(q_i^r)})) + \sum_{i \in N_r(x)} t_{i,r} \leq \sum_{q \in N_r(x)} (c_r(x^{\sigma_r(q)}) + t_{i,r}) \quad (3.25)$$

$$= \sum_{q \in N_r(x)} (c_r(x^q) + t_{i,r}) \quad (3.26)$$

$$\leq \sum_{q \in N_r(x) \cap N_r(y)} (c_r(y) + t_{i,r}) + \sum_{q \in N_r(x) \setminus N_r(y)} \left(\frac{1}{p(q)} c_{y(q)}(y) + t_{q,y(q)} \right), \quad (3.27)$$

where equation (3.24) follows from (3.19) and inequality (3.25) follows from (3.22) and (3.23). Equation (3.26) holds because σ_r is an automorphism on $N_r(x)$ and finally inequality (3.27) follows from our definition of the intermediate profiles x^q and our results regarding these profiles as in (3.12). \square

We now continue the proof of Lemma 3.20 where we left off with (3.13) and conclude across both sets,

$$\begin{aligned} C(z) &\leq \sum_{\substack{r \in R \\ N_r^1(z) = \emptyset}} \sum_{i \in N_r(x)} (\xi_{i,r}(x_i, z_{-i}) + t_{i,r}) + \sum_{\substack{r \in R \\ N_r^1(z) \neq \emptyset}} \sum_{i \in N_r(x)} t_{i,r} \\ &\leq \sum_{\substack{r \in R \\ N_r^1(z) = \emptyset}} \left(\sum_{q \in N_r(x) \cap N_r(y)} (c_r(y) + t_{q,r}) + 2 \cdot \sum_{q \in N_r(x) \setminus N_r(y)} \left(\frac{1}{p(q)} c_{y(q)}(y) + t_{q,y(q)} \right) \right) + \sum_{\substack{r \in R \\ N_r^1(z) \neq \emptyset}} \sum_{q \in N_r(x)} t_{q,r} \end{aligned} \quad (3.28)$$

$$\leq \sum_{r \in R} \left(\sum_{q \in N_r(x) \cap N_r(y)} (c_r(y) + t_{q,r}) + \sum_{q \in N_r(y) \setminus N_r(x)} \left(\frac{2}{p(q)} c_r(y) + 2 \cdot t_{q,r} \right) \right) \quad (3.29)$$

$$\leq \sum_{r \in R} \left(|N_r(y) \cap N_r(x)| \cdot c_r(y) + \sum_{p=|N_r(y) \cap N_r(x)|+1}^{|N_r(y)|} \frac{2}{p} \cdot c_r(y) + 2 \cdot \sum_{q \in N_r(y)} t_{q,r} \right) \quad (3.30)$$

$$\leq \sum_{r \in R} \left(|N_r(y)| \cdot c_r(y) + 2 \cdot \sum_{q \in N_r(y)} t_{q,r} \right) \leq n \cdot C(y).$$

Here, (3.28) follows from Claims 3.21 and 3.22. In (3.29), we change the order of summation: instead summing up the $q \in N_r(x) \setminus N_r(y)$ that were moved *from* other resources by the algorithm, we sum up the $q \in N_r(y) \setminus N_r(x)$ that were moved *to* other resources by the algorithm. At the same time we extend summation across all resources $r \in R$ where before we only summed up across resources with $N_r^1(z) = \emptyset$. For (3.30), recall how we introduced $p(q)$: the first job that is

moved away has $p(q) = |N_r(y)|$, the next has $p(q) = |N_r(y) - 1|$ until the last job that is moved away has $p(q) = |N_r(y) \cap N_r(x)| + 1$. \square

Corollary 3.23. *The price of anarchy for facility location games with basic or separable protocols is n ,*

$$\min_{\Xi \in \mathcal{B}_n} POA(\Xi) = n, \quad \min_{\Xi \in \mathcal{S}_n} POA(\Xi) = n.$$

3.5.1 Distinction between Basic and Separable Protocols: Symmetric Matroid Games without Delays

Above, we have shown that both basic and separable protocols can give a price of anarchy of n in the very general setting of asymmetric matroid games with delays. Our matching lower bounds covered asymmetric singleton games without delays and symmetric singleton games with delays. The case left open is relaxing the singleton assumption to matroid strategy spaces while maintaining the symmetry and no-delay assumptions, i.e., symmetric matroid games without delays. Such instances expose the first distinction in the performance of basic and separable protocols: We propose in this section a lower bound with price of anarchy of n for separable protocols and introduce a basic protocol that guarantees price of anarchy of \mathcal{H}_n .

In the following, we denote a *symmetric matroid model* by $I = (N, R, X, d, c)$, which is similar to a facility location model with $X_i = X_j$ for all $i, j \in N$ and $t_{i,r} = 0$ for all $i \in N$ and $r \in R$.

Lemma 3.24. *For symmetric matroid games,*

$$\min_{\Xi \in \mathcal{S}_n} POA(\Xi) \geq n.$$

This lower bound holds even for models with unit demands and uniform matroids.

Proof. Consider the symmetric matroid model (N, R, X, d, c) with n players that have unit demand $d_i = 1$ for all $i \in N$ and resources $R = \{r_0, r_1, \dots, r_n\}$ with cost functions as in Table 3.2. All players $i \in N$ have identical strategy sets $X_i = \{z_i \subset R : |z_i| = n\}$. Note that the sets in X_i are the bases of the uniform matroid on R of rank n .

load	$c_{r_0}(\ell)$	$c_{r_1}(\ell)$...	$c_{r_n}(\ell)$
$\ell = 0$	0	0	...	0
$0 < \ell < n$	1	0	...	0
$\ell = n$	1	1	...	1

Table 3.2: Cost functions for resources used in the proof of Lemma 3.24.

Consider the profile x with $x_i = \{r_1, \dots, r_n\}$ for all players $i \in N$. We show that x is a pure Nash equilibrium for any separable protocol. First, note that any strategy $z_i \in X_i$, $z_i \neq x_i$, can be

written as $z_i = x_i + r_0 - r_j$ for some $j \leq n$. Then,

$$\sum_{r \in x_i} \xi_{i,r}(x) \leq \sum_{r \in x_i} \xi_{i,r}(x) - \xi_{i,r_j}(x) + \xi_{i,r_0}(z_i, x_{-i}) \quad (3.31)$$

$$= \sum_{r \in z_i} \xi_{i,r}(z_i, x_{-i}), \quad (3.32)$$

where (3.31) holds because $\xi_{i,r_j}(x) \leq c_{r_j}(x) = 1 = \xi_{i,r_0}(z_i, x_{-i})$ due to budget balance of the protocol and (3.32) because $\xi_{i,r}(x) = \xi_{i,r}(z, x_{-i})$ for all $r \in x_i, r \neq r_j$ due to separability of the protocol. Clearly, y with $y_i = R \setminus \{r_i\}$ is an optimal profile with $C(y) = 1$. Hence, the fact that x with $C(x) = n$ is a pure Nash equilibrium under any separable protocol proves that the price of anarchy is at least n . \square

For basic protocols and symmetric games without delay, we introduce *strongly decharged* strategy profiles, an extension of the decharged concept. We outline a cost sharing protocol that makes a strongly decharged profile not only a pure Nash equilibrium but even the most expensive pure Nash equilibrium of the game. Using this structural result, an upper bound on the price of anarchy only needs existence guarantees for strongly decharged profiles. For these, we again provide an algorithm that transforms an optimal profile into a strongly decharged profile.

Definition 3.25 (strongly decharged strategy profile)

In a symmetric matroid model (N, R, X, d, c) , a profile $x \in X$ is *strongly decharged*, if it is decharged, that is,

$$c_r(x) \leq \sum_{i \in N_r(x)} \Delta_i^r(x) \text{ for all } r \in R.$$

and additionally the cost of any resources used by multiple players is strictly less than the sum of the cost of alternatives,

$$c_r(x) < \sum_{i \in N_r(x)} \Delta_i^r(x) \text{ for all } r \in R : c_r(x) > 0 \text{ and } |N_r(x)| > 1.$$

Lemma 3.26. *Consider a symmetric matroid model (N, R, X, d, c) and a profile $x \in X$. If x is strongly decharged, there is a basic protocol that assigns cost share functions such that x is the most expensive pure Nash equilibrium of the induced game.*

For the proof of this lemma, we introduce the anarchy eliminating protocol that just like the enforcing protocol relies on a decharged profile x .

Definition 3.27 (anarchy eliminating protocol)

The anarchy eliminating protocol takes as input a strongly decharged profile x . For any profile z , denote a global set of foreign players by $N^1(z) = \{i \in N : z_i \neq x_i\}$. Then, the anarchy eliminating protocol assigns for all $i \in N$ and $r \in R$ the cost share functions

$$\xi_{i,r}(z) = \begin{cases} \frac{\Delta_i^r(x)}{\sum_{j \in N_r(x)} \Delta_j^r(x)} \cdot c_r(x), & \text{if } r \in z_i \text{ and } z = x, \\ c_r(z), & \text{if } z \neq x, N^1(z) \cap N_r(z) \neq \emptyset \text{ and } i = \min N^1(z) \cap N_r(z), \\ c_r(z), & \text{if } z \neq x, N^1(z) \cap N_r(z) = \emptyset \text{ and } i = \min N_r(z). \\ 0, & \text{else.} \end{cases}$$

Proof of Lemma 3.26. We first show that the profile x is a pure Nash equilibrium under the anarchy eliminating protocol. To this end, let ξ^* be the cost share functions assigned by the enforcing protocol and note that the cost shares relevant for a Nash equilibrium are identical under both protocols, that is, $\xi_{i,r}(x) = \xi_{i,r}^*(x)$ and $\xi_{i,r}(z_i, x_{-i}) = \xi_{i,r}^*(z_i, x_{-i})$ for any $i \in N$ and any $z_i \in X_i$. Then,

$$\sum_{r \in x_i} \xi_{i,r}(x) = \sum_{r \in x_i} \xi_{i,r}^*(x) \leq \sum_{r \in z_i} \xi_{i,r}^*(z_i, x_{-i}) \quad (3.33)$$

$$\leq \sum_{r \in z_i} c_r(z_i, x_{-i}) = \sum_{r \in z_i} \xi_{i,r}(z_i, x_{-i}), \quad (3.34)$$

where (3.33) holds because x is a pure Nash equilibrium under the enforcing protocol and (3.34) holds because the enforcing protocol is budget balanced.

To complete the proof, we show that x is the most expensive pure Nash equilibrium, that is, $C(z) \leq C(x)$ for any pure Nash equilibrium z . We denote the foreign player with the smallest index by $i^* = \min N^1(z)$. Then,

$$\xi_j(z) \leq \xi_j(z_{i^*}, z_{-j}) = 0 \quad \text{for all } j \in N^1(z), j \neq i^*, \quad (3.35)$$

because z is a pure Nash equilibrium. In order to estimate $\xi_{i^*}(z)$, we need to examine three cases:

- (1) $N^1(z) = \{i^*\}$ and $N_r(x) = \{i^*\}$ for all $r \in x_{i^*}$,
- (2) $N^1(z) = \{i^*\}$ and $|N_r(x)| > 1$ for some $r \in x_{i^*}$,
- (3) $|N^1(z)| > 1$.

In Case (1), we have $c_r(z) = c_r(x)$ for all $r \in R \setminus (z_{i^*} \cup x_{i^*})$ and $N_r(z) = \emptyset$ for all $r \in x_{i^*} \setminus z_{i^*}$ and, hence, with

$$\begin{aligned} C(z) &= \sum_{r \in R \setminus (z_{i^*} \cup x_{i^*})} c_r(z) + \sum_{r \in x_{i^*} \setminus z_{i^*}} c_r(z) + \sum_{r \in z_{i^*}} c_r(z) \\ &= \sum_{r \in R \setminus (z_{i^*} \cup x_{i^*})} c_r(z) + \sum_{r \in z_{i^*}} \xi_{i^*,r}(z) \end{aligned} \quad (3.36)$$

$$\leq \sum_{r \in R \setminus (z_{i^*} \cup x_{i^*})} c_r(x) + \sum_{r \in x_{i^*}} \xi_{i^*,r}^*(x_{i^*}, z_{-i^*}) \quad (3.37)$$

$$= \sum_{r \in R \setminus (z_{i^*} \cup x_{i^*})} c_r(x) + \sum_{r \in x_{i^*}} c_r(x) = C(x) \quad (3.38)$$

we are done. Equation (3.36) holds because $c_r(z) = 0$ for all $r \in x_{i^*} \setminus z_{i^*}$, inequality (3.37) holds because z is a pure Nash equilibrium and (3.38) because $N_r(x) = \{i^*\}$ for all $r \in x_{i^*}$.

In Case (2), as $z_i = x_i$ for all $i \neq i^*$, we have $x = (x_{i^*}, z_{-i^*})$ and $z = (z_{i^*}, x_{-i^*})$. Hence,

$$\begin{aligned}\xi_{i^*}(z_{i^*}, x_{-i^*}) &= \xi_{i^*}(x) \\ &< \xi_{i^*}(z_{i^*}, x_{-i^*}) = \xi_{i^*}(z),\end{aligned}$$

because x is strongly decharged. This contradicts that z is a pure Nash equilibrium, i.e., Case (2) is a contradiction.

In Case (3), for resources that are used by players from $N^1(z)$, we have

$$\sum_{\substack{r \in R \\ N_r(z) \cap N^1(z) \neq \emptyset}} c_r(z) = \sum_{i \in N^1(z)} \xi_i(z) \quad (3.39)$$

$$= \xi_{i^*}(z) \quad (3.40)$$

$$\leq \xi_{i^*}(x_{i^*}, z_{-i^*}) \quad (3.41)$$

$$= \sum_{\substack{r \in x_{i^*} \\ i^* = \min N_r(x) \\ N_r(z) \cap N^1(x_{i^*}, z_{-i^*}) = \emptyset}} c_r(x_{i^*}, z_{-i^*}) \quad (3.42)$$

$$\leq \sum_{\substack{r \in x_{i^*} \\ i^* = \min N_r(x) \\ N_r(z) \cap N^1(x_{i^*}, z_{-i^*}) = \emptyset}} c_r(x), \quad (3.43)$$

where (3.39) is given by the protocol, (3.40) follows from (3.35) and (3.41) holds because z is a pure Nash equilibrium. Because there are multiple foreign players (Case (3)), only the second case from the definition of the protocol applies for i^* and (3.42) follows. Finally, (3.43) holds because for resources $r \in x_{i^*}$ without foreign players, i.e., $N_r(z) \cap N^1(x_{i^*}, z_{-i^*}) = \emptyset$, we have $\ell_r(x_{i^*}, z_{-i^*}) \leq \ell_r(x)$.

We conclude,

$$C(z) = \sum_{\substack{r \in R \\ N_r(z) \cap N^1(z) = \emptyset \\ c_r(z) > 0}} c_r(z) + \sum_{\substack{r \in R \\ N_r(z) \cap N^1(z) \neq \emptyset}} c_a(z) \quad (3.44)$$

$$\leq \sum_{\substack{r \in R \\ N_r(z) \cap N^1(z) = \emptyset \\ c_r(z) > 0}} c_r(x) + \sum_{\substack{r \in x_{i^*} \\ i^* = \min N_r(x) \\ N_r(z) \cap N^1(x_{i^*}, z_{-i^*}) = \emptyset}} c_r(x) \quad (3.44)$$

$$\leq \sum_{r \in R} c_r(x) = C(x), \quad (3.45)$$

where the first part of (3.44) holds because for resources $r \in R$ with $N_r(z) \cap N^1(z) = \emptyset$ we have $\ell_r(z) \leq \ell_r(x)$ and the second part follows from (3.43). The two summands in (3.44) are disjoint, as for all resources $r \in x_{i^*}$ with $i^* = \min N_r(x)$ and $i^* \notin N_r(z)$ we have $c_r(z) = 0$, because $\xi_j(z) \leq \xi_j(z_{i^*}, z_{-j}) = 0$ for all $j \in N_r(z)$. Hence, (3.45) follows. \square

Algorithm 2 Find strongly decharged profile x

Input: symmetric matroid congestion model (N, R, X, d, c) , profile y
Output: strongly decharged profile x

```

1:  $k \leftarrow 1$       {step number}
2:  $x^1 \leftarrow y$     {starts with optimal profile  $y$ }
3: while there are charged resources do
4:   select the most expensive charged resource  $r^k \leftarrow \operatorname{argmax} \{c_r(x^k) : r \in R \text{ is charged}\}$ 
5:   if there is a player  $i \in N_{r^k}(x^k)$  that can switch to cost-free resource, i.e.,  $\Delta_i^{r^k}(x^k) = 0$  then
        {case called Zero-move}
6:     select player this player  $i^k \leftarrow i$ 
7:   else if some player on  $r^k$  was moved before then    {case called Shuffle}
8:     select player  $i^k \in N_{r^k}(x^k)$  that was moved last
9:   else    {case called Kickoff}
10:    select player who is willing to pay the least
11:     $i^k \leftarrow \operatorname{argmin}_{i \in N_{r^k}(x^k)} \Delta_i^{r^k}(x^k)$ 
12:    select cheapest (1,1)-exchange  $s^k \leftarrow \operatorname{argmin}_{\substack{s \in R \\ x_{i^k}^k + s - r^k \in X_{i^k}}} c_s(x_{i^k}^k + s - r^k, x_{-i^k}^k)$ 
13:    execute (1,1)-exchange  $x^{k+1} \leftarrow (x_{i^k}^k + s^k - r^k, x_{-i^k}^k)$ 
14:    iterate  $k \leftarrow k + 1$ 
15: end while
16: return  $x \leftarrow x^k$ 

```

It follows from the above that a cost guarantee for strongly decharged profiles implies an upper bound on the price of anarchy. We present Algorithm 2, closely related to Algorithm 1, that takes an optimal profile of a symmetric matroid model without delays as input and returns a strongly decharged profile that costs at most \mathcal{H}_n times as much as the input.

To prove that Algorithm 2 works as desired, we follow the same steps as in the proof of Algorithm 1. First, we show that the algorithm terminates, then we deal with the cost of the strategy profile x .

Lemma 3.28. *Algorithm 2 terminates.*

Proof. Similar to the proof of Lemma 3.13, where we showed that Algorithm 1 terminates, we use the interpretation of players scheduling jobs on the resources, fix a player i and follow one of her jobs over the course of the algorithm, showing that it can only be moved finitely often. The difference to the proof of Lemma 3.13 is that we need a different reasoning to bound the number of Shuffles and that are (possibly many) Zero-moves.

There can be at most one Kickoff involving the tracked job, afterwards the job is only moved by Shuffles and Zero-moves. Say the algorithm is in iteration k and a Shuffle involving our job is performed. Then,

$$\Delta_j^{r^k}(x^k) > 0 \quad \text{for all } j \in N_{r^k}(x^k), \tag{3.46}$$

because otherwise the algorithm would perform a Zero-move. We now consider two cases: For $|N_{r^k}(x^k)| = 1$, we obtain

$$c_{r^k}(x^k) > \Delta_i^{r^k}(x^k) \quad (3.47)$$

$$= c_{s^k}(x_i^k + s - r^k, x_{-i}^k) = c_{s^k}(x^{k+1}), \quad (3.48)$$

where (3.47) follows because r^k is charged in x^k . Equality (3.48) follows since Algorithm 2 moves i to the cheapest available resource.

If $|N_{r^k}(x^k)| > 1$, then we obtain

$$c_{r^k}(x^k) \geq \sum_{j \in N_r(x^k)} \Delta_j^{r^k}(x^k) \quad (3.49)$$

$$> \Delta_i^{r^k}(x^k) \\ = c_{s^k}(x_i^k + s^k - r^k, x_{-i}^k) = c_{s^k}(x^{k+1}). \quad (3.50)$$

where (3.49) is valid because r^k is charged in x^k . The second inequality (3.50) holds because of (3.46), that is, even if the sum of the cheapest alternatives is equal to the cost of r^k , the cheapest alternative for our tracked job is strictly cheaper than the cost of r^k as the cheapest alternatives of the other jobs in r^k have strictly positive cost. The equalities follow as in (3.48). In both cases, a Shuffle moves the player to a strictly cheaper resource and the last-in/first-out mechanism of the algorithm makes sure that in between two Shuffles the cost of the resource does not increase more than it decreases. Hence, the number of Shuffles is limited.

The last type of move to be considered is Zero-moves. After a Zero-move, a player could only be considered for a Shuffle when the resource cost is zero (last-in-first-out), however, then the resource is decharged by definition. Hence, once a job was moved by a Zero-move, only further Zero-moves are possible. A job can indeed be moved multiple times by Zero-moves if the cost of its resource becomes positive, i.e., when other jobs are added to its resource through Kickoff or Shuffle moves. However, the number of Kickoffs and Shuffles is finite, hence the number of Zero-moves is finite as well and the algorithm terminates after a finite number of iterations. \square

Lemma 3.29. *The strategy profile x returned by Algorithm 2 has at most \mathcal{H}_n times the cost of the input profile y .*

Proof. We follow the reasoning of the proof of Lemma 3.14, adapting where Algorithm 2 differs from Algorithm 1. In Algorithm 2, a job's first move is always a Kickoff or a Zero-move. For a Kickoff of a job q in iteration k , we have

$$c_{y(q)}(y) \geq p(q)c_{s^k}(x^{k+1})$$

as in (3.8). For a Zero-move, $c_{y(q)}(y) \geq 0 = c_{s^k}(x^{k+1})$. Since Shuffles take job q to strictly cheaper resources as shown in the previous lemma, we have for any iteration l in which q is

moved, $c_s(x^{l+1}) \leq \frac{1}{p(q)} \cdot c_{y(q)}(y)$. Altogether, in the final profile x , the cost of a resource $r \in R$ to which jobs have been moved is determined by the last job that was moved there. Thus,

$$C(x) \leq \mathcal{H}_n \cdot C(y)$$

as in (3.11). \square

Corollary 3.30. *The price of anarchy of symmetric matroid games induced by basic cost sharing protocols is \mathcal{H}_n ,*

$$\min_{\Xi \in \mathcal{B}_n} POA(\Xi) = \mathcal{H}_n.$$

3.6 Concave and Convex Costs

We now restrict the set of cost functions to be either concave or convex. Concave functions are frequently used to model economies of scale, that is, situations in which marginal costs are decreasing. Examples include network design games, where cost functions are modeled by fixed or concave costs, cf. Anshelevich et al. (2008), Bilò et al. (2010), Chekuri et al. (2007), Chen and Roughgarden (2009). On the other hand, convex costs are used to model the sharp increase of costs if resources are scarce or if there are hard capacities on the amount of resources available. For instance in telecommunication networks, relevant cost functions are the so-called $M/M/1$ -delay functions (see Bertsekas and Gallagher 1992, Orda et al. 1993). These are convex functions of the form $c_a(x) = 1/(u_a - x)$, where u_a represents the physical capacity of arc a .

3.6.1 Concave Cost Functions

For concave cost functions, we show that any optimal strategy profile of a facility location model is decharged, and, thus, the price of stability is one in this case.

Theorem 3.31. *The price of stability of facility location games with concave costs induced by basic or separable protocols is one,*

$$\min_{\Xi \in \mathcal{B}_n} POS(\Xi) = 1, \quad \min_{\Xi \in \mathcal{S}_n} POS(\Xi) = 1.$$

Proof. We prove that any socially optimal strategy profile $x \in X$ satisfies (D1) and (D2) and, thus, by Theorem 3.11 the enforcing protocol induces x as a pure Nash equilibrium.

Assume for a contradiction that x does not satisfy (D1) for some player i and resource r . Then we can simply reassign player i to her cheapest alternative and strictly reduce the total cost, which is a contradiction to our assumption that x is socially optimal.

Assume that x does not satisfy (D2). Hence, there is a resource $r \in R$ with

$$c_r(x) > \sum_{i \in N_r(x)} \Delta_i^r(x) - t_{i,r}. \quad (3.51)$$

We construct a strategy profile z with $C(z) < C(x)$, showing a contradiction. To this end, for each player $i \in N_r(x)$, let

$$s_i = \operatorname{argmin}_{\substack{s \in R \\ x_i + s - r \in X_i}} (c_s(x_i + s - r, x_{-i}) + t_{i,s})$$

be the cheapest alternative to r in x . Let $z_i = x_i + s_i - r$ be the strategy where i switches from r to s_i , such that $\Delta_i^r(x) = c_{s_i}(z_i, x_{-i}) + t_{i,s_i}$. For $i \notin N_r(x)$, let $z_i = x_i$.

Inequality (3.51) implies that $r \notin z_i$ for all $i \in N_r(x)$ and hence that $N_r(z) = \emptyset$, i.e., $c_r(z) = 0$. Resources $s \neq r$ either have the same users in z as in x , i.e., $c_s(z) = c_s(x)$, or they have users $N_s(z) \setminus N_s(x)$ that have switched from r to s . In the latter case, as the cost functions are concave,

$$c_s(z) \leq \sum_{i \in N_s(z) \setminus N_s(x)} c_s(z_i, x_{-i}) = \sum_{i \in N_s(z) \setminus N_s(x)} (\Delta_i^r(x) - t_{i,s}). \quad (3.52)$$

Combining these observations, we conclude

$$\begin{aligned} C(z) &= \sum_{\substack{s \in R \\ N_s(z) = N_s(x)}} \left(c_s(z) + \sum_{i \in N_s(z)} t_{i,s} \right) + \sum_{\substack{s \in R \\ N_s(z) \supset N_s(x)}} \left(c_s(z) + \sum_{i \in N_s(z)} t_{i,s} \right) \\ &\stackrel{(3.52)}{\leq} \sum_{\substack{s \in R \\ N_s(z) = N_s(x)}} \left(c_s(x) + \sum_{i \in N_s(x)} t_{i,s} \right) + \sum_{i \in N_r(x)} \Delta_i^r(x) + \sum_{\substack{s \in R \\ N_s(z) \supset N_s(x)}} \sum_{i \in N_s(x)} t_{i,s} \\ &\stackrel{(3.51)}{<} \sum_{\substack{s \in R \\ N_s(z) \supseteq N_s(x)}} \left(c_s(x) + \sum_{i \in N_s(x)} t_{i,s} \right) + c_r(x) + \sum_{i \in N_r(x)} t_{i,r} = C(x). \end{aligned}$$

This implies that x is not an optimal strategy profile, a contradiction. \square

3.6.2 Convex Cost Functions

We now consider cost functions that are non-negative, non-decreasing and whose per-unit costs $\frac{c(\ell(z))}{\ell(z)}$ are non-decreasing with respect to the load $\ell(z)$. Such functions are quite rich and contain non-negative, non-decreasing and convex functions.

We introduce the *opt-enforcing protocol*, for which we prove a price of anarchy of 1. The intuition behind this protocol is similar to the enforcing protocols presented before: in all undesired outcomes some player is assigned high cost that allows a profitable deviation.

Definition 3.32 (opt-enforcing protocol)

For a given model (N, R, X, d, c, t) , the opt-enforcing protocol takes as input an optimal strategy profile y . We again denote for any profile z and resource r the set of foreign players on r by $N_r^1(z) = N_r(z) \setminus N_r(y)$. Then, the opt-enforcing protocol assigns the cost shares

$$\xi_{i,r}(z) = \begin{cases} d_i \cdot \frac{c_r(z)}{\ell_r(z)}, & \text{if } r \in z_i \text{ and } N_r^1(z) = \emptyset, \\ c_r(z), & \text{if } r \in z_i, N_r^1(z) \neq \emptyset \text{ and } i = \min N_r^1(z), \\ 0, & \text{else.} \end{cases} \quad (3.53)$$

$$\xi_{i,r}(z) = \begin{cases} d_i \cdot \frac{c_r(z)}{\ell_r(z)}, & \text{if } r \in z_i \text{ and } N_r^1(z) = \emptyset, \\ c_r(z), & \text{if } r \in z_i, N_r^1(z) \neq \emptyset \text{ and } i = \min N_r^1(z), \\ 0, & \text{else.} \end{cases} \quad (3.54)$$

$$\xi_{i,r}(z) = \begin{cases} d_i \cdot \frac{c_r(z)}{\ell_r(z)}, & \text{if } r \in z_i \text{ and } N_r^1(z) = \emptyset, \\ c_r(z), & \text{if } r \in z_i, N_r^1(z) \neq \emptyset \text{ and } i = \min N_r^1(z), \\ 0, & \text{else.} \end{cases} \quad (3.55)$$

Under the opt-enforcing protocol, the players share the cost proportional to their demands on all resources without foreign players. On resources with foreign players, the foreign player with the smallest index pays the entire cost of the resource.

Theorem 3.33. *The opt-enforcing protocol is separable and the price of anarchy of facility location games with convex costs induced by the opt-enforcing protocol is one. Thus, for games with convex costs,*

$$\min_{\Xi \in \mathcal{B}_n} POA(\Xi) = 1, \quad \min_{\Xi \in \mathcal{S}_n} POA(\Xi) = 1.$$

Proof. Budget balance and separability are clear from the definition of the protocol. For stability it can easily be verified that for an instance (N, R, X, d, c, t) , the optimal outcome y is a Nash equilibrium. We only prove the bound on the price of anarchy, showing that all Nash equilibria have the same cost as a socially optimal profile y . More precisely, we show for every player i and Nash equilibrium x that

$$\sum_{r \in x_i} (\xi_{i,r}(x) + t_{i,r}) \leq \sum_{r \in y_i} (\xi_{i,r}(y) + t_{i,r}). \quad (3.56)$$

By definition of a Nash equilibrium,

$$\sum_{r \in x_i} (\xi_{i,r}(x) + t_{i,r}) \leq \sum_{r \in y_i} (\xi_{i,r}(y_i, x_{-i}) + t_{i,r}). \quad (3.57)$$

Two cases are to be considered for every resource $r \in y_i$: either it hosts a nonempty set $N_r^1(x)$ of foreign players or $N_r^1(x) = \emptyset$. If there are foreign players on r , then one of them pays for the entire cost there and hence (3.55) gives $\xi_{i,r}(y_i, x_{-i}) = 0$ (note that $i \notin N_r^1(y_i, x_{-i})$ by Definition 3.32). If there are no foreign players on r , then $\ell_r(y_i, x_{-i}) \leq \ell_r(y)$ yields

$$\frac{c_r(y_i, x_{-i})}{\ell_r(y_i, x_{-i})} \leq \frac{c_r(y)}{\ell_r(y)},$$

because the cost per unit is non-decreasing. Plugging this into (3.53), we have $\xi_{i,r}(y_i, x_{-i}) \leq \xi_{i,r}(y)$. With $\xi_{i,r}(y_i, x_{-i}) \leq \xi_{i,r}(y)$ in both cases for all resources $r \in y_i$,

$$\sum_{r \in y_i} (\xi_{i,r}(y_i, x_{-i}) + t_{i,r}) \leq \sum_{r \in y_i} (\xi_{i,r}(y) + t_{i,r}),$$

which combined with (3.57) yields (3.56). \square

3.7 Computational Complexity of Cost Shares

The protocols introduced in the previous sections involve decharged profiles which in turn rely on the computation of optimal strategy profiles, which are generally *NP*-hard to compute. A natural question is whether there are different approaches that allow to compute optimal cost shares or cost shares with good approximation guarantees in polynomial time. We answer this question negatively by showing that the problem of computing optimal cost shares is strongly NP-hard

and not approximable by any constant factor even for instances with unweighted players, zero delays and singleton strategies. In light of this hardness we restrict our model to unweighted players, zero delays and singleton strategies, and switch our objective: previously we looked for cost shares such that the best/worst Nash equilibrium is a good approximation of the social optimum, now we ask for *polynomial time computable* cost shares that approximate the optimal cost shares in terms of cost of the best/worst Nash equilibrium.

3.7.1 Hardness and Inapproximability

Let I be an instance of a matroid facility location model and $\Xi(I)$ the set of possible cost shares for I . We investigate the computational complexity of the following two optimization problems

$$\begin{aligned} \min_{\xi \in \Xi(I)} \min_{x \in PNE(I, \xi)} \frac{C(x)}{C(y)} & \quad (\text{BEST NASH}) \\ \min_{\xi \in \Xi(I)} \max_{x \in PNE(I, \xi)} \frac{C(x)}{C(y)}, & \quad (\text{WORST NASH}) \end{aligned}$$

where $PNE(I, \xi)$ denotes the set of pure Nash equilibria of the game induced by I and ξ , and y is some optimal strategy profile.

Theorem 3.34. *The problem **BEST NASH** is strongly NP-complete and there are no clogn approximation algorithms for any $c < 1$, unless $P = NP$. This holds even for instances with unweighted players, zero delays, singleton strategies and unit fixed costs.*

Proof. By Theorem 3.11, problem **BEST NASH** can be reformulated as

$$\min_{x \in X} C(x) \text{ s.t. } x \text{ is decharged.}$$

Now, given any $x \in X$, we can check in polynomial time whether x is feasible (i.e., it fulfills (D1) and (D2)) and whether $C(x) = k$ for some given k . Thus, **BEST NASH** is in NP.

We prove inapproximability by providing an approximation preserving reduction from HITTING SET. An instance of HITTING SET consists of a collection \mathcal{C} of subsets of a finite set of elements E . A hitting set for \mathcal{C} is a subset $S \subseteq E$ such that S contains at least one element from each subset in \mathcal{C} . The goal is to minimize the cardinality of the hitting set.

Given an instance of HITTING SET, we create an instance of **BEST NASH** as follows: we identify E with R and for each subset in \mathcal{C} , the instance has a player with this subset as strategy set, i.e., $|N| = |\mathcal{C}|$ and $X = \mathcal{C}$. Each resource of the instance has fixed unit cost when used.

Claim 3.35. *For $k \in \mathbb{N}$, there is a hitting set of cardinality less or equal to k if and only if there is a decharged profile with cost less or equal k .*

Proof. In the constructed cost sharing instance, every strategy profile is decharged: with no delays (D1) is always fulfilled, and with fixed unit resource costs (D2), too. Given a strategy profile, the used resources form a hitting set and the cost of the profile is equal to the cardinality of the set. Given a hitting set, a strategy profile can be constructed by assigning each player i to one of the resources in X_i that is in the hitting set. \square

The proof follows now from the above claim and the fact that HITTING SET is equivalent to the SET COVER problem (Ausiello et al. 1980) which is known not to be approximable within $c \log n$ for any $c < 1$ unless $P = NP$ (Raz and Safra 1997). \square

We now prove an inapproximability result for computing cost shares minimizing the cost of the worst Nash equilibrium

Theorem 3.36. *Problem WORST NASH is strongly NP-hard and there are no $c \log n$ approximation algorithms for any $c < 1$ unless $P = NP$. This holds even for instances with unweighted players, zero delays, singleton strategies and unit fixed costs.*

Proof. We again reduce from HITTING SET. Given an instance of HITTING SET, we construct a cost sharing game like above and add two more players denoted player a and player b . We set $X_a = X_b = R$.

Claim 3.37. *There is a hitting set of cardinality less or equal to k if and only if*

$$\min_{\xi \in \Xi(I)} \max_{x \in PNE(I, \xi)} C(x) \leq k.$$

Proof. For the direction \Leftarrow , observe that if there are cost shares ξ such that the most expensive Nash equilibrium x has cost $C(x) \leq k$, then the resources used in x form a hitting set of cardinality less than or equal to k . For the \Rightarrow direction, given a hitting set of size at most k , we construct a strategy profile x by first assigning the players except a and b to resources that correspond to elements of the hitting set and then assigning players a and b both to a resource used by other players. We assign the following cost shares for all $r \in R$ and $z \in X$:

$$\begin{aligned} \xi_{a,r}(z) &= \begin{cases} 1, & \text{if } N_r(z) = \{a\} \text{ or } N_r(z) = \{a, b\}, \\ \frac{1}{2}, & \text{if } a \in N_r(z) \text{ and } N_r(z) \subseteq N_r(x) \cup \{a, b\} \text{ and } N_r(z) \not\subseteq \{a, b\}, \\ 0, & \text{if } a \in N_r(z) \text{ and } N_r(z) \not\subseteq N_r(x) \cup \{a, b\} \text{ and } b \notin N_r(z), \\ 1, & \text{if } a \in N_r(z) \text{ and } N_r(z) \not\subseteq N_r(x) \cup \{a, b\} \text{ and } b \in N_r(z), \\ 0, & \text{if } a \notin N_r(z), \end{cases} \\ \xi_{b,r}(z) &= \begin{cases} 1, & \text{if } b \in N_r(z) \text{ and } a \notin N_r(z), \\ 0, & \text{else,} \end{cases} \end{aligned}$$

and for all $i \neq a, b$

$$\xi_{i,r}(z) = \begin{cases} 1 - \xi_{a,r}(z) - \xi_{b,r}(z), & \text{if } i = \min N_r(z) \setminus N_r(x), \\ 1 - \xi_{a,r}(z) - \xi_{b,r}(z), & \text{if } N_r(z) \setminus N_r(x) = \emptyset \text{ and } i = \min N_r(z), \\ 0, & \text{else.} \end{cases}$$

The profile x is a pure Nash equilibrium under these cost shares, as player a pays $\frac{1}{2}$ and would pay either $\frac{1}{2}$ or 1 on any other resource, player b pays 0 and all other players pay 1 or less and would pay 1 if they switched.

For any profile $z \neq x$ with $C(z) > C(x)$, there is some resource r with $N_r(x) = \emptyset$ and $N_r(z) \neq \emptyset$. We show that such a profile cannot be an equilibrium. For an equilibrium, the definition of the

cost shares for player b implies that $z_b = z_a$, as otherwise player b could immediately reduce her cost by switching to z_a . We, hence, assume from now on $z_b = z_a$. If $N_r(z) = \{a, b\}$, then player a has cost 1 and, if $N_{x_a}(z) \neq \emptyset$ could reduce her cost to $\frac{1}{2}$ by switching to x_a or, in case $N_{x_a}(z) = \emptyset$, could reduce her cost to 0 by switching to some other resource s with $N_s(z) \not\subseteq N_s(x)$. If $N_r(z) \neq \{a, b\}$, then there are other players than a and b using r as $N_r(z) \neq \emptyset$ and $z_b = z_a$. Hence, either $\{a, b\} \cap N_r(z) = \emptyset$ and player a could reduce her cost by switching to r , or $\{a, b\} \subset N_r(z)$ and player a could reduce her cost by switching somewhere else as above – unless $N_s(z) = \emptyset$ for all $s \neq r$, in which case $C(z) = 1$, a contradiction to $C(z) > C(x)$. Consequently, x with $C(x) \leq k$ is the most expensive pure Nash equilibrium of the game. \square

 \square

3.7.2 Approximation Algorithms

We have shown that it is computationally hard to find cost shares that approximate the cost of the optimal cost shares within a logarithmic factor in n . This hardness result even holds for instances with unweighted players, zero delays, singleton strategies and unit fixed costs. In light of this hardness even for this restricted class of problems, we study approximation algorithms for the case of unweighted players, zero delays and singleton strategies, still assuming general non-decreasing costs. Such instances are given as $I = (N, R, X, c)$, with $d_i = 1$ and $t_{i,r} = 0$ for all $i \in N$ and $r \in R$. As the players' strategies in this setting consist of a single resource, we denote the cost shares in a strategy profile $x \in X$ simply by $\xi_i(x) = \xi_{i,r}(x)$ with $r = x_i$.

We present in the following an approximation algorithm that computes in polynomial time a decharged profile whose cost is bounded from above by \mathcal{H}_n times the cost of an optimal profile – matching the performance bound presented in Section 3.4.

For a given instance $I = (N, R, X, c)$, the algorithm starts with an empty strategy profile x and updates x as it iteratively assigns the n players to the resources in R . While the algorithm runs, let $\bar{n}(x)$ be the number of players not yet assigned to a resource and let $G(x)$ be a directed graph representing the current allocation x . The graph has vertices for all players, all resources and additional source and sink vertices s and t . $G(x)$ is bipartite with arcs according to the following rules:

- arc (i, r) from player i to resource r if $r \in X_i$
- arc (r, i) from resource r to player i if $r = x_i$, i.e., i is assigned to r in x
- arc (s, i) if player i is not assigned to any resource in x .

All above arcs have capacity 1.

Using the cost shares assigned by Algorithm 3, the cost of a resource r is shared proportionally among its users, unless there are players $N_r(z) \setminus N_r(x)$, in which case only these players share the cost.

Theorem 3.38. *Algorithm 3 computes in polynomial time cost shares that guarantee a price of stability of at most \mathcal{H}_n and a price of anarchy of at most n .*

Proof. We prove the theorem in four lemmas.

Algorithm 3 Compute semi-proportional cost shares

Input: instance $I = (N, R, X, c)$ **Output:** cost shares ξ

- 1: start with empty profile x , no players assigned to resources
- 2: **while** not all players N assigned to resources R **do**
- 3: compute $G(x)$
- 4: **for all** combinations $r \in R$ and integers $1 \leq v \leq \bar{n}(x)$ **do**
- 5: compute, if existent, an integer s, r flow f with flow value $v(f) = v$.
- 6: **end for**
- 7: from all combinations of r and $v(f)$ for which a flow was found, choose the one with the lowest resulting cost-per-unit $\frac{1}{v(f)}(c_r(|N_r(x)| + v(f)) - c_r(x))$. For each (i, r) arc used in this flow, update x such that player i is assigned to resource r .
- 8: **end while**
- 9: assign cost shares for all $i \in N$ and $z \in X$

$$\xi_i(z) = \begin{cases} \frac{c_{z_i}(z)}{|N_{z_i}(z)|} & , \text{if } N_{z_i}(z) \subseteq N_{z_i}(x) \\ \frac{c_{z_i}(z)}{|N_{z_i}(z) \setminus N_{z_i}(x)|} & , \text{if } N_{z_i}(z) \not\subseteq N_{z_i}(x) \text{ and } i \in N_{z_i}(z) \setminus N_{z_i}(x) \\ 0 & , \text{otherwise.} \end{cases}$$

- 10: **return** ξ
-

Lemma 3.39. *The profile x used in the definition of the cost shares is a pure Nash equilibrium.*

Proof. We use a counter k to enumerate the iterations of the algorithm's main while loop. Denote by x^k the profile at the beginning of iteration k and by f^k the selected flow. We observe that on each resource r , the sequence of the per-unit costs of the load increments is non-decreasing, that is, if in step k the load on r is increased and the next load increase on r is in step l , then

$$\frac{c_r(x^{k+1}) - c_r(x^k)}{|N_r(x^{k+1})| - |N_r(x^k)|} \leq \frac{c_r(x^{l+1}) - c_r(x^l)}{|N_r(x^{l+1})| - |N_r(x^l)|}$$

because otherwise Algorithm 3 had increased the load on r in iteration k directly to $N_r(x^{l+1})$ instead of $N_r(x^{k+1})$. Thus, if on resource r the last load increment was in step k^* , then the per unit cost of r in x , which is the average over all such increments, is no greater than the per unit cost of the increment in k^* . This in turn is no greater than the cost of adding one player to any other resource \bar{r} in x^{k^*} because otherwise the algorithm had increased the load on \bar{r} instead of increasing it on r . Hence, for all $i \in N_r(x)$ and $\bar{r} \in X_i$,

$$\begin{aligned} \xi_i(x) &= \frac{c_r(x)}{|N_r(x)|} \leq \frac{c_r(x) - c_r(x^{k^*})}{|N_r(x)| - |N_r(x^{k^*})|} \leq \frac{c_{\bar{r}}(|N_{\bar{r}}(x^{k^*})| + 1) - c_{\bar{r}}(|N_{\bar{r}}(x^{k^*})|)}{1} \\ &\leq c_{\bar{r}}(|N_{\bar{r}}(x^{k^*})| + 1) \leq c_{\bar{r}}(|N_{\bar{r}}(x)| + 1) = \xi_i(\bar{r}, x_{-i}). \end{aligned}$$

□

Lemma 3.40. *Let y be some optimal strategy profile. Then, the profile x used to define the cost shares has cost $C(x) \leq \mathcal{H}_n \cdot C(y)$.*

Proof. We first check that updating x^k corresponding to the flow f^k as done in line 7 works as desired. Clearly, updating x^k this way is feasible with regard to the strategy space X . Moreover, when we add $v(f^k)$ players going from x^k to x^{k+1} , the load only changes on one resource r and the cost difference of x^k and x^{k+1} is $C(x^{k+1}) - C(x^k) = c_r(x^{k+1}) - c_r(x^k)$.

To bound the per-unit cost increase in a given iteration k , note that for each resource \bar{r} with less users than in the optimal profile y , i.e., with $|N_{\bar{r}}(x^k)| < |N_{\bar{r}}(y)|$, there is an integer s, \bar{r} flow with flow value $|N_{\bar{r}}(y)| - |N_{\bar{r}}(x^k)|$. As the algorithm does not choose this flow, the corresponding per-unit cost is no less than that on the resource r ,

$$\frac{c_r(x^{k+1}) - c_r(x^k)}{v(f^k)} \leq \frac{c_{\bar{r}}(y) - c_{\bar{r}}(x^k)}{|N_{\bar{r}}(x^k)| - |N_{\bar{r}}(y)|}. \quad (3.58)$$

Noting that the number of players missing on such resources is in total at least $\bar{n}(x^k)$, i.e.,

$$\bar{n}(x^k) \leq \sum_{\substack{\bar{r} \in R \\ |N_{\bar{r}}(x^k)| < |N_{\bar{r}}(y)|}} (|N_{\bar{r}}(x^k)| - |N_{\bar{r}}(y)|),$$

we sum up (3.58) over all such resources,

$$\begin{aligned} \frac{c_r(x^{k+1}) - c_r(x^k)}{v(f^k)} &\leq \sum_{\substack{\bar{r} \in R \\ |N_{\bar{r}}(x^k)| < |N_{\bar{r}}(y)|}} \frac{|N_{\bar{r}}(x^k)| - |N_{\bar{r}}(y)|}{\bar{n}(x^k)} \frac{c_{\bar{r}}(y) - c_{\bar{r}}(x^k)}{|N_{\bar{r}}(x^k)| - |N_{\bar{r}}(y)|} \\ &= \sum_{\substack{\bar{r} \in R \\ |N_{\bar{r}}(x^k)| < |N_{\bar{r}}(y)|}} \frac{c_{\bar{r}}(y) - c_{\bar{r}}(x^k)}{\bar{n}(x^k)} \leq \frac{C(y)}{\bar{n}(x^k)}. \end{aligned} \quad (3.59)$$

This bound on the per-unit cost of the load increases is equivalent to $c_r(x^{k+1}) - c_r(x^k) \leq \frac{v(f^k)}{\bar{n}(x^k)} \cdot C(y)$, and, thus, allows to estimate the cost of x as follows.

$$\begin{aligned} C(x) &= \sum_k \left(C(x^{k+1}) - C(x^k) \right) \\ &\leq \sum_k \frac{v(f^k)}{\bar{n}(x^k)} \cdot C(y) = \sum_k \frac{v(f^k)}{\sum_{j \geq k} v(f^j)} \cdot C(y) = \sum_k \sum_{i=1}^{v(f^k)} \frac{1}{\sum_{j \geq k} v(f^j)} \cdot C(y) \\ &\leq \sum_k \sum_{i=1}^{v(f^k)} \frac{1}{\sum_{j \geq k} v(f^j) - i + 1} \cdot C(y) = \mathcal{H}_n \cdot C(y), \end{aligned} \quad (3.60)$$

where we use in (3.60) that $\bar{n}(x^k)$ is equal to the total of players added in iteration k and later, i.e., $\sum_{j \geq k} v(f^j) = \bar{n}(x^k)$. \square

Lemma 3.41. *Let y be some optimal strategy profile and $z \in PNE(I, \xi)$ a pure Nash equilibrium. Then, $C(z) \leq n \cdot C(y)$.*

Proof. First observe that from (3.59) we can derive $c_r(x^{l+1}) - c_r(x^l) \leq C(y)$ for every iteration l where the load on a resource r is increased. Particularly, $c_r(x) \leq \sum_{i \in N_r(x)} C(y)$. We estimate the cost of z by estimating the cost shares of groups of players in z .

For players that are in x on resources r with $N_r(x) = N_r(z)$, we have by the above observation

$$\sum_{i \in N_r(x)} \xi_i(z) = c_r(x) \leq \sum_{i \in N_r(x)} C(y).$$

For players that are in x on resources r with $N_r(z) \subset N_r(x)$, we choose the smallest k such that $|N_r(z)| + 1 \leq |N_r(x^k)|$ and let q be the number of times the algorithm has increased the load on r up to iteration k . Then, $q \leq |N_r(z)| + 1$ and by our above observation the q load increases will increase the cost by up to $q \cdot C(y)$. Then, for all $i \in N_r(x)$,

$$\xi_i(z) \leq \xi_i(r, z_{-i}) = \frac{c_r(|N_r(z)| + 1)}{|N_r(z)| + 1} \quad (3.61)$$

$$\leq \frac{c_r(x^k)}{q} \leq \frac{q \cdot C(y)}{q} = C(y). \quad (3.62)$$

Here, inequality (3.61) holds because z is a Nash equilibrium, and inequality (3.62) follows from our previous observations $|N_r(z)| + 1 \leq |N_r(x^k)|$ and $q \leq |N_r(z)| + 1$. For the last inequality observe that $c_r(x^k)$ is the sum of the costs of the q load increases up to iteration k .

For players that are in x on resources r with $N_r(z) \not\subseteq N_r(x)$, we have $\xi_i(z) \leq \xi_i(x_i, z_{-i}) = 0$ for all players $i \in N_r(x)$ because z is a Nash equilibrium. Summing up across all player gives the desired

$$C(z) = \sum_{i \in N} \xi_i(z) \leq \sum_{i \in N} C(y) = n \cdot C(y).$$

□

Lemma 3.42. *The runtime of Algorithm 3 is polynomial in the size of the input instance.*

Proof. For an instance I with n players and m resources, the algorithm's main while loop (lines 2 to 8) can run at most n times. Computing $G(x)$ can be done in $\mathcal{O}(mn)$ time, the $\mathcal{O}(mn)$ flows of an iteration can each be computed in $\mathcal{O}(m^3n^3)$ with the Edmonds-Karp algorithm (Edmonds and Karp 1972), and updating x can be done in $\mathcal{O}(n)$ time. Hence, the algorithm's runtime is bounded by $\mathcal{O}(m^4n^5)$. □

□

3.8 Non-Matroid Strategy Spaces

The matroid strategy spaces used in this chapter capture many combinatorial settings, including singleton models, models where players choose an arbitrary subset of fixed size from the ground set of resources, and spanning trees in graphs. However, our framework of decharged strategy profiles may be generalized even further. We present in this section a definition of decharged profiles for models where strategy spaces are without combinatorial structure, that is, the strategy set X_i of a player i consists of arbitrary subsets of R .

Given such a *non-matroid* model (N, R, X, d, c, t) and a strategy profile $z \in X$, denote for any set of resources $R' \subseteq R$ all possible partitions into disjoint sets by $\mathcal{P}(R')$. Then, for a partition $P \in \mathcal{P}(R')$, the direct union of all sets $p \in P$ is $\dot{\cup}_{p \in P} p = R'$. We extend our definition of the cheapest alternative to such sets of resources as follows: for a subset $p \subseteq z_i$ of the strategy of player i , denote

$$\Delta_i^p(z) = \min_{\substack{x_i \in X_i \\ x_i \cap p = \emptyset}} \sum_{s \in x_i \setminus z_i} (c_s(x_i, z_{-i}) + t_{i,s}).$$

Definition 3.43 (decharged strategy profile)

A strategy profile z of a non-matroid model (N, R, X, d, c, t) is *decharged* if it fulfills the following two properties:

- (1) Each player i is willing to pay the delays of any subset $p \subseteq z_i$ of her strategy:

$$\sum_{r \in p} t_{i,r} \leq \Delta_i^p(z). \quad (\text{DG1})$$

- (2) For each subset of resources $R' \subseteq R$, its users, after having paid their delays, are willing to share the cost of the resources,

$$\sum_{r \in R'} c_r(z) \leq \min_{i \in N_{R'}(z)} \sum_{P \in \mathcal{P}(z_i \cap R')} \sum_{p \in P} \left(\Delta_i^p(z) - \sum_{r \in p} t_{i,r} \right), \quad (\text{DG2})$$

where the minimization is for each player over all partitions of $z_i \cap R'$.

In Definition 3.6, the $(1, 1)$ -exchange property of matroid bases made it sufficient to compare each single resource of a strategy with its cheapest alternative. In this non-matroid setting, however, an entire subset of a player's strategy can possibly cheaply be replaced by a single resource. Hence, for (DG1), we need to consider all subsets of each player's strategy. In (DG2), we need to minimize over all partitions, as possibly multiple subsets have the same cheapest alternative and, hence, cannot be exchanged at the same time. To make sure that the cost of each single resource is collected by possible cost shares, we need to require that (DG2) holds for all subsets $R' \subseteq R$.

We give an extension of our characterization from Theorem 3.11.

Lemma 3.44. *Any pure Nash equilibrium x in a non-matroid facility location game with a separable cost sharing protocol is decharged.*

Proof. The proof is analogous to the proof of Lemma 3.7, using the Nash inequality in conjunction with the fact that, using a separable protocol, the cost shares of a deviating player do not change for resources that were not exchanged. \square

Definition 3.45 (x -enforcing cost shares for non-matroid models)

Given a non-matroid facility location model and and a decharged strategy profile x , cost shares ξ are x -enforcing if for every player $i \in N$ and every $p \subseteq x_i$

$$\sum_{r \in p} \xi_{i,r}(x) \leq \Delta_i^p(x) - \sum_{r \in p} t_{i,r},$$

and for any $z_i \in X_i$ and $r \in z_i \setminus x_i$

$$\xi_{i,r}(z_i, x_{-i}) = c_r(z_i, x_{-i}).$$

For a given decharged profile such cost shares can be found with a linear program.

Lemma 3.46. *Given a decharged profile x and x -enforcing separable cost shares ξ , x is a pure Nash equilibrium.*

Proof. For a given player $i \in N$ and alternative strategy $z_i \in X_i$, denote by $p = x_i \setminus z_i$ the resources that were exchanged. Then,

$$\sum_{r \in p} (\xi_{i,r}(x) + t_{i,r}) \leq \Delta_i^p(x) \quad (3.63)$$

$$\leq \sum_{r \in z_i \setminus x_i} (c_r(z_i, x_{-i}) + t_{i,r}) \quad (3.64)$$

$$= \sum_{r \in z_i \setminus x_i} (\xi_{i,r}(z_i, x_{-i}) + t_{i,r}), \quad (3.65)$$

where (3.63) and (3.65) hold because the cost shares are x -enforcing, and (3.64) holds by definition of $\Delta_i^p(x)$. Additionally, because the cost shares are separable, $\xi_{i,r}(x) = \xi_{i,r}(z_i, x_{-i})$ for all $r \in x_i \cap z_i$. Combined,

$$\sum_{r \in x_i} (\xi_{i,r}(x) + t_{i,r}) \leq \sum_{r \in z_i} (\xi_{i,r}(z_i, x_{-i}) + t_{i,r}),$$

that is, x is a pure Nash equilibrium. \square

Theorem 3.47 (Characterization of Pure Nash Equilibria for Separable Protocols). *Given a strategy profile $x \in X$ of a non-matroid facility location model, there is a separable protocol such that x is a pure Nash equilibrium in the induced game if and only if x is decharged.*

Proof. Follows from Lemma 3.44 and Lemma 3.46. \square

Convex Cost. For non-matroid models with convex cost functions, note that the opt-enforcing protocol from Section 3.6 does not make use of the matroid structure of the strategy spaces. Hence, it also achieves a price of stability and price of anarchy of 1 in the non-matroid setting.

Concave Cost. Chen et al. (2010) have shown that the price of stability in this setting is at least $\frac{3}{2}$, while our results from this chapter show that the price of anarchy is at least n . For models without delays, i.e., where $t_{i,r} = 0$ for all $i \in N, r \in R$, it can be shown that the semi-ordered protocol from von Falkenhausen and Harks (2011) gives an upper bound of n on the price of anarchy, even without the restriction on the players' weights required in von Falkenhausen and Harks (2011). Anshelevich et al. (2008) have further shown the price of stability to be at most logarithmic in n , again for models without delays.

3.9 Conclusion

In this chapter, we considered the design of cost sharing protocols for matroid facility location games. Previously, Chen et al. (2010) had used the price of anarchy as a design metric for cost sharing protocols in network games, and we followed and extended their axiomatization, differentiating between basic and separable protocols. Our first question was whether there are cost sharing protocols that can even in our very general setting give efficient worst-case bounds on the price of anarchy and price of stability. An additional goal was a tractable way of computing cost shares for a given instance that are optimal with respect to the price of anarchy and price of stability.

We exposed that only a subset of strategy profiles can be pure Nash equilibria and their precise characterization through easily verifiable properties termed ‘decharged’ allowed us to answer the first question positively. For a constructive proof of the existence of cheap decharged strategies, we introduced an algorithm that modifies any strategy profile into a decharged profile while increasing the cost by at most a factor of \mathcal{H}_n . Based on such profiles and the algorithm, we further introduced the enforcing protocol, which is separable and has an optimal worst-case

price of anarchy of n and price of stability of \mathcal{H}_n . For the special case of symmetric games without delays, we further proposed the anarchy eliminating protocol with a worst-case price of anarchy of \mathcal{H}_n , based on an extension to the original decharged concept.

While games modeled with player-specific matroids contain many interesting classes (such as the facility location games), other classes such as general multi-commodity networks are not covered by our model. We gave a characterization of pure Nash equilibria with an extension of the decharged concept to general, non-matroid strategy spaces. However, the problem of designing optimal cost sharing protocols for these general strategy spaces still eludes us.

For the second question, we found that it is NP-complete to compute optimal cost shares for a given facility location model, both when optimizing with respect to the price of anarchy or with respect to the price of stability. We proved this with a reduction to HITTING SET. For a limited setting, however, it is possible to compute in polynomial time cost shares that – while not being instance-wise optimal – match our worst-case efficiency bounds. We introduced an algorithm that computes the corresponding semi-proportional cost sharing protocol in polynomial time. While this algorithm achieves the best possible approximation factor for singleton strategies and delays in $\{0, \infty\}$, the case of general matroids with arbitrary delays remains unresolved.

Finally, there is no generic choice of a suitable protocol design space. For applications that allow different means of cost sharing, other axioms defining the protocol space seem reasonable. For instance, requiring the protocols to guarantee convergence of best-response dynamics or to ensure stability against deviations of coalitions of players seem equally natural.

Chapter 4

Quantitative Comparative Statics for a Multimarket Paradox

When a policy maker considers introducing a subsidy for a good, a major question is how this affects the good's market, for example in terms of the equilibrium profit of the firms selling the good. The market can be modeled as a strategic maximization game played by these firms, and the subsidy can be regarded as a parameter that affects the firms' utility functions.

The analysis of equilibrium changes induced by exogenous parameter changes has a long standing tradition in economic literature and is generally referred to as *comparative statics*.¹ As we discuss in detail in Section 4.1, prevailing approaches in comparative statics are of a mostly qualitative nature, examining the monotonicity of marginal parameter changes. Our goal is to use the quantitative worst-case style of the price of anarchy to explore the effect of parameter changes on equilibria. The price of anarchy itself can in certain cases provide a rough quantitative bound for such parameter changes, but it is not the approach of choice. Instead, we develop an approach that shares the price of anarchy's quantitative nature and worst-case perspective, but is tailored to the problems targeted with comparative statics. We apply this new approach to a classic multimarket Cournot oligopoly model, where we model price shocks as parameter changes.

Observing the *maximum possible* effect of a parameter change – shift of the inverse demand function in our case – with a worst-case approach exhibits both

- (1) significance: are changes in a given parameter worth considering?
- (2) robustness: how sensitive is the game to changes of a parameter?

Significance is a crucial motivation of both the analysis of an effect and discussion of whether it can be put to use (à la ‘should a new tax be introduced?’). Robustness on the other hand is important when there is uncertainty about the values of parameters and when parameters change over time.

¹‘Statics’ is the analysis of a system’s static states, i.e., its equilibria.

A Paradox in Multimarket Cournot Oligopoly. We apply our quantitative approach to the multimarket oligopoly model introduced by [Bulow, Geanakoplos, and Klemperer \(1985\)](#). They investigated how “changes in one market have ramifications in a second market” and discovered that a positive price shock in a firm’s monopoly market can have a negative effect on the firm’s profit by influencing competitors’ strategies in a different market. This counterintuitive phenomenon led them to the classification of markets in terms of strategic substitutes and strategic complements.²³ Using our new methodology, we give the first quantification of the profit effects induced by price shocks in multimarket Cournot oligopolies.

Motivating Example. To illustrate the following abstract definition, we recall the original example of [Bulow et al. \(1985\)](#). There are two markets $\{1, 2\}$ and two firms $\{a, b\}$, firm a is a monopolist on market 1 and competes with firm b on market 2. Demand is infinitely elastic on market 1 for the constant price $p_1 = 50$. On market 2, there is an affine price function given by $p_2(q_{a,2} + q_{b,2}) = 200 - (q_{a,2} + q_{b,2})$, where $q_{a,2}, q_{b,2}$ denote the quantities offered by the respective firms on market 2. Production costs are symmetric and given by $c(q) = \frac{1}{2}q^2$, where q is the total quantity produced by a firm.⁴ In the Cournot equilibrium, we obtain $q_{a,1} = 0$ and $q_{a,2} = q_{b,2} = 50$ and each firm earns profits 3750.

Suppose now that market 1 experiences a positive price shock raising its constant price by five units to 55. The Cournot equilibrium changes to $q_{a,1} = 8$ and $q_{a,2} = 47$, $q_{b,2} = 51$. Under this new equilibrium, firm b increases its profit to 3901.5 while firm 1 obtains after the price shock a profit of 3721.5. As noted by [Bulow et al.](#), the *positive* price shock to market 1 has hurt firm a . The actual profit reduction for firm a amounts to 0.76% of the original profit. A natural question to ask is: how much can a firm lose from a positive price shock – especially in a monopoly market?

Quantitative Comparative Statics. Our goal is an analysis of parameter changes that yields a quantitative result while only using functional properties of the underlying model as precise information may not be known, e.g. to a policy maker. The below formulation of quantitative comparative statics was developed to formalize such an approach.

Following the framework of [Milgrom and Roberts \(1994\)](#), we denote by \mathcal{G} a class of strategic maximization games describing what the modeler knows about the economic environment, e.g. Cournot competition on a given number of markets with convex cost technologies. Suppose there is an objective function $f : \mathcal{G} \rightarrow \mathbb{R}$ (e.g., welfare of the unique equilibrium outcome) to evaluate these games and denote for any $G \in \mathcal{G}$ by Δ_G all parameter changes that are to be considered. We express the effect of a parameter change $\delta \in \Delta_G$ on a game $G \in \mathcal{G}$ (assuming for

²See Section VII in [Bulow et al. \(1985\)](#).

³In a market with *strategic substitutes*, more aggressive play by a firm leads to less aggressive play of the competitors on that market; with *strategic complements*, more aggressive play results in more aggressive play of the competitors.

⁴[Bulow et al. \(1985\)](#) assume additional fixed cost $F > 0$ to prevent firms from setting up multiple plants. Fixed costs, however, do not change the resulting equilibria assuming that the access to markets is exogenously determined.

the moment $f(G) \neq 0$ as the value

$$\gamma^f(G, \delta) = \frac{f(G(\delta))}{f(G)},$$

where $G(\delta) \in \mathcal{G}$ denotes the changed game. For $f(G) = 0$, we set $\gamma^f(G, \delta) = 1$, if $f(G(\delta)) \geq 0$, and $\gamma^f(G, \delta) = -\infty$, if $f(G(\delta)) < 0$. The maximum possible effect across all games in \mathcal{G} and their respective parameter changes $\Delta_{\mathcal{G}}$ is then defined as

$$\gamma^f(\mathcal{G}, \Delta_{\mathcal{G}}) = \inf_{G \in \mathcal{G}} \inf_{\delta \in \Delta_G} \gamma^f(G, \delta).$$

Note that if we are interested in the opposite direction of f , we can replace the infima with suprema in the above definition.

4.1 Comparative Statics

Comparative statics have received wide attention in economics, an overview is provided in [Dixit \(1986\)](#). Examples of the types of parameter changes studied with comparative statics are the exposure to international trade ([Krugman 1980, Melitz 2003](#)), a forced reduction of produced quantity ([Gaudet and Salant 1991](#)), changes of the payoff functions via a demand shift ([Quirmbach 1988](#)), a cost shift ([Février and Linnemer 2004](#)), or the introduction of export taxes/subsidies ([Brander and Spencer 1985, Eaton and Grossman 1986](#)).

The prevailing approaches in comparative statics up to date rely on using either the implicit function theorem⁵ (e.g. [Bulow et al. 1985](#)) or the powerful machinery of lattice theory applicable to supermodular games, see [Amir \(1996\)](#), [Edlin and Shannon \(1998\)](#), [Kukushkin \(1994\)](#), [Milgrom and Shannon \(1994\)](#), [Milgrom and Roberts \(1990, 1994\)](#), [Shannon \(1995\)](#) and [Topkis \(1979, 1998\)](#), also [Athey \(2002\)](#), [Quah \(2007\)](#), and [Vives \(1990, 2005\)](#).

Since [Bulow et al. \(1985\)](#) introduced the concepts of strategic substitutes and strategic complements, a considerable amount of research has used comparative statics to investigate situations where strategic substitutes or complements occur. Notably, [Brander and Spencer \(1985\)](#) found that an export subsidy can increase welfare in Cournot competition with strategic substitutes. [Eaton and Grossman \(1986\)](#) extended this model to a two-stage game where first governments set policies and then firms engage in competition. [Gaudet and Salant \(1991\)](#) looked at how a forced marginal change of production quantity for a subset of firms impacts profit in a situation with strategic substitutes. [Février and Linnemer \(2004\)](#) gave a decomposition of price shocks in Cournot oligopoly into an average effect common to all firms and a heterogeneity effect that is firm-specific. [Acemoglu and Jensen \(2013\)](#) present a framework for comparative statics results for a superclass of Cournot oligopoly called aggregative games, see also [Cochón \(1994\)](#).

⁵The application of the implicit function theorem for instance in oligopoly models requires some regularity assumptions such as convexity and smoothness of cost and inverse demand functions, see the discussion in [Milgrom and Roberts \(1994\)](#).

Difference to Quantitative Comparative Statics

There are three conceptual differences between comparative statics analysis and our approach. First, classic comparative statics analysis is concerned with the *monotonicity* of an effect, for example, whether a subsidy increases a firm's profits and under what circumstances it decreases a firm's profits. However, this leaves the quantification of the effect open, motivating our quantitative approach.

Second, our quantitative comparative statics allow for *discrete* parameter changes, contrasting approaches using the implicit function theorem or standard sensitivity analysis in optimization based on first and second order conditions that rely on a *local analysis* of parameter changes, and, thus, do not allow for meaningful conclusions if parameter changes happens to be discrete.

Finally, in practice, the economic model including the endogenous variables and parameters may not be known precisely but only approximately by knowing some functional property of the relations of endogenous variables or the space of parameters. The nature of a worst-case analysis such as the proposed quantitative comparative statics is that it provides results across a class of games \mathcal{G} described for example by functional properties and a set $\Delta_{\mathcal{G}}$ of parameter changes. Both \mathcal{G} and $\Delta_{\mathcal{G}}$ need to be chosen reasonably to yield a meaningful result (any combination of positive price shocks in our case). This is contrasted by the assumption of *perfect knowledge* of all elements of the precise instance G including the precise parameter change δ common to most existing approaches in comparative statics.⁶

4.2 Contribution and Outline

We conduct quantitative comparative statics analysis for multimarket oligopolies with affine price functions and convex cost technologies of firms, a class of games that contains the above example. We consider positive price shocks as parameter changes and three different objective functions: the individual profit of a firm, welfare measured by summing up the firm's profits, and social surplus defined by integrating the price functions and subtracting the firm's costs. In Section 4.4 we find that a firm's profit can be reduced at most 25% by a positive price shock. We give the bound as a function of the number of firms, showing e.g. that in the two firm case the profit reduction is a most 6.25% and provide a matching lower bound instance. When we relax our assumptions and allow either non-convex cost functions or concave price functions, we find that the profit loss can be arbitrarily high. In Section 4.5, we turn to the effects of parameter changes on welfare and social surplus. We show that the effect on welfare (aggregate profit) can be as strong as on individual firm profit. Additionally, we formalize the relationship between quantitative comparative statics for welfare and price of anarchy results. Finally, we find that a game's social surplus is more resilient to parameter changes, as it can only be reduced by up to 16.7% by a positive price shock.

Our results give the first quantitative comparative statics result of an important paradoxical phenomenon previously only qualitatively analyzed. Each of our bounds is complemented by an example instance that actually attains the worst case ratio. Moreover, our results immediately

⁶An exception are the papers Milgrom and Roberts (1994) and Milgrom (1994), where monotone comparative statics results are derived for classes of functions without knowing the precise instantiation.

extend to negative price shocks, as any negative price shock can be seen as taking back a positive price shock. The profit and welfare gain from a negative price shock is no more than 33.4%, and the social surplus can increase by up to 20%.

Example Application. Profit gains from negative price shocks can occur in international trade, as noted by [Bulow et al. \(1985, Sec. VI \(C\)\)](#). Consider two markets located in separate countries with convex cost technologies, one of which is a monopoly market for some firm a . A tax change in the country of the monopolist can be considered a price shock. A government may decide to increase domestic taxes in order to increase firm a 's profitability in the foreign market. Our results imply that this positive effect can be significant as it may increase the profitability by up to 33% of current profits.

Bibliographic Information. The results presented in this chapter are joint work with Tobias Harks. Parts of the results contained in this chapter appeared in *Proceedings of the 9th International Conference on Web and Internet Economics*, see [Harks and von Falkenhausen \(2013\)](#).

4.3 Multimarket Cournot Oligopoly Model

In this section we introduce the specific model for which we investigate quantitative comparative statics. As outlined above, the three ingredients for such analysis are a set of games \mathcal{G} , for each such game $G \in \mathcal{G}$ as set of feasible parameter changes Δ_G and an objective function $f : \mathcal{G} \rightarrow \mathbb{R}$ to evaluate the games.

Class of Games \mathcal{G} . We consider multimarket oligopoly competition defined as follows: A set N of n firms competes on a set of markets M . Each firm $i \in N$ has access to some subset $M_i \subseteq M$ of these markets. A strategy of a firm i is to choose production quantities $q_{i,m} \geq 0$ for all markets $m \in M_i$ that it serves. We set $q_{i,m} = 0$ for any market not served by firm i and denote the total quantity of firm i by $q_i = \sum_{m \in M} q_{i,m}$. Correspondingly, the total quantity on any market m is denoted by $q_m = \sum_{i \in N} q_{i,m}$.

Assumption 4.1. The price of a market is an affine function of the total quantity produced, i.e., if on market $m \in M$ a quantity q_m is produced the price is $p_m(q_m) = p_m(0) - r_m q_m$, where $p_m(0)$ is an initial positive price and the coefficient $r_m > 0$ describes how the price decreases as the demand is satisfied.

Assumption 4.2. Firm i 's cost for producing the quantity q_i is given by the function $c_i(q_i)$ which we assume to be non-decreasing, convex and differentiable in q_i with $c_i(0) = 0$.

Multimarket Cournot oligopoly can be embedded into the strategic game framework.

Definition 4.3 (multimarket Cournot oligopoly games)

Denoted in the following by \mathcal{G} , multimarket Cournot competition is the class of strategic maximization games given as a tuple $G = (N, M, (M_i)_{i \in N}, (c_i)_{i \in N}, (p_m)_{m \in M})$, where the strategy space

of a firm $i \in N$ is the space of non-negative production quantity vectors, i.e., $X_i = \mathbb{R}_{\geq 0}^{|M_i|}$. The utility of firm i is defined as

$$u_i(q) = \sum_{m \in M} p_m(q_m) q_{i,m} - c_i(q_i).$$

In reference to [Cournot \(1838\)](#), pure Nash equilibria in these games are called *Cournot equilibrium*. In a Cournot equilibrium q , no firm can increase its utility by unilaterally deviating to a different strategy, i.e., $u_i(q) \geq u_i(q'_i, q_{-i})$ for all strategies q'_i available to firm i . As the games introduced above involve convex and compact strategy spaces together with quasi-concave utilities, standard fixed-point arguments of type Kakutani (cf. [Debreu 1952](#), [Fan 1952](#), [Glicksberg 1952](#), [Kakutani 1941](#)) imply the existence of an equilibrium. As shown in Section 4.4, the assumptions on the price and cost functions further imply that there is a unique equilibrium.

We denote the marginal revenue of firm i on market m by

$$\pi_{i,m}(q_{i,m}, q_{-i,m}) = \frac{\partial (p_m(q_{i,m} + q_{-i,m}) q_{i,m})}{\partial q_{i,m}} = p_m(q_m) - r_m q_{i,m},$$

where $q_{-i,m}$ is the quantity sold by firms $j \neq i$. The marginal cost is $c'_i(q_i)$, and we often write $\pi_{i,m}(q)$ and $c'_i(q)$. In an equilibrium x , the marginal revenue of any served market m equals the marginal cost: $\pi_{i,m}(x) = c'_i(x)$ for all $i \in N$ with $x_{i,m} > 0$.

Parameter Changes Δ_G . The parameter changes analyzed are positive shocks to the price functions: on every market m of a game $G \in \mathcal{G}$ the price increases by some amount $\delta_m \geq 0$ such that $p_m^\delta(q_m) = p_m(q_m) + \delta_m$. We denote the set of feasible parameter changes as $\Delta_G = \mathbb{R}_{\geq 0}^{|M|}$, where $|M|$ is the number of markets in G .

Objective Functions f . The games in \mathcal{G} have unique Cournot equilibria (as shown in Section 4.4), and our three objective functions are equilibrium firm profit, equilibrium welfare and equilibrium social surplus. Given a game G and a price shock δ , we denote the unique equilibrium of G by x and the unique equilibrium of $G(\delta)$ by y .

- (1) *firm profit*: the profit of an individual firm is $u_i(q) = \sum_{m \in M} p_m(q_m) q_{i,m} - c_i(q_i)$ and we minimize across the firms of a game to obtain an overall measure:

$$\gamma^u(G, \delta) = \min_{i \in N} \frac{u_i^\delta(y)}{u_i(x)}.$$

- (2) *welfare*: we consider utilitarian welfare, that is, the welfare is obtained as the sum of the profits of all firms $U(q) = \sum_{i \in N} u_i(q)$, such that

$$\gamma^U(G, \delta) = \frac{U^\delta(y)}{U(x)}.$$

- (3) *social surplus*: this measure assumes that the price function of a market expresses the marginal value that the buyers in the market have from an additional quantity of the good. We denote $S(q) = \sum_{m \in M} \int_0^{q_m} p_m(z) dz - \sum_{i \in N} c_i(q_i)$, such that

$$\gamma^S(G, \delta) = \frac{S^\delta(y)}{S(x)}.$$

While the first measure has been analyzed by [Bulow et al. \(1985\)](#), the second two measures have been used among others by [Anderson and Renault \(2003\)](#), [Ushio \(1985\)](#) and [Tsitsiklis and Xu \(2013\)](#). For each measure, we strive to provide a infimum bound across all multimarket Cournot games in \mathcal{G} and price shocks $\Delta_{\mathcal{G}}$ combined with concrete games that converge to this bound. The following example instance exhibits the basic intuition for the quantitative analysis in Sections 4.4 and 4.5.

Example Instance. Consider a game with two markets. Market 1 is served only by a monopolist (denoted as firm a), while all firms compete in market 2. Marginal profit and cost of both firm a and a competitor $i \neq a$ are illustrated in Figure 4.1a. Given a positive price shock on market 1, firm a reduces its quantity on market 2 in favor of selling more on its more profitable monopoly market. In effect, the competitors' marginal revenue strictly increases and leads to a new equilibrium in which these competitors increase their quantities, see Figure 4.1b. Markets in which a less aggressive play of one firm leads to a more aggressive play of competitors are called markets with strategic substitutes ([Bulow et al. 1985](#)). The more aggressive competition experienced by firm a in market 2 reduces its profit below the original level, or said differently, after the price shock a part of firm a 's quantity on market 2 has been substituted by the competitors, see Figure 4.1c.

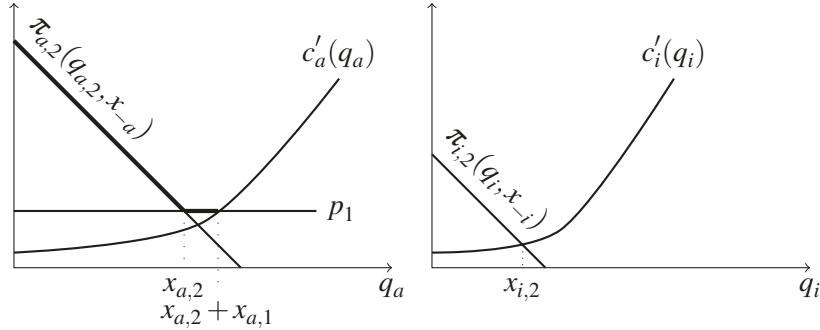
Negative Price Shocks. While we restrict our analysis to positive price shocks, the results immediately extend to negative price shocks. If for example a subsidy (i.e., a positive price shock) causes the equilibrium to shift such that the firms' profits decrease, then taking back the subsidy (i.e., a negative price shock) restores the old equilibrium and thus increase the firm's profits. In this sense the two effects are dual: any negative price shock can be seen as taking back a positive price shock, and the profit gain from the negative price shock is equal to the profit loss from positive price shock. This is true for any objective function f and, denoting negative price shocks to a game by $\Delta_{\mathcal{G}}^-$ we have $\gamma^f(\mathcal{G}, \Delta_{\mathcal{G}}^-) = (\gamma^f(\mathcal{G}, \Delta_{\mathcal{G}}))^{\text{-1}}$.

4.4 Maximum Profit Loss of an Individual Firm

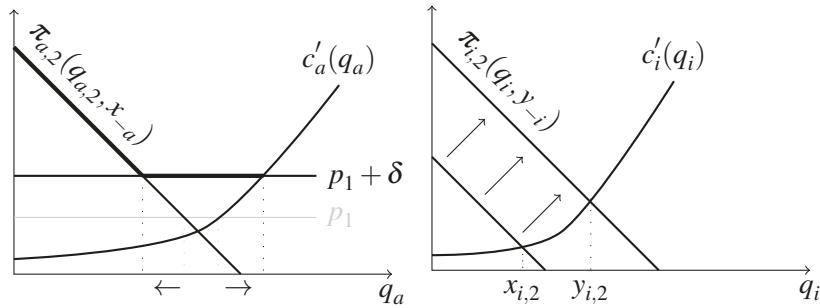
We use quantitative comparative statics to investigate the worst case profit loss of a firm from a positive price shock as expressed by $\gamma^u(\mathcal{G}, \Delta_{\mathcal{G}})$.

Theorem 4.4. *Given a game G with n firms, no firm loses more than a $\frac{(n-1)^2}{4n^2}$ fraction of its profit from a price shock $(\delta_m)_{m \in M}$ with $\delta_m \geq 0$ for all $m \in M$,*

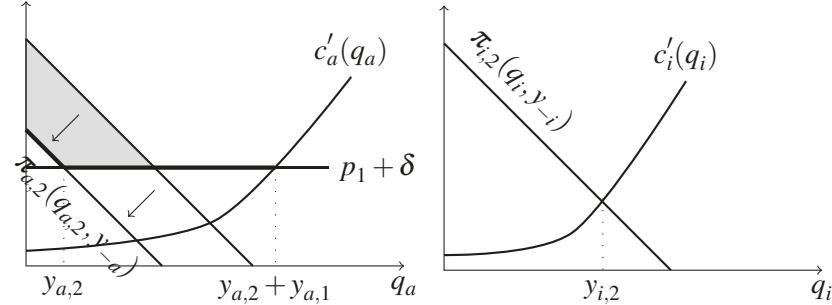
$$\gamma^u(G, \delta) \geq 1 - \frac{(n-1)^2}{4n^2} \geq \frac{3}{4}.$$

Figure 4.1: Example instance with two markets

- (a) Initial equilibrium x :** The price on market 1 is constant at p_1 (thus also firm a 's marginal revenue). On market 2 the price is decreasing, such that firm a sees marginal revenue $\pi_{a,2}(\cdot, x_{-a})$, given its competitors' equilibrium quantities x_{-a} . Firm a produces $x_{a,2}$ on market 2 such that the marginal revenue there is equal to the marginal revenue p_1 on market 1. Its production on market 1 is such that the marginal cost of the aggregate quantity $x_{a,1} + x_{a,2}$ is equal to p_1 . The competitors $i \neq a$ have marginal revenue $\pi_{i,2}(\cdot, x_{-i})$ on market 2. They choose their equilibrium quantity $x_{i,2}$ at the intersection of marginal cost and marginal revenue.



- (b) Price shock triggers strategic substitution:** The shock leads firm a to shift production from market 2 to market 1. This increases the marginal revenue of competitors on market 2, causing increased production $y_{i,2} > x_{i,2}$.



- (c) Effect on firm a 's profit:** This derogates firm a 's marginal revenue on market 2, causing it to further withdraw from the market. The shaded area indicates the profit loss from the strategic substitution, which is partially compensated by profit gain from the increased price on market 1.

This is the main result of this chapter. It shows that no firm loses more than 25% of current profits. This bound is *robust* in the sense that it holds for an entire class of games and parameters, that is, in order to arrive at this bound the modeler only needs to know that price shocks are non-negative, inverse demand functions are affine, and cost technologies are convex.

To prove the statement, we first establish uniqueness of equilibria and that a price shock $\delta \geq 0$ causes the price on every market to increase, and that in this favorable setting every firm increases its total quantity. Using these insights into the effect of the price shock, for any given firm i we can identify a market where this firm suffers the relatively strongest loss and use this to bound $\frac{u_i^\delta(y)}{u_i(x)}$. This proves the theorem, as $\gamma^\mu(G, \delta)$ is the minimum of this fraction across all firms.

Lemma 4.5 (Uniqueness of Equilibria). *Let x and y be equilibria of some game G . Then $x = y$.*

Proof. As a first step of the proof, we show that on every market m , $x_m = y_m$. Let $M^+ = \{m \in M : x_m < y_m\}$ and assume for a contradiction that $M^+ \neq \emptyset$. Then, there is a firm i with $\sum_{m \in M^+} (y_{i,m} - x_{i,m}) > 0$ and a market $m \in M^+$ with $y_{i,m} > x_{i,m} \geq 0$. It follows from the equilibrium definition that in y firm i 's marginal cost and marginal revenue on m are equal, i.e.,

$$c'_i(y_i) = p_m(y) - r_m y_{i,m} < p_m(x) - r_m x_{i,m} \leq c'_i(x), \quad (4.1)$$

where we used that $p_m(y) < p_m(x)$ because $m \in M^+$. From $c'_i(y) < c'_i(x)$ we follow that

$$y_i < x_i. \quad (4.2)$$

Also, for all markets $m' \in M$ where $y_{i,m'} < x_{i,m'}$, we follow again from the equilibrium definition that

$$p_{m'}(y) - r_{m'} y_{i,m'} \leq c'_i(y_i) < c'_i(x) \leq p_{m'}(x) - r_{m'} x_{i,m'},$$

and hence, $p_{m'}(y) < p_{m'}(x)$, i.e., $m' \in M^+$. Then, we find a contradiction as

$$0 > y_i - x_i = \sum_{m \in M} (y_{i,m} - x_{i,m}) > \sum_{m \in M^+} (y_{i,m} - x_{i,m}) > 0.$$

Here, we can limit the summation from all $m \in M$ to $m \in M^+$ because we found that $m' \in M^+$ for all markets with $y_{i,m'} < x_{i,m'}$.

As the next step of the proof, we use $x_m = y_m$ for all $m \in M$ to show $x_{i,m} = y_{i,m}$ for all firms i . For a contradiction, assume there are $i \in N$ and $m \in M$ such that $x_{i,m} < y_{i,m}$. Then, we can again apply (4.1) to obtain $y_i < x_i$ and there must be some market $m' \in M$ with $y_{i,m'} < x_{i,m'}$, which leads with the same reasoning to $x_i < y_i$, a contradiction. Altogether, $x = y$. \square

We now show that the prices on all markets increase after the positive price shock.

Lemma 4.6. *Let x and y be equilibria of a game G before and after a price shock $(\delta_m)_{m \in M}$ with $\delta_m \geq 0$ for all $m \in M$. Then, on all markets $m \in M$ the price in y is higher than in x , that is, $p_m^\delta(y_m) \geq p_m(x_m)$.*

Proof. While in the proof of Lemma 4.5 we compared two equilibria of the same game, we now compare the equilibria before and after the price shock. The analysis remains largely unchanged: if we denote by $M^+ = \{m \in M : x_m + \frac{\delta_m}{r_m} < y_m\}$ the set of markets where the price decreases, then $M^+ \neq \emptyset$ still implies that there is some firm i with $\sum_{m \in M^+} (y_{i,m} - x_{i,m}) > 0$ and $y_i < x_i$ as in (4.2), leading to the same contradiction as before. We follow $M^+ = \emptyset$, i.e., $p_m^\delta(y_m) \geq p_m(x_m)$ for all $m \in M$. \square

Given increasing prices, all firms increase their quantity.

Lemma 4.7. *Let x and y be equilibria of a game G before and after a price shock $(\delta_m)_{m \in M}$ with $\delta_m \geq 0$ for all $m \in M$. Then, each firm $i \in N$ produces more in y than in x , $y_i \geq x_i$.*

Proof. Assume there is $i \in N$ with $y_i < x_i$. Then

$$\pi_{i,m}^\delta(y) \leq c'_i(y_i) \leq c'_i(x_i) = \pi_{i,m}(x)$$

on every market $m \in M$ with $x_{i,m} > 0$. With $p_m^\delta(y) \geq p_m(x)$ as found in the previous lemma, we follow $y_{i,m} \geq x_{i,m}$ on every market with $x_{i,m} > 0$, a contradiction to $y_i < x_i$. Hence $x_i \leq y_i$ for all $i \neq a$. \square

We are now ready to prove the main theorem.

Proof of Theorem 4.4. Let G be a game with n firms. We show that for any given firm i , $\frac{u_i^\delta(y)}{u_i(x)} \geq 1 - \frac{(n-1)^2}{4n}$. The theorem follows because $\gamma^\mu(G, \delta)$ is the minimum of this quantity across all firms.

Denote the set of markets where firm i decreases their quantity after the price shock by $M^- = \{m \in M : y_{i,m} < x_{i,m}\}$ and similarly the set where i increases their quantity after the price shock by $M^+ = M \setminus M^-$.

We assume that the markets are indexed such that market 1 is a solution to

$$\operatorname{argmin}_{m \in M^-} \frac{(p_m^\delta(y) - c'_i(x))x_{i,m} - r_m y_{i,m}(x_{i,m} - y_{i,m})}{(p_m(x) - c'_i(x))x_{i,m}}. \quad (4.3)$$

Note that the denominator of the above fraction is always positive as any market $m \in M^-$ has a non-zero quantity $x_{i,m} > 0$ and thus by the first order equilibrium condition also $p_m(x) > c'_i(x)$.

We find the following relations that will be helpful later on: The quantity added on markets in M^+ corresponds exactly to the quantity taken away from markets in M^- and the additional quantity $y_i - x_i$, i.e.

$$\sum_{m \in M^+} y_{i,m} - x_{i,m} = \sum_{m \in M^-} x_{i,m} - y_{i,m} + y_i - x_i. \quad (4.4)$$

Also, the price on markets $m \in M^+$ in y is related to the marginal cost, i.e.,

$$p_m^\delta(y) \geq p_m^\delta(y) - r_m y_{i,m} = \pi_{i,m}^\delta(y) = c'_i(y) \quad (4.5)$$

which in turn is related to the price on markets $m \in M^-$, i.e.,

$$c'_i(y) \geq \pi_{i,m}^\delta(y) = p_m^\delta(y) - r_m y_{i,m}. \quad (4.6)$$

As the cost function is convex,

$$c'_i(y)(y_i - x_i) \geq c_i(y) - c_i(x). \quad (4.7)$$

We combine the above to a statement that relates the profit of quantity added to M^+ to the profit lost by reducing quantity in M^- ,

$$\begin{aligned} \sum_{m \in M^+} p_m^\delta(y)(y_{i,m} - x_{i,m}) &\stackrel{(4.5)}{\geq} \sum_{m \in M^+} c'_i(y)(y_{i,m} - x_{i,m}) \\ &\stackrel{(4.4)}{\geq} \sum_{m \in M^-} c'_i(y)(x_{i,m} - y_{i,m}) + c'_i(y)(y_i - x_i) \\ &\stackrel{(4.6),(4.7)}{\geq} \sum_{m \in M^-} (p_m^\delta(y) - r_m y_{i,m})(x_{i,m} - y_{i,m}) + c_i(y) - c_i(x). \end{aligned} \quad (4.8)$$

We further assume that $\frac{u_i^\delta(y)}{u_i(x)} < 1$, as we are interested in worst case instances and our lower bounds show that such instances exist. Note that for a fraction with value less than 1, subtracting the same amount from both numerator and denominator decreases the value of the fraction. We estimate

$$\begin{aligned} \frac{u_i^\delta(y)}{u_i(x)} &= \frac{\sum_{m \in M^-} p_m^\delta(y)y_{i,m} + \sum_{m \in M^+} p_m^\delta(y)y_{i,m} - c_i(y)}{\sum_{m \in M^-} p_m(x)x_{i,m} + \sum_{m \in M^+} p_m(x)x_{i,m} - c_i(x)} \\ &\stackrel{(4.8)}{\geq} \frac{\sum_{m \in M^-} p_m^\delta(y)x_{i,m} - r_m y_{i,m}(x_{i,m} - y_{i,m}) + \sum_{m \in M^+} p_m^\delta(y)x_{i,m} - c_i(x)}{\sum_{m \in M^-} p_m(x)x_{i,m} + \sum_{m \in M^+} p_m(x)x_{i,m} - c_i(x)} \end{aligned} \quad (4.9)$$

$$\geq \frac{\sum_{m \in M^-} (p_m^\delta(y) - c'_i(x))x_{i,m} - r_m y_{i,m}(x_{i,m} - y_{i,m}) + \sum_{m \in M^+} (p_m(x) - c'_i(x))x_{i,m}}{\sum_{m \in M^-} (p_m(x) - c'_i(x))x_{i,m} + \sum_{m \in M^+} (p_m(x) - c'_i(x))x_{i,m}} \quad (4.9)$$

$$\geq \frac{\sum_{m \in M^-} (p_m^\delta(y) - c'_i(x))x_{i,m} - r_m y_{i,m}(x_{i,m} - y_{i,m})}{\sum_{m \in M^-} (p_m(x) - c'_i(x))x_{i,m}} \quad (4.10)$$

$$\stackrel{(4.3)}{\geq} \frac{(p_1^\delta(y) - c'_i(x))x_{i,1} - r_1 y_{i,1}(x_{i,1} - y_{i,1})}{(p_1(x) - c'_i(x))x_{i,1}}. \quad (4.11)$$

In (4.9) we use that the cost function is convex and hence $-c_i(x) \geq \sum_{m \in M^-} c'_i(x)x_{i,m}$ and in (4.10) that the price on a market with positive quantity is at least the marginal cost, i.e., $p_m(x) \geq c'_i(x)$ for a market with $x_{i,m} > 0$.

We now need to further examine the relation of $p_1^\delta(y)$, $p_1(x)$ and $c'_i(x)$. For any firm $j \neq i$ that has increased its quantity on market 1, i.e., $y_{j,1} > x_{j,1}$,

$$p_1(x) - r_1 x_{j,1} = \pi_{j,1}(x) \leq c'_j(x) \leq c'_j(y) = \pi_{j,1}^\delta(y) = p_1^\delta(y) - r_1 y_{j,1},$$

that is,

$$r_1(y_{j,1} - x_{j,1}) \leq p_1^\delta(y) - p_1(x). \quad (4.12)$$

Then, considering $\sum_{j:y_{j,1} > x_{j,1}} (y_{j,1} - x_{j,1}) + y_{i,1} - x_{i,1} \geq y_1 - x_1$, we can rather precisely observe how the price on market 1 changes with the price shock.

$$\begin{aligned} p_1^\delta(y) &= p_1(x) + \delta_1 - r_1(y_1 - x_1) \geq p_1(x) + r_1(x_{i,1} - y_{i,1}) - r_1 \sum_{j:y_{j,1} > x_{j,1}} (y_{j,1} - x_{j,1}) \\ &\stackrel{(4.12)}{\geq} p_1(x) + r_1(x_{i,1} - y_{i,1}) - (n-1)(p_1^\delta(y) - p_1(x)), \end{aligned}$$

as there are at most $n-1$ firms with $y_{j,1} > x_{j,1}$. This can be rearranged to

$$p_1^\delta(y) \geq p_1(x) + \frac{r_1}{n}(x_{i,1} - y_{i,1}). \quad (4.13)$$

Observe further that $x_{i,1} > 0$ because market 1 is in M^- and hence

$$p_1(x) - r_1 x_{i,1} = \pi_{i,1}(x) = c'_i(x). \quad (4.14)$$

We continue the proof from (4.11),

$$\begin{aligned} \frac{u_i^\delta(y)}{u_i(x)} &\stackrel{(4.11)}{\geq} \frac{(p_1^\delta(y) - c'_i(x))x_{i,1} - r_1 y_{i,1}(x_{i,1} - y_{i,1})}{(p_1(x) - c'_i(x))x_{i,1}} \\ &\stackrel{(4.13)}{\geq} \frac{(p_1(x) - c'_i(x))x_{i,1} + \frac{r_1}{n}(x_{i,1} - y_{i,1})x_{i,1} - r_1 y_{i,1}(x_{i,1} - y_{i,1})}{(p_1(x) - c'_i(x))x_{i,1}} \\ &= 1 + \frac{r_1(\frac{1}{n}x_{i,1} - y_{i,1})(x_{i,1} - y_{i,1})}{(p_1(x) - c'_i(x))x_{i,1}} \\ &\stackrel{(4.14)}{=} 1 + \frac{(\frac{1}{n}x_{i,1} - y_{i,1})(x_{i,1} - y_{i,1})}{x_{i,1}^2} \\ &\geq 1 - \frac{(n-1)^2}{4n^2} = \frac{3}{4} + \frac{2n-1}{4n^2}, \end{aligned} \quad (4.15)$$

using that $y_{i,1} = \frac{n+1}{2n}x_{i,1}$ minimizes (4.15). \square

4.4.1 Lower Bound

To show that the bound of the previous theorem is tight, we construct a simple instance with matching profit loss.

Proposition 4.8. *For any n , there is a game G with n firms where a positive price shock δ decreases the profit of some firm a by a factor $\frac{(n-1)^2}{4n^2}$, that is,*

$$\gamma^\mu(G, \delta) = 1 - \frac{(n-1)^2}{4n^2}.$$

Proof. The instance has two markets $M = \{1, 2\}$ and there are n firms. All firms serve market 2 while market 1 is only served by some firm $a \in N$. We fix the price on market 1 before the price shock to 0, i.e., $p_1(q_1) \equiv 0$ and $r_1 = 0$. On market 2 the price is $p_2(q_2) = 2 - q_2$, where q_2 is the total quantity sold in market 2. The cost of firm a for any total quantity $q_a = q_{a,1} + q_{a,2}$ is 0 if $q_a \leq \frac{2}{n+1}$; for any larger quantity the cost is prohibitively high. For firms $i \neq a$, the cost is always 0.

In the initial Cournot equilibrium x , no quantity is sold on market 1 and on market 2, $x_{a,2} = x_{i,2} = \frac{2}{n+1}$. Firm a 's profit is $u_a(x) = (2 - n\frac{2}{n+1})\frac{2}{n+1} = \frac{4}{(n+1)^2}$.

A price shock that increases the price on market 1 to $\frac{n-1}{n^2}$ leads firm a to shift to market 1. In the new equilibrium y , $y_{a,1} = \frac{n-1}{n(n+1)}$, $y_{a,2} = \frac{1}{n}$, and $y_{i,2} = \frac{2n-1}{n^2}$. The profit of firm a is

$$u_a(y) = \left(2 - \frac{1}{n} - (n-1)\frac{2n-1}{n^2}\right)\frac{1}{n} + \frac{n-1}{n^2}\frac{n-1}{n(n+1)} = \frac{3n-1}{n^2(n+1)}.$$

Then, the ratio of profit before and after the price shock is

$$\gamma^\mu(G, \delta) = \frac{u_a(y)}{u_a(x)} = \frac{(3n-1)(n+1)}{4n^2} = 1 - \frac{(n-1)^2}{4n^2}.$$

□

Remark 4.9. Note that this lower bound is quite generic in the sense that such an instance can be constructed for any price function on market 2 and any linear cost function for competitors $i \neq a$. In general, the profit loss of a firm can be large when it has a strongly convex cost function, such that a positive price shock in one market causes it to decrease quantity in another market, and when this is met by competitors with linear (or not “too convex”) cost functions.

4.4.2 Non-convex Cost

If we relax Assumption 4.2 and allow non-convex cost functions, we possibly lose uniqueness of equilibria and we would have to redefine our objective function, e.g. involving equilibrium selection.⁷ Moreover, one can easily construct examples where a positive price shock completely eliminates the profit of firm a in all equilibria. If e.g. fixed costs are allowed, in the example of [Bulow et al. \(1985\)](#) one can set the fixed cost of the monopolist equal to their revenue after the price shock using the fact that fixed costs do not change the equilibria of the game as long as non-negative profits are guaranteed. Similar examples are possible if we mix between economies of scale and diseconomies of scale among firms' cost technologies.

⁷Cournot equilibria continue to exist for non-convex costs if inverse demand functions satisfy rather mild assumptions, see, e.g., [Novshek \(1985\)](#) and [Amir \(1996\)](#) and [Roberts and Sonnenschein \(1976\)](#).

4.4.3 Concave Inverse Demand Functions

Relaxing Assumption 4.1 toward concave prices reveals another counterintuitive phenomenon: Very small price shocks may decrease the profit of a firm by an arbitrary amount. Consider the class $\bar{\mathcal{G}} \supseteq \mathcal{G}$ that allows for concave inverse demand functions. We obtain the following value for $\gamma^u(\bar{\mathcal{G}}, \Delta_{\bar{\mathcal{G}}})$.

Proposition 4.10. *For any $n \geq 4$, there is a game with only two markets and concave price functions such that the profit ratio of one of the firms before and after a positive price shock is less than $\frac{2}{n}$. Thus, $\gamma^u(\bar{\mathcal{G}}, \Delta_{\bar{\mathcal{G}}}) = 0$.*

Proof. Fix some $n \geq 4$. We construct a game $G \in \bar{\mathcal{G}}$ with n firms and two markets $M = \{1, 2\}$ to fulfill the above claim. The firm whose profit ratio we observe, denoted by a , has cost $c_a(q_a) = 0$ for quantities $q_a \leq 1$ and prohibitively high cost for larger quantities. All other firms $i \neq a$ have cost $c_i(q_i) = q_i$ for any quantity.

Market 1 is only served by firm a and has constant price $p_1(q_1) \equiv 0$. Market 2 is served by all firms and has a concave price function satisfying $p_2(1) = 1$ with $p'_2(1) = -1$ and $p_2(1 + \frac{1}{n}) = 1 - \frac{2}{n}$ with $p'_2(1 + \frac{1}{n}) = -n$.

The initial equilibrium x of this game is $x_{a,1} = 0$, $x_{a,2} = 1$ and $x_{j,2} = 0$ for all $j \neq a$. To verify this, observe that marginal revenue and cost of firm a are all equal as $\pi_{a,1}(x) = 0$, $\pi_{a,2}(x) = 1 - 1 = 0$ and $c'_a(x) = 0$ as well as for competitors $j \neq a$ we have $\pi_{j,2}(x) = 1 = c'_j(x)$.

Let the price shock be $\delta_1 = \frac{1}{n}$ and $\delta_2 = \frac{3}{n}$. The new equilibrium is $y_{a,1} = 1 - \frac{1}{n}$, $y_{a,2} = \frac{1}{n}$ and $y_{j,2} = \frac{1}{n^2}$ for all competitors $j \neq a$. We again verify $\pi_{a,1}^\delta(y) = \frac{1}{n}$, $\pi_{a,2}^\delta(y) = 1 + \frac{1}{n} - n \frac{1}{n} = \frac{1}{n}$ which are equal and greater than 0. For competitors $j \neq a$, $\pi_{j,2}^\delta(y) = 1 + \frac{1}{n} - n \frac{1}{n^2} = 1 = c'_j(y)$.

We calculate the profit of firm a in both equilibria: $u_a(x) = 1$ and thus

$$\begin{aligned}\gamma(G, \delta) &= u_a^\delta(y) = p_1^\delta(y)y_{a,1} + p_2^\delta(y)y_{a,2} - c_a(y) \\ &= \frac{1}{n}(1 - \frac{1}{n}) + (1 + \frac{1}{n})\frac{1}{n} = \frac{2}{n}.\end{aligned}$$

□

The above bound largely depends on the number of firms in the game. For games with only two firms we show that the profit loss can be up to 25%, as opposed to the 6.25% in the case of affine price functions.

Proposition 4.11. *For any value $\gamma > \frac{3}{4}$, there is a game $G \in \bar{\mathcal{G}}$ with only two firms and concave price functions and a price shock $\delta \in \Delta_G$ such that $\gamma \geq \gamma^u(G, \delta)$ is a bound on the profit ratio of one of the firms before and after a positive price shock.*

Proof. Fix some $k > 4$. We construct a game G with two firms a and b and two markets $M = \{1, 2\}$, such that the profit ratio $\gamma(G, \delta)$ of firm a before and after a price shock δ is dependent on k .

The firm a has cost $c_a(q_a) = (k - 1)q_a$ for quantities $q_a \leq 1$ and prohibitively high cost for larger quantities. Firm b has cost $c_b(q_b) \equiv 0$ for any quantity.

Market 1 is only served by firm a and has price $p_1(q_1) = k - 1 - \frac{1}{k}q_1$. Market 2 is served by both firms and has a concave price function satisfying $p_2(k + 1 - \frac{1}{k}) = k + \frac{1}{k}$ and $p_2(k + 1) = k$ with slopes $p'_2(k + 1 - \frac{1}{k}) = -\frac{k+1}{k+\frac{1}{2}-\frac{1}{k}}$ and $p'_2(k + 1) = -1$. Note that $k \geq 4$, hence, such a concave function exists.

The initial equilibrium x of this game is $x_{a,1} = 0$, $x_{a,2} = 1$ and $x_{b,2} = k$. To verify this, observe that marginal revenue and cost of firm a are all equal as $\pi_{a,1}(x) = k - 1$, $\pi_{a,2}(x) = k - 1$ and $c'_a(x) = k - 1$ as well as for firm b we have $\pi_{b,2}(x) = 0$.

Let the price shock be $\delta_1 = 1 + 2\frac{1}{k} - \frac{k+\frac{1}{k}}{2(k-\frac{1}{k}+\frac{1}{2})}$ and $\delta_2 = 0$. The new equilibrium is $y_{a,1} = \frac{1}{2}$, $y_{a,2} = \frac{1}{2}$ and $y_{b,2} = k - \frac{1}{k} + \frac{1}{2}$. We again verify $\pi_{a,1}^\delta(y) = k + \frac{1}{k} - \frac{k+\frac{1}{k}}{2(k-\frac{1}{k}+\frac{1}{2})}$, $\pi_{a,2}^\delta(y) = k + \frac{1}{k} - \frac{k+\frac{1}{k}}{2(k-\frac{1}{k}+\frac{1}{2})}$ which are equal and higher than $k - 1$. For firm b , again $\pi_{b,2}^\delta(y) = 0$.

We calculate the profit of firm a in both equilibria: $u_a(x) = k - (k - 1) = 1$ and thus

$$\begin{aligned}\gamma(G, \delta) &= u_a^\delta(y) = p_1^\delta(y)y_{a,1} + p_2^\delta(y)y_{a,2} - c_a(y) \\ &= (k + 2\frac{1}{k} - \frac{k + \frac{1}{k}}{2(k - \frac{1}{k} + \frac{1}{2})} - \frac{1}{k}\frac{1}{2})\frac{1}{2} + (k + \frac{1}{k})\frac{1}{2} - (k - 1) \\ &= 1 + \frac{5}{4}\frac{1}{k} - \frac{1}{4}\frac{k + \frac{1}{k}}{k - \frac{1}{k} + \frac{1}{2}} \xrightarrow{k \rightarrow \infty} \frac{3}{4}.\end{aligned}$$

□

4.5 Effect of Price Shocks on Aggregates

Theorem 4.4 shows that an individual firm can lose no more than 25% of its profit as a result of a positive price shock. The lower bound, however, had the property that one firm loses while all competitors gained in their total profits. In this section, we study effects of price shocks on aggregate measures: the welfare and the social surplus.

4.5.1 Welfare

The relative loss of welfare $\gamma^U(G, \delta)$ is by definition a weighted average of the relative profit losses of the individual firms, and, thus, bounded by $\gamma^u(G, \delta)$. Consequently, for any game G , $\gamma^U(G, \delta) \geq \gamma^u(G, \delta) \geq \frac{3}{4}$. We propose instances that converge to this bound, yielding a tight bound.

Theorem 4.12. *Positive price shocks can decrease welfare in multimarket Cournot competition with affine price functions and convex cost by 25%,*

$$\gamma^U(\mathcal{G}, \Delta_{\mathcal{G}}) = \gamma^u(\mathcal{G}, \Delta_{\mathcal{G}}) = \frac{3}{4}.$$

Proof. The upper bound follows from Theorem 4.4. As a lower bound be construct instances with n firms where a positive price shock decreases welfare by a $\frac{(n-1)^2}{4(n^2+n-1)}$ fraction, converging to 25% as n grows.

The lower bound games are similar to those from the proof of Lemma 4.8, except that firm a can produce a quantity of $q_a \leq 1$ at cost 0 and its competitors $i \neq a$ have a per-unit cost of 1, i.e., the cost is $c_i(q_i) = q_i$. In the initial Cournot equilibrium x , no quantity is sold on market 1 and on market 2, $x_{a,2} = 1$ and $x_{i,2} = 0$ for all $i \neq a$. The equilibrium welfare is $U(x) = 1$.

A price shock that increases the price on market 1 to $\delta = \frac{n^2 - 1}{2(n^2 + n - 1)}$ leads firm a to shift to market 1. In the new equilibrium y , $y_{a,1} = \frac{n^2 - n}{2(n^2 + n - 1)}$, $y_{a,2} = \frac{n^2 + 3n - 2}{2(n^2 + n - 1)}$, and $y_{i,2} = \frac{n - 1}{2(n^2 + n - 1)}$ for all $i \neq a$. The new equilibrium welfare is

$$\begin{aligned} U^\delta(y) &= \delta y_{a,1} + (2 - (y_{a,2} + \sum_{i \neq a} y_{i,2})) (y_{a,2} + \sum_{i \neq a} y_{i,2}) - \sum_{i \neq a} y_{i,2} \\ &= \frac{3n^2 + 6n - 5}{4(n^2 + n - 1)} = 1 - \frac{(n - 1)^2}{4(n^2 + n - 1)}. \end{aligned}$$

□

Relationship to the Price Of Anarchy

Known price of anarchy bounds can effortlessly be used to roughly estimate quantitative comparative statics for worst-case equilibrium welfare in maximization games and for worst-case equilibrium social cost in minimization games. This holds for any (possibly non-unique) type of equilibrium as long as the set of parameter changes considered is non-decreasing in a maximization game's welfare or non-increasing in a minimization game's social cost.

For a given equilibrium concept (e.g. mixed Nash equilibria), denote the set of equilibria of a game G by $EQ(G)$. Also, for a maximization game (N, X, u) with welfare $U(x) = \sum_{i \in N} u_i(x)$, denote the welfare optimum by $U_{\max}(G) = \sup_{x \in X} U(x)$.

Lemma 4.13. *A price of anarchy bound $POA(\mathcal{G}, EQ)$ for a class of maximization games \mathcal{G} and an equilibrium concept EQ is also a bound for the maximum welfare loss caused by a parameter change that is monotone in the game's welfare optimum,*

$$\gamma^U(\mathcal{G}, \Delta_{\mathcal{G}}) \geq POA(\mathcal{G}, EQ),$$

where f with $f(G) = \inf_{x \in EQ(G)} U(x)$ is a game's worst-case equilibrium welfare and where $\Delta_{\mathcal{G}}$ is a set of parameter changes with $U_{\max}(G) \leq U_{\max}(G(\delta))$ for all $G \in \mathcal{G}$ and $\delta \in \Delta_{\mathcal{G}}$.

Proof. With our definition of f and U_{\max} , we can express the price of anarchy bound as

$$POA(\mathcal{G}, EQ) = \inf_{G \in \mathcal{G}} \frac{f(G)}{U_{\max}(G)}.$$

Then, for any $G \in \mathcal{G}$ and $\delta \in \Delta_G$,

$$\begin{aligned} \gamma^f(G, \delta) &= \frac{f(G(\delta))}{f(G)} = \frac{f(G(\delta))}{\inf_{x \in EQ(G)} U(x)} \\ &\geq \frac{f(G(\delta))}{U_{\max}(G)} \end{aligned} \tag{4.16}$$

$$\geq \frac{f(G(\delta))}{U_{\max}(G(\delta))} \geq POA(\mathcal{G}, EQ), \tag{4.17}$$

where (4.16) holds by definition of U_{\max} and (4.17) holds because of our assumption on the parameter changes. \square

A similar reduction can be made for minimization games where f is the game's worst-case equilibrium social cost and the parameter changes are non-increasing in the game's social cost optimum. Also, price of stability results imply quantitative comparative statics when f is the best-case equilibrium welfare or best-case equilibrium social cost.

[Johari and Tsitsiklis \(2005\)](#) showed that the price of anarchy for single-market Cournot oligopoly is $\frac{2}{3}$ and that this bound is tight. Their lower bound also holds for our multimarket setting and, hence, implies $\gamma^U(\mathcal{G}, \Delta_{\mathcal{G}}) \geq \frac{2}{3}$, which is as expected an estimation of the tight bound $\gamma^U(\mathcal{G}, \Delta_{\mathcal{G}}) = \frac{3}{4}$ presented in this work.

4.5.2 Aggregate Social Surplus

We now consider the effect of price shocks on the social surplus. The assumption that price functions are affine allows to express the social surplus as

$$S(q) = \sum_{m \in M} \int_0^{q_m} p_m(z) dz - \sum_{i \in N} c_i(q_i) = \sum_{m \in M} \left(p_m(q_m) q_m + \frac{r_m q_m^2}{2} \right) - \sum_{i \in N} c_i(q_i). \quad (4.18)$$

For a given game G , we want to bound the ratio $\gamma^S(G, \delta) = \frac{S^\delta(y)}{S(x)}$.

Theorem 4.14. *Given a game G , a positive price shock δ can decrease the social surplus by at most a factor $\frac{1}{6}$, that is,*

$$\gamma^S(G, \delta) \geq \frac{5}{6}.$$

Before we prove the theorem, we characterize the equilibrium y with a *variational inequality*. Variational inequalities have been used before, e.g., in the context of characterizing equilibria (cf. [Dafermos 1980](#), [Haurie and Marcotte 1985](#)) and for bounding the price of anarchy (cf. [Cominetti et al. 2009](#), [Correa et al. 2004, 2008](#), [Harks and Miller 2011](#), [Harks 2011](#), [Roughgarden 2002](#), [Roughgarden and Schoppmann 2011](#)).

Lemma 4.15. *Let y be the equilibrium for the game with price shock $\delta_m \geq 0, m \in M$ and let x be the original equilibrium. Then, for all $i \in N$ it holds*

$$\sum_{m \in M_i} (p_m^\delta(y) - r_m y_{i,m} - c'_i(y_i)) (x_{i,m} - y_{i,m}) \leq 0. \quad (4.19)$$

Proof. For every firm i , given the equilibrium quantities y_{-i} of its competitors, solving

$$\max_{(q_{i,m})_{m \in M_i}} u_i(q_i, y_{-i})$$

is a convex program. Thus, at an optimal solution $(y_{i,m})_{m \in M_i}$, the gradient $\nabla u_i(y)$ only decreases along any feasible direction. In particular, $(x_{i,m} - y_{i,m})_{m \in M_i}$ is a feasible direction. \square

Before we prove the theorem, we give another technical lemma, estimating the surplus gained at x .

Lemma 4.16. *For a game G , the social surplus in the Cournot equilibrium x is at least*

$$S(x) \geq \sum_{i \in N} \sum_{m \in M} \frac{3}{2} r_m x_{i,m}^2.$$

Proof. For every firm $i \in N$, the cost function c_i is convex, hence,

$$\begin{aligned} \sum_{i \in N} c_i(x) &\leq \sum_{i \in N} c'_i(x)x_i \\ &= \sum_{i \in N} \sum_{m \in M} (p_m(x)x_{i,m} - r_m x_{i,m}^2) \\ &= \sum_{m \in M} p_m(x)x_m - \sum_{i \in N} \sum_{m \in M} r_m x_{i,m}^2 \end{aligned} \quad (4.20)$$

$$(4.21)$$

Here, (4.20) holds because x is an equilibrium for the unperturbed game and, hence, the first order conditions give $c'_i(x) = p_m(x)x_{i,m} - r_m x_{i,m}$ for every market $m \in M$ with $x_{i,m} > 0$. We combine this with the definition of the surplus (4.18),

$$\begin{aligned} S(x) &= \sum_{m \in M} \left(p_m(x)x_m + \frac{r_m x_m^2}{2} \right) - \sum_{i \in N} c_i(x_i) \\ &\stackrel{(4.21)}{\geq} \sum_{m \in M} \left(p_m(x)x_m + \frac{r_m x_m^2}{2} \right) - \sum_{m \in M} p_m(x)x_m + \sum_{i \in N} \sum_{m \in M} r_m x_{i,m}^2 \\ &= \sum_{m \in M} \left(\frac{r_m x_m^2}{2} + \sum_{i \in N} r_m x_{i,m}^2 \right) \geq \sum_{m \in M} \sum_{i \in N} \frac{3}{2} r_m x_{i,m}^2, \end{aligned}$$

as $x_m^2 = (\sum_{i \in N} x_{i,m})^2 \geq \sum_{i \in N} x_{i,m}^2$. □

We now prove the theorem.

Proof of Theorem 4.14. First note that on any market $m \in M$, the price difference between the equilibria is $p_m(x) \leq p_m^\delta(y) + r_m(y_m - x_m)$ and, thus, by rearranging,

$$\begin{aligned} p_m(x)x_m + \frac{r_m x_m^2}{2} &\leq p_m^\delta(y)x_m + r_m(y_m - x_m)x_m + \frac{r_m x_m^2}{2} \\ &= p_m^\delta(y)x_m - \frac{r_m(x_m - y_m)^2}{2} + \frac{r_m y_m^2}{2} \\ &\leq p_m^\delta(y)x_m + \frac{r_m y_m^2}{2}, \end{aligned} \quad (4.22)$$

where (4.22) holds because $(x_m - y_m)^2 \geq 0$. We use this to estimate the difference between the social surplus of the equilibria as in (4.18),

$$\begin{aligned} S(x) - S^\delta(y) &= \sum_{m \in M} \left(p_m(x)x_m - p_m^\delta(y)y_m + \frac{r_m x_m^2}{2} - \frac{r_m y_m^2}{2} \right) + \sum_{i \in N} (c_i(y) - c_i(x)) \\ &\stackrel{(4.22)}{\leq} \sum_{m \in M} p_m^\delta(y)(x_m - y_m) + \sum_{i \in N} c'_i(y)(y_i - x_i), \end{aligned} \quad (4.23)$$

where in (4.23) we also use the fact that $c_i(y) - c_i(x) \leq c'_i(y)(y_i - x_i)$ as the cost functions are convex. The variational inequality (4.19) summed up across all firms is

$$0 \geq \sum_{m \in M} p_m^\delta(y)(x_m - y_m) - \sum_{m \in M} \sum_{i \in N} r_m y_{i,m}(x_{i,m} - y_{i,m}) - \sum_{i \in N} c'_i(y)(x_i - y_i). \quad (4.24)$$

Subtracting (4.24) from (4.23) and estimating $y_{i,m}(x_{i,m} - y_{i,m}) \leq \frac{1}{4}x_{i,m}^2$, we find

$$\begin{aligned} S(x) - S^\delta(y) &\leq \sum_{m \in M} \sum_{i \in N} r_m y_{i,m}(x_{i,m} - y_{i,m}) \leq \sum_{m \in M} \sum_{i \in N} \frac{1}{4} r_m x_{i,m}^2 \\ &\leq \frac{1}{6} S(x), \end{aligned}$$

using the estimation from Lemma 4.16. Rearranging completes the proof of the theorem. \square

We use the instances constructed in the proof of Theorem 4.12 to show that this bound is tight.

Proposition 4.17. *There is a game $G \in \mathcal{G}$ with n firms where a positive price shock δ decreases the equilibrium consumer surplus by a factor $\frac{2n^4+2n^3-5n^2+1}{12(n^2+n-1)^2}$, that is,*

$$\gamma^S(G, \delta) = 1 - \frac{2n^4+2n^3-5n^2+1}{12(n^2+n-1)^2} \xrightarrow{n \rightarrow \infty} \frac{5}{6}.$$

Proof. Note that $S(q) = U(q) + \frac{1}{2} \sum_{m \in M} r_m q_m^2$. We use the instances and equilibria from the proof of Theorem 4.12 and find $S(x) = U(x) + \frac{1}{2} 1^2 = \frac{3}{2}$ and

$$S(y) = U(y) + \frac{1}{2} (y_{a,2} + \sum_{i \neq a} y_{i,2})^2 = \frac{10n^4+22n^3-7n^2-24n+11}{8(n^2+n-1)^2}$$

Combined, this gives

$$\gamma^S(G, \delta) \leq \frac{2}{3} \cdot \frac{10n^4+22n^3-7n^2-24n+11}{8(n^2+n-1)^2} = 1 - \frac{2n^4+2n^3-5n^2+1}{12(n^2+n-1)^2}.$$

\square

4.6 Conclusion

We introduced the quantitative comparative statics approach as a way to apply the worst-case quantitative nature of the price of anarchy to problems that were previously qualitatively analyzed with comparative statics. The goal of our new approach is to provide an understanding of the significance and robustness of the effects of exogenous parameter changes on a system's equilibria.

As a first application, we considered multimarket Cournot competition, where [Bulow et al. \(1985\)](#) showed that a positive price shock can reduce a monopolist's profit. Our results – a positive price shock can reduce profit and welfare by 25% and a negative price shock can increase profit and welfare by 33% – imply that the effect may be significant. For example, the possible 33% increase in profit may be enough for a government to consider imposing a domestic sales tax in order to force domestic companies to compete more aggressively abroad. For a market participant on the other hand, our results imply that equilibrium profit is robust in the sense that at least 75% of current profit is maintained in case of a positive price shock. We further showed that social surplus is more robust against positive price shocks: the worst case loss is bounded by 16.7%.

The structure of our results, especially the bound for the social surplus, hint that the same bounds hold also for mixed and correlated equilibria. In multimarket Cournot competition, such equilibria are not necessarily unique. However, the variational inequality from Lemma 4.15 holds for any such equilibrium and we conjecture that our bounds hold for any two such equilibria – both for two equilibria within the same game as well as for equilibria before and after a positive price shock.

Another natural extension of our analysis would be the generalization of the Cournot market model to aggregative games with strategic substitutes as considered by [Acemoglu and Jensen \(2013\)](#). Additionally, a quantification for a qualitatively similar paradoxical effect for markets with strategic complements and joint economies of scale remains open.

Chapter 5

Convergence of Regret-Minimizing Algorithms in Atomic Splittable Routing

In the previous chapters, we assumed the perspective of a game designer who is concerned with the efficiency of pure Nash equilibria. However, it is not a priori clear how, or even if, players of a game will reach a pure Nash equilibrium. Indeed, computing a Nash equilibrium is intractable (formally, PPAD-complete) in general games (Chen et al. 2009, Daskalakis et al. 2009). This issue suggests two ways to make Nash equilibrium analyses more meaningful for a class of games. The first is to propose a set of learning algorithms – procedures that players use to select actions over time – and prove that the collective behavior of players that use such learning algorithms converges to a Nash equilibrium over time. The second is to extend conclusions about Nash equilibrium to weaker equilibrium concepts that are more computationally tractable.

In this chapter, we analyze the collective behavior of network users that employ no-regret algorithms. When a player uses a no-regret algorithm in a repeated game, the player’s average (over time steps) regret, i.e., the difference between her average cost and the average cost of the best single action in hindsight, converges to 0. Several computationally efficient no-regret algorithms are known (Kalai and Vempala 2005, Zinkevich 2003), even for the ‘bandit’ setting where only the costs of edges used (rather than the costs of all edges) are revealed to a player after a given time step (Awerbuch and Kleinberg 2008).

For the analysis of this chapter, we choose atomic splittable routing games and their natural extension, atomic splittable congestion games. In such routing games, each self-interested end user routes a prescribed amount of traffic from its source to its sink in a network with the goal to minimize her own cost. Atomic games, as studied here, have a finite set of users, each of whom controls a non-negligible amount of flow. We further assume that a user can split her traffic over multiple paths – every fractional flow that sends the correct amount of traffic is a viable strategy. Such atomic splittable routing games naturally model, among other things, multipath routing in communication networks and fleet scheduling in transportation networks. Atomic splittable *congestion* games are a superclass of atomic splittable routing games, where players split their demand across predefined subsets of resources with arbitrary (non-network) combinatorial structure.

Much is now known about pure Nash equilibria of atomic splittable routing games. Such equilibria always exist (Orda et al. 1993), and several necessary and sufficient conditions for uniqueness are known (Bhaskar et al. 2009, Neyman 1997, Richman and Shimkin 2007). The price of anarchy — the worst-case overall cost of Nash equilibria, relative to that of an optimal solution — was studied in Correa et al. (2005), Cominetti et al. (2009), Harks (2011), Roughgarden (2005a), Roughgarden and Schoppmann (2011). See Roughgarden (2005b, §4.8) for further references on splittable routing games.

The purpose of the introduction of no-regret learning to this line of research is to establish a broader foundation. In general games, even with only two players, the behavior of players that use no-regret algorithms need not converge to a Nash equilibrium. Are positive results possible for the special case of atomic splittable routing games? In the best-case scenario, the use of no-regret algorithms would guarantee the quick convergence to a Nash equilibrium. In the second-best scenario, laudable properties of Nash equilibria – like having small cost, relative to an optimal flow – would hold more generally for sequences of outcomes generated by no-regret learners (even if a Nash equilibrium is not reached).

The only known result for no-regret learning in atomic splittable routing is by Even-Dar et al. (2009), who have shown the convergence of no-regret learners for socially concave games, which include atomic splittable routing restricted to affine latency functions. We improve and extend their results for this setting in Section 5.3. Additionally, several previous works have studied the convergence of learning dynamics in other types of routing games. Most closely related to our work is the paper by Blum et al. (2006), who studied analogous questions in nonatomic selfish routing networks, where there is a continuum of players of negligible size. Atomic splittable routing networks are technically more complicated than nonatomic ones. One difference relevant for learning dynamics is that equilibria are essentially unique in all nonatomic routing networks, but not in all atomic splittable routing networks (Bhaskar et al. 2009). Another difference is that atomic splittable routing games are only ‘locally smooth’ games in the sense of Roughgarden and Schoppmann (2011), while nonatomic networks are ‘smooth’ games in the sense of Roughgarden (2009). For this reason, it is non-trivial to prove bounds on the average cost of no-regret sequences in atomic splittable networks.

Other works on convergence in nonatomic selfish routing networks, which focus on specific dynamics that possess good properties, include Fischer and Vöcking (2009) and Fischer et al. (2010). A similar style of result, for atomic unsplittable networks, is Kleinberg et al. (2009).

As far as we know, the only previous work on learning dynamics in atomic splittable routing games with non-affine latency functions concern best-response dynamics, which seem more ill-behaved – or at least much more difficult to analyze – than no-regret dynamics (Mertzios 2008, Orda et al. 1993).¹ An additional disadvantage of best-response dynamics is the assumption that every player is aware of the current actions taken by all other players – with no-regret dynamics, each player merely keeps track of observed costs, and does not need to know anything about other players.

¹Best-response dynamics have been exhaustively studied in atomic unsplittable networks and more generally in potential games (see e.g. Tardos and Wexler 2007); atomic splittable games are not potential games, however.

5.1 Contribution and Outline

We first state our results for atomic splittable selfish routing networks with affine latency functions, which are stronger than what is possible for networks with general latency functions. These games are known to belong to the class of *socially concave games*, in which no-regret play converges to a Nash equilibrium (Even-Dar et al. 2009). As our first result, we improve the known convergence bound, proving that if the aggregate regret across all players is at most ε , then the time-averaged flow is an ε -approximate Nash equilibrium (the total benefit of players' best responses is at most ε). Second, we show that this implies that, on all but an ε' fraction of time steps, the flow at a given time step is an ε' -approximate Nash equilibrium (where ε' depends polynomially on ε , the network size, and the maximum slope of a latency function). Third, we prove that the average cost of every sequence with aggregate regret ε is at most $\frac{3}{2}$ times that of an optimal flow, plus an error term $f(\varepsilon)$ that tends to 0 as ε goes to 0. We prove this by showing that the average cost of such a sequence is not much more than the cost of the corresponding time-averaged flow, and then extending the “local smoothness” bounds in Roughgarden and Schoppmann (2011) from exact to approximate Nash equilibria. Finally, we show that our analysis yields a new “essential uniqueness” result for the coarse correlated equilibria (i.e., limits of no-regret sequences) of atomic splittable routing games with affine latency functions.

Atomic splittable networks with general latency functions exhibit fundamentally different behavior. One clue is that networks with affine latency functions have a convex potential function and unique pure Nash (even correlated) equilibria (Neyman 1997), while networks with general latency functions do not (Bhaskar et al. 2009). More concretely, a key issue in the analysis of no-regret sequences in selfish routing networks is to control the difference between the average cost (over time) of an edge and the cost of the average (over time) flow value on that edge. These two quantities are guaranteed to be identical if and only if the latency function is affine.

We provide approximate convergence results in networks with general latency functions that are close to the best possible. Precisely, we give two examples that delineate what we can hope for: first, a game with cubic cost functions with a coarse correlated equilibrium that costs 5.29 times as much as an optimal flow, showing that price of anarchy bounds obtained through local smoothness (Roughgarden and Schoppmann 2011) do not apply to coarse correlated equilibria. Second, a game where even when $\varepsilon = 0$, the cost of a no-regret sequence can exceed that of all Nash equilibria by an error term that depends linearly on the network size k and the maximum slope a . On the positive side, we prove that in networks with latency functions drawn from a set \mathcal{L} , every no-regret sequence with aggregate regret ε has average cost at most $\rho(\mathcal{L})$ times that of an optimal flow – where $\rho(\mathcal{L})$ is the worst-case price of anarchy of pure Nash equilibria in such networks – plus an error term $f(\varepsilon) + \mathcal{O}(ak)$, where $f(\varepsilon)$ tends to 0 with ε .

We present our results in this chapter for routing games, because their network structure allows for stronger bounds. The results can, however, also be stated for atomic splittable congestion games by replacing the factors m and k in the bounds (representing the number of edges and nodes in the graph of the routing game, respectively) by the number m of resources of the congestion game.

Bibliographic Information. The results presented in this chapter are joint work with Tim Roughgarden.

5.2 Atomic Splittable Routing and No-regret Play

The context of this chapter are atomic splittable routing games. In these games, a set of players $i \in N$ routes flow through a graph (V, E) , each from a source node $s_i \in V$ to a terminal node $t_i \in V$. As usual, $n = |N|$ denotes the number of players, while $m = |E|$ and $k = |V|$ denote the number of edges and vertices of the graph, respectively. The players individually minimize the latency they incur on the edges used by their flow. Each player routes $\frac{1}{n}$ units of flow, such that the total flow is normalized to 1 unit. We denote the set of s_i, t_i paths by \mathcal{P}_i and a player i may split her flow across multiple paths $P \in \mathcal{P}_i$, such that any fractional allocation $x_i = (x_{i,P})_{P \in \mathcal{P}_i}$ of $\frac{1}{n}$ to the s_i, t_i paths $P \in \mathcal{P}_i$ is a feasible pure strategy. Given a strategy x_i for every player, we denote the joint strategy $x = (x_i)_{i \in N}$. The flow $x_{i,e}$ of player i on edge e is the sum of all $x_{i,P}$ for the paths $P \in \mathcal{P}_i$ with $e \in P$ and the total flow on edge e is $x_e = \sum_{i \in N} x_{i,e}$.

Every edge e has latency as a function ℓ_e of x_e . In Section 5.3 we deal with affine latency functions of the form $\ell_e(x_e) = a_e x_e + b_e$ and in Section 5.4 with general latency functions that are non-decreasing, continuously differentiable, semi-convex (i.e. $\ell_e(x_e)x_e$ is convex in x_e). In both sections we assume that the latency functions' slopes are bounded, that is, there is some $a \in \mathbb{R}$ such that on all edges $|\ell_e(x_e) - \ell_e(y_e)| \leq a|x_e - y_e|$ for all x_e, y_e . Given a flow x , the cost of player i is $c_i(x) = \sum_{e \in E} \ell_e(x_e)x_{i,e}$ and the social cost is $C(x) = \sum_{i \in N} c_i(x)$. We give a formal definition that embeds atomic splittable routing games into the strategic game framework.

Definition 5.1 (atomic splittable routing game)

An atomic splittable routing game is a strategic minimization game that is given as a tuple $(N, (V, E), (s_i, t_i)_{i \in N}, (\ell_e)_{e \in E})$. N is the set of players, the strategy set of a player $i \in N$ is $X_i = \{(x_{i,P})_{P \in \mathcal{P}_i} : \sum_{P \in \mathcal{P}_i} x_{i,P} = \frac{1}{n}\}$ and her cost function is $c_i(x) = \sum_{e \in E} \ell_e(x_e)x_{i,e}$.

No-Regret Play and Price of Total Anarchy. We study repeated play of an atomic splittable routing game and ask the question if and how efficiently players will reach a pure Nash equilibrium. We assume that all players employ no-regret algorithms (see Definition 2.7) to choose their strategy, generating a sequence x^1, x^2, \dots of strategy profiles such that at time step T each player i has bounded regret at most $\mathcal{R}_i(T)$ with $\mathcal{R}_i(T) = o(T)$. We denote for $T \in \mathbb{N}$ by $\varepsilon^T = \frac{1}{T} \sum_{i \in N} \mathcal{R}_i(T)$ a bound on the time-average overall regret. From $\mathcal{R}_i(T) = o(T)$ it follows that ε^T converges to 0 as T grows large. We further denote by $\hat{x}^T = \frac{1}{T} \sum_{t=1}^T x^t$ the average strategy profile, which we will find to be a good approximation of the actual flows in the sequence. For ease of exposition, we will sometimes consider a fixed T , write $(x^t)_{t=1}^T$ and set $\varepsilon = \varepsilon^T$ and $\hat{x} = \hat{x}^T$.

No-regret sequences are dual to coarse correlated equilibria as described in Chapter 2. To measure the efficiency of such sequences, Blum et al. (2008) proposed the *price of total anarchy* that extends the price of anarchy from equilibria to no-regret sequences. That is, it relates the average cost of the most expensive no-regret sequence to a social optimum and converges to the price of anarchy for coarse correlated equilibria as the regret goes to zero.

5.3 Affine Latency: Learning Converges

We begin by studying atomic splittable routing games with affine latency functions. We treat this important special case separately because it exhibits qualitatively different behavior than the general case, and because it allows stronger results. In addition, the analysis of this section acts as a warm-up for the more complicated analysis in Section 5.4.

5.3.1 Convergence to an Equilibrium

[Even-Dar et al. \(2009\)](#) have shown that atomic splittable routing games with affine latency are *socially concave games*² for which the average strategy of a no-regret sequence is an approximate equilibrium.

Theorem 5.2 (Even-Dar et al. 2009). *Consider repeated play of an instance of atomic splittable routing with affine latency. Given a no-regret sequence $(x^t)_{t=1}^T$ with average flow \hat{x} and time-average regret at most ε , no player $i \in N$ can improve by more than ε through a unilateral deviation from \hat{x} , i.e., \hat{x} is a strict ε -approximate Nash equilibrium.*

As discussed in Section 2.1.2, this implies that the maximum unilateral improvement of all players combined is $n\varepsilon$, i.e., that the average flow \hat{x} is a $(n\varepsilon)$ -approximate Nash flow. We find in the upcoming lemma that \hat{x} is in fact an ε -approximate Nash flow. This not only improves the previously known bound, but will also allow us to make statements about the cost of no-regret sequences.

Theorem 5.3. *Consider repeated play of an instance of atomic splittable routing with affine latency. The average flow \hat{x} of a no-regret sequence $(x^t)_{t=1}^T$ is an ε -approximate Nash flow.*

Proof. When ℓ_e is affine, $\ell_e(x_e)x_e$ is a convex function of x_e . Hence, $\ell_e(\hat{x}_e)\hat{x}_e \leq \frac{1}{T} \sum_{t=1}^T \ell_e(x_e^t)x_e^t$, and we have

$$\sum_{i \in N} c_i(\hat{x}) = \sum_{e \in E} \ell_e(\hat{x}_e)\hat{x}_e \leq \sum_{e \in E} \frac{1}{T} \sum_{t=1}^T \ell_e(x_e^t)x_e^t = \frac{1}{T} \sum_{t=1}^T \sum_{i \in N} c_i(x^t).$$

We now use that the time-average regret of the sequence $(x^t)_{t=1}^T$ is bounded by ε .

$$\begin{aligned} \sum_{i \in N} c_i(\hat{x}) &\leq \sum_{i \in N} \frac{1}{T} \sum_{t=1}^T c_i(x^t) \leq \varepsilon + \sum_{i \in N} \min_{y_i} \frac{1}{T} \sum_{t=1}^T c_i(y_i, x_{-i}^t) \\ &= \varepsilon + \sum_{i \in N} \min_{y_i} c_i(y_i, \hat{x}_{-i}), \end{aligned} \tag{5.1}$$

where in the last step we use that for affine latency $\frac{1}{T} \sum_{t=1}^T \ell_e(x_{-i,e}^t + y_{i,e}) = \ell_e(\hat{x}_{-i,e} + y_{i,e})$. \square

²Socially concave games are defined by two characteristics. First, the utility function of each agent is a convex in the actions of all other agents. Second, there exists a strict convex combination of the utility functions which is a concave function. See [Even-Dar et al. \(2009, §2\)](#) for an exact definition.

For linear latency functions, we give an extension of the above theorem, showing that not only the average flow but also ‘most’ of the flows in the sequence are approximate Nash flows.

Lemma 5.4. *Consider repeated play of an instance of atomic splittable routing with linear latency. Given a no-regret sequence $(x^t)_{t=1}^T$, for all but an $\sqrt[4]{am^3\epsilon^T}$ fraction of time steps up to T , x^t is an $(\epsilon^T + 3\sqrt[4]{am^3\epsilon^T} + \sqrt{am\epsilon^T})$ -approximate Nash flow, where n is the number of players, m is the number of edges in the graph and the latencies have maximum slope a . In other words, for all but an ρ^T fraction of time steps up to T , x^t is an ρ^T -approximate Nash flow for $\rho^T = \Theta(\sqrt[4]{\epsilon^T})$.*

We prove Lemma 5.4 in two steps. First, we express the average difference between the latency on an edge and its latency in the average flow as the *jitter* of the edge. We quantify this jitter, showing that the flows of a no-regret sequence approach their average as T grows large. This implies that on a given edge e ‘on most days’ t (i.e., in most time steps t in the sequence), the flow x_e^t is almost similar to the average flow \hat{x}_e , which will be key to the proof of Lemma 5.4.

Lemma 5.5. *Consider repeated play of an instance of atomic splittable routing with affine latency. Given a no-regret sequence $(x^t)_{t=1}^T$ with average flow \hat{x} and time-average regret at most ϵ , the jitter of latency on the edges is bounded. That is, $\frac{1}{T} \sum_{t=1}^T |\ell_e(x_e^t) - \ell_e(\hat{x}_e)| \leq \sqrt{a\epsilon_e}$ with $\sum_{e \in E} \epsilon_e \leq \epsilon$.*

Proof. Our maximum slope assumption implies that $|\ell_e(x_e^t) - \ell_e(\hat{x}_e)| \leq a|x_e^t - \hat{x}_e|$. Then, on every edge e ,

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T |\ell_e(x_e^t) - \ell_e(\hat{x}_e)| &\leq \frac{1}{T} \sum_{t=1}^T \sqrt{a(x_e^t - \hat{x}_e)(\ell_e(x_e^t) - \ell_e(\hat{x}_e))} \\ &\leq \sqrt{a \frac{1}{T} \sum_{t=1}^T (x_e^t - \hat{x}_e) \ell_e(x_e^t) - a \underbrace{\frac{1}{T} \sum_{t=1}^T (x_e^t - \hat{x}_e) \ell_e(\hat{x}_e)}_{=0}} \quad (5.2) \\ &= \sqrt{\frac{1}{T} \sum_{t=1}^T a(x_e^t - \hat{x}_e) \ell_e(x_e^t)} = \sqrt{a\epsilon_e}, \end{aligned}$$

where we use the Cauchy-Schwarz inequality to bring the summation into the root and define $\epsilon_e = \frac{1}{T} \sum_t \ell_e(x_e^t)(x_e^t - \hat{x}_e)$. It remains to bound ϵ_e . We use that the time-average regret of the

sequence $(x^t)_{t=1}^T$ is bounded by ε ,

$$\begin{aligned} \sum_{e \in E} \varepsilon_e &= \sum_{i \in N} \sum_{e \in E} \frac{1}{T} \sum_{t=1}^T \left(\ell_e(x_e^t) x_{i,e}^t - \ell_e(x_e^t) \hat{x}_{i,e} \right) \\ &\leq \varepsilon + \sum_{i \in N} \sum_{e \in E} \frac{1}{T} \sum_{t=1}^T \left(\ell_e(x_{-i,e}^t + \hat{x}_{i,e}) \hat{x}_{i,e} - \ell_e(x_e^t) \hat{x}_{i,e} \right) \\ &= \varepsilon + \sum_{i \in N} \sum_{e \in E} \underbrace{\left(\hat{x}_{i,e} \frac{1}{T} \sum_{t=1}^T (\ell_e(x_{-i,e}^t + \hat{x}_{i,e}) - \ell_e(x_e^t)) \right)}_{=0} = \varepsilon, \end{aligned} \quad (5.3)$$

where in the last step we use that the latency functions are affine. \square

Proof of Lemma 5.4. Note that for linear latencies, $c_i(x) = \sum_{e \in E} a_e x_e x_{i,e}$ and $C(x) = \sum_{e \in E} a_e x_e^2$. From (5.2) we follow that for a given edge e on all but at most a $\sqrt[4]{a\varepsilon_e}$ fraction of time steps, $|\ell_e(\hat{x}_e) - \ell_e(x_e^t)| \leq \sqrt[4]{a\varepsilon_e}$. Then, on all but a $\sum_{e \in E} \sqrt[4]{a\varepsilon_e} \leq \sqrt[4]{a\varepsilon m^3}$ fraction of time steps for all edges e simultaneously $a_e |\hat{x}_e - x_e^t| = |\ell_e(\hat{x}_e) - \ell_e(x_e^t)| \leq \sqrt[4]{a\varepsilon_e}$. Fix such a time step t and a flow y and denote quantities $\alpha = \sum_{i \in N} (c_i(x^t) - c_i(\hat{x}))$ and $\beta = \sum_{i \in N} (c_i(y_i, \hat{x}_{-i}) - c_i(y_i, x_{-i}^t))$ such that

$$\begin{aligned} \sum_{i \in N} c_i(x^t) &= \sum_{i \in N} c_i(\hat{x}) + \alpha \\ &\leq \sum_{i \in N} c_i(y_i, \hat{x}_{-i}) + \varepsilon + \alpha \\ &= \sum_{i \in N} c_i(y_i, x_{-i}^t) + \varepsilon + \alpha + \beta, \end{aligned}$$

where we use that \hat{x} is an ε -approximate Nash flow. To bound α , we find that on any edge $e \in E$,

$$\begin{aligned} \sum_{i \in N} a_e (x_e^t x_{i,e}^t - \hat{x}_e \hat{x}_{i,e}) &= a_e ((x_e^t)^2 - \hat{x}_e^2) \\ &= 2\hat{x}_e a_e (x_e^t - \hat{x}_e) + a_e (x_e^t - \hat{x}_e)^2 \\ &\leq 2\sqrt[4]{a\varepsilon_e} + \sqrt{a\varepsilon_e}, \end{aligned} \quad (5.4)$$

where (5.4) uses our assumption $a_e |\hat{x}_e - x_e^t| \leq \sqrt[4]{a\varepsilon_e}$ and that the total flow is scaled to one unit. Then,

$$\begin{aligned} \alpha &= \sum_{i \in N} (c_i(x^t) - c_i(\hat{x})) \\ &\leq \sum_{e \in E} (2\sqrt[4]{a\varepsilon_e} + \sqrt{a\varepsilon_e}) \end{aligned} \quad (5.5)$$

$$\begin{aligned} &\leq 2\sqrt[4]{am^3 \sum_{e \in E} \varepsilon_e} + \sqrt{am \sum_{e \in E} \varepsilon_e} \\ &\leq 2\sqrt[4]{am^3 \varepsilon} + \sqrt{am\varepsilon} \end{aligned} \quad (5.6)$$

Here, (5.5) follows from (5.4) and in (5.6) we use the Cauchy-Schwarz inequality to bring the summation into the roots before we apply the jitter bound (5.3).

For β , we use that no player can have more than $\frac{1}{n}$ flow on a given edge, i.e., $y_{i,e} \leq \frac{1}{n}$, and that

$$\sum_{i \in N} (\hat{x}_{-i,e} - x_{-i,e}^t) = (n-1)(\hat{x}_e - x_e^t). \quad (5.7)$$

Then,

$$\begin{aligned} \beta &= \sum_{i \in N} c_i(y_i, \hat{x}_{-i}) - c_i(y_i, x_{-i}^t) \\ &= \sum_{i \in N} \sum_{e \in E} a_e (\hat{x}_{-i,e} + y_{i,e} - x_{-i,e}^t - y_{i,e}) y_{i,e} \\ &\stackrel{(5.7)}{\leq} \sum_{e \in E} (n-1) a_e (\hat{x}_e - x_e^t) \frac{1}{n} \\ &\leq \frac{n-1}{n} \sum_{e \in E} \sqrt[4]{a \varepsilon_e} \leq \frac{n-1}{n} \sqrt[4]{am^3 \sum_{e \in E} \varepsilon_e} \end{aligned} \quad (5.8)$$

$$\leq \frac{n-1}{n} \sqrt[4]{am^3 \varepsilon} \quad (5.9)$$

In (5.8) we apply the assumption $a_e |\hat{x}_e - x_e^t| \leq \sqrt[4]{a \varepsilon_e}$ and use the Cauchy-Schwarz inequality to bring the summation into the root. Then, in (5.9), we again apply the jitter bound.

Altogether, $\sum_{i \in N} c_i(x^t) \leq \sum_{i \in N} c_i(y_i, x_{-i}^t) + \varepsilon + 3\sqrt[4]{am^3 \varepsilon} + \sqrt{am\varepsilon}$. \square

5.3.2 Price of Total Anarchy

We establish a price of total anarchy statement via the fact that a no-regret sequence's average flow \hat{x} is an ε -approximate Nash flow. First we show how the average cost of the sequence relates to the cost of \hat{x} , then we extend a known price of anarchy result for pure, mixed and correlated equilibria to ε -approximate Nash flows.

Lemma 5.6. *Consider repeated play of an instance of atomic splittable routing with affine latency. The average cost of a no-regret sequence $(x^t)_{t=1}^T$ is at most ε more than the cost of its average flow \hat{x} , i.e., $\frac{1}{T} \sum_{t=1}^T C(x^t) \leq \varepsilon + C(\hat{x})$.*

Proof. For sequences in games with affine latency, $\frac{1}{T} \sum_{t=1}^T \ell_e(x_{-i,e}^t + \hat{x}_{i,e}) = \ell_e(\hat{x}_e)$ and hence also $\frac{1}{T} \sum_{t=1}^T c_i(\hat{x}_i, x_{-i}^t) = c_i(\hat{x})$. Then, using the regret bound,

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T C(x^t) &= \sum_{i \in N} \frac{1}{T} \sum_{t=1}^T c_i(x^t) \leq \varepsilon + \sum_{i \in N} \frac{1}{T} \sum_{t=1}^T c_i(\hat{x}_i, x_{-i}^t) \\ &= \varepsilon + \sum_{i \in N} c_i(\hat{x}) = \varepsilon + C(\hat{x}). \end{aligned}$$

\square

For a price of total anarchy statement it remains to relate the cost of an ε -approximate Nash equilibrium to the cost of a socially optimal flow. We extend on a theorem for pure, mixed and correlated Nash equilibria established by Roughgarden and Schoppmann (2011) through local smoothness.

Definition 5.7 (locally smooth games)

A cost-minimization game (N, X, c) is locally (λ, μ) -smooth with respect to the outcome y if for every outcome x ,

$$C(x) - \sum_{i \in N} \sum_{e \in E} c'_{i,e}(x)(x_{i,e} - y_{i,e}) \leq \lambda \cdot C(y) + \mu \cdot C(x)$$

Here, $c'_{i,e}(x)$ is the marginal cost of player i on e with respect to $x_{i,e}$. That is, if a_e is the slope of ℓ_e ,

$$c'_{i,e}(x) = \frac{\partial \ell_e(x_{-i,e} + x_{i,e})}{\partial x_{i,e}} \leq \ell_e(x_e) + a_e x_{i,e}, \quad (5.10)$$

which holds for affine latency functions with equality.

Theorem 5.8 (Roughgarden and Schoppmann 2011). *Atomic splittable routing with affine latency is locally $(1, \frac{1}{3})$ -smooth with respect to optimal flows y . That is, for all flows x ,*

$$C(x) \leq \frac{3}{2}C(y) + \frac{3}{2} \sum_{i \in N} \sum_{e \in E} c'_{i,e}(x)(x_{i,e} - y_{i,e}).$$

We now extend this theorem to approximate Nash flows.

Lemma 5.9. *Consider an instance of atomic splittable routing with n players, k nodes and affine latency with maximum slope a . Given an ε -approximate Nash flow x and a socially optimal flow y , $C(x) \leq \frac{3}{2} \cdot C(y) + 3\sqrt{2\frac{ak\varepsilon}{n}} + 3\varepsilon$.*

Proof. Let x be an ε -approximate Nash flow and y socially optimal. We reduce the proof of the lemma to Theorem 5.8 by showing that for each player i the value of the gradient $\gamma_i = \sum_{e \in E} c'_{i,e}(x)(x_{i,e} - y_{i,e})$ is bounded. As shown below, $\gamma_i > 0$ if and only if there is some $\delta_i \in (0, 1]$ such that $x^{\delta_i} = (x_i + \delta_i(y_i - x_i), x_{-i})$ is an improving move, i.e. $c_i(x) > c_i(x^{\delta_i})$. The cost of player i in x^{δ_i} is

$$\begin{aligned} c_i(x^{\delta_i}) &= \sum_{e \in E} \ell_e(x^{\delta_i}) x_{i,e}^{\delta_i} = \sum_{e \in E} (\ell_e(x_e) + a_e \delta_i(y_{i,e} - x_{i,e})) (x_{i,e} + \delta_i(y_{i,e} - x_{i,e})) \\ &= \sum_{e \in E} (\ell_e(x_e) x_{i,e} + (\ell_e(x_e) + a_e x_{i,e}) \delta_i(y_{i,e} - x_{i,e}) + a_e \delta_i^2 (y_{i,e} - x_{i,e})^2) \\ &= c_i(x) + \sum_{e \in E} (-\delta_i c'_{i,e}(x)(x_{i,e} - y_{i,e}) + a_e \delta_i^2 (y_{i,e} - x_{i,e})^2). \end{aligned} \quad (5.11)$$

Hence, the cost difference of x and x^{δ_i} is for player i

$$c_i(x) - c_i(x^{\delta_i}) = \delta_i \gamma_i - \delta_i^2 \sum_{e \in E} a_e (y_{i,e} - x_{i,e})^2. \quad (5.12)$$

It follows that if $\gamma_i > 0$, there is a $\delta_i \in (0, 1]$ such that $c_i(x) - c_i(x^{\delta_i}) > 0$, and we choose the δ_i such that x^{δ_i} has the least cost:

- Case 1: $c_i(x^{\delta_i})$ minimized by $\delta_i \in (0, 1)$ such that $\frac{\partial}{\partial \delta_i} c_i(x^{\delta_i}) = 0$
- Case 2: $c_i(x^{\delta_i})$ minimized by $\delta_i = 1$ and $\frac{\partial}{\partial \delta_i} c_i(x^{\delta_i}) \leq 0$

Denote by N_1 and N_2 the sets of players for which Case 1 or Case 2 apply, respectively. For these players let $\varepsilon_i = c_i(x) - c_i(x^{\delta_i})$ and for those with $\gamma_i \leq 0$ let $\varepsilon_i = 0$. Then, $\sum_i \varepsilon_i \leq \varepsilon$ because the flow x is an ε -approximate Nash equilibrium. For the derivative we follow from (5.12),

$$\frac{\partial}{\partial \delta_i} c_i(x^{\delta_i}) = -\gamma_i + 2\delta_i \sum_{e \in E} a_e (y_{i,e} - x_{i,e})^2. \quad (5.13)$$

Then, we have for $i \in N_1$,

$$\varepsilon_i \stackrel{(5.12)}{=} \delta_i \gamma_i - \delta_i^2 \sum_{e \in E} a_e (y_{i,e} - x_{i,e})^2 \stackrel{(5.13)=0}{=} \delta_i^2 \sum_{e \in E} a_e (y_{i,e} - x_{i,e})^2, \quad (5.14)$$

Using (5.13) = 0 and estimating $\sum_{e \in E} a_e (y_{i,e} - x_{i,e})^2 \leq 2\frac{ak}{n^2}$, we have for $i \in N_1$,

$$\gamma_i = 2\delta_i \sum_{e \in E} a_e (y_{i,e} - x_{i,e})^2 \stackrel{(5.14)}{=} 2\sqrt{\varepsilon_i} \cdot \sqrt{\sum_{e \in E} a_e (y_{i,e} - x_{i,e})^2} \leq 2\sqrt{\varepsilon_i} \sqrt{2\frac{ak}{n^2}}.$$

For $i \in N_2$ we have $\delta_i = 1$ and can rearrange (5.12) to

$$\gamma_i = \varepsilon_i + \sum_{e \in E} a_e (y_{i,e} - x_{i,e})^2.$$

Also, with $\frac{\partial}{\partial \delta_i} c_i(x^{\delta_i}) \leq 0$ we derive from (5.13) that

$$\sum_{e \in E} a_e (y_{i,e} - x_{i,e})^2 \leq \frac{1}{2} \gamma_i.$$

Consequently, $\gamma_i \leq 2\varepsilon_i$ for $i \in N_2$. Combining both cases and using the Cauchy-Schwarz inequality to move the summation inside the root,

$$\begin{aligned} \sum_{i \in N} \gamma_i &\leq \sum_{i \in N_1} 2\sqrt{2\frac{ak}{n^2} \varepsilon_i} + \sum_{i \in N_2} 2\varepsilon_i \\ &\leq 2\sqrt{2\frac{ak}{n} \sum_{i \in N_1} \varepsilon_i} + 2\varepsilon \\ &\leq 2\sqrt{2\frac{ak\varepsilon}{n}} + 2\varepsilon. \end{aligned}$$

gives the desired result. \square

To show that the above lemma is tight, we propose a routing game with an ε -approximate Nash equilibrium with cost that matches the bound given in Lemma 5.9

Lemma 5.10. *For any values a, k, n , there are instances of atomic splittable routing with an ε -approximate Nash equilibrium that costs $\mathcal{O}(\sqrt{\frac{ak}{n}}\varepsilon + \varepsilon)$ more than the Nash equilibrium.*

Proof. Consider the following game: there are two parallel edges e_1 and e_2 and two players. The latency on e_1 is $\ell_{e_1}(x_{e_1}) = x_{e_1}$, i.e. linear in x_{e_1} , and the latency on e_2 is $\ell_{e_2}(x_{e_2}) = 1$, i.e. constant. Then, the unique Nash flow x^N is $x_{i,e_1}^N = \frac{1}{3}$ and $x_{i,e_2}^N = \frac{1}{6}$ for $i = 1, 2$. The flow x^ε with $x_{i,e_1}^\varepsilon = \frac{1}{3} + \frac{\sqrt{2\varepsilon}}{3}$ and $x_{i,e_2}^\varepsilon = \frac{1}{6} - \frac{\sqrt{2\varepsilon}}{3}$ for $i = 1, 2$ is an ε -approximate Nash flow, as a simple calculation shows $c_1(x^\varepsilon) + c_2(x^\varepsilon) = \varepsilon + c_1(y_1, x_2^\varepsilon) + c_2(y_2, x_1^\varepsilon)$ for any flows y_1 and y_2 . The cost difference between x^ε and x^N is $C(x^\varepsilon) - C(x^N) = \frac{1}{9}(2\sqrt{2\varepsilon} + 8\varepsilon) \in \mathcal{O}(\sqrt{\varepsilon})$. This can be generalized for slope a, n players and paths of k edges, which gives $\mathcal{O}(\sqrt{\frac{ak}{n}}\varepsilon + \varepsilon)$. \square

Corollary 5.11. *The price of total anarchy for atomic splittable routing with affine latency is $\frac{3}{2}$. Moreover, for an instance with affine latency with maximum slope a , k nodes, and n players, a no-regret sequence $(x^t)_{t=1}^T$ and an optimal flow y ,*

$$\frac{1}{T} \sum_{t=1}^T C(x^t) \leq \frac{3}{2} C(y) + 3\sqrt{2\frac{ak}{n}\varepsilon} + 4\varepsilon.$$

5.3.3 Uniqueness of Coarse Correlated Equilibria

Every instance of atomic splittable routing with affine latency has a unique pure, mixed and correlated Nash equilibrium (Neyman 1997). We find that coarse correlated equilibria (the convergence limit of no-regret sequences) are *essentially unique* in the sense that on every edge the flow is almost surely equal to the unique pure Nash flow. The flow on a given edge can, however, be split arbitrarily among the players, as long as it is on average the players' pure Nash flow.

Theorem 5.12. *Consider an instance of atomic splittable routing with affine latency, a probability distribution σ over the strategy space X and a flow x distributed according to σ with expected flow $\bar{x} = \mathbb{E}_{x \sim \sigma}[x]$. The distribution σ is a coarse correlated equilibrium if and only if \bar{x} is a pure Nash equilibrium and on every edge e the flow x_e is almost surely equal to \bar{x}_e .*

Proof. Let x be a flow distributed according to a coarse correlated equilibrium and $\bar{x} = \mathbb{E}_{x \sim \sigma}[x]$. We start the forward direction of the proof by showing that on every edge e , $x_e \stackrel{\text{a.s.}}{=} \bar{x}_e$. For a coarse correlated equilibrium by definition,

$$\begin{aligned} \sum_{i \in N} \mathbb{E}_{x \sim \sigma}[c_i(x)] &\leq \sum_{i \in N} \mathbb{E}_{x \sim \sigma}[c_i(\bar{x}_i, x_{-i})] \\ &= \sum_{i \in N} \sum_{e \in E} \mathbb{E}_{x \sim \sigma}[\ell_e(\bar{x}_{i,e} + x_{-i,e})] \bar{x}_{i,e} \\ &= \sum_{i \in N} c_i(\bar{x}), \end{aligned}$$

where we use in the last step that the latency functions are affine.

Summing up across edges instead of summing up across players, this is equivalent to

$$\sum_{e \in E} \mathbb{E}_{x \sim \sigma} [\ell_e(x_e) x_e] \leq \sum_{e \in E} \ell_e(\bar{x}_e) \bar{x}_e.$$

For affine latency, $\ell_e(x_e) x_e$ is a convex function of x_e , hence

$$\mathbb{E}_{x \sim \sigma} [\ell_e(x_e) x_e] \geq \ell_e(\bar{x}_e) \bar{x}_e$$

on every edge. Together, these two inequalities imply that x_e is almost surely equal to \bar{x}_e on every edge. To complete the forward direction of the proof, it remains to show that \bar{x} is a pure Nash flow. Because the flows on the edges are almost surely constant, for any player i the cost in \bar{x} is the same as in x ,

$$c_i(\bar{x}) = \sum_{e \in E} \ell_e(\bar{x}_e) \bar{x}_{i,e} = \sum_{e \in E} \mathbb{E}_{x \sim \sigma} [\ell_e(\bar{x}_e) x_{i,e}] \stackrel{\text{a.s.}}{=} \sum_{e \in E} \mathbb{E}_{x \sim \sigma} [\ell_e(x_e) x_{i,e}] = \mathbb{E}_{x \sim \sigma} [c_i(x)].$$

Since x is an equilibrium, $\mathbb{E}_{x \sim \sigma} [c_i(x)] \leq \mathbb{E}_{x \sim \sigma} [c_i(y_i, x_{-i})]$ for all pure strategies y_i and players i , and, because the latencies are affine, $\mathbb{E}_{x \sim \sigma} [\ell_e(y_{i,e} + x_{-i,e})] = \ell_e(y_{i,e} + \bar{x}_{-i,e})$ and hence $\mathbb{E}_{x \sim \sigma} [c_i(y_i, x_{-i})] = c_i(y_i, \bar{x}_{-i})$. Altogether, $c_i(\bar{x}) \leq c_i(y_i, \bar{x}_{-i})$.

To proof the converse, assume x is distributed such that $\bar{x} = \mathbb{E}_{x \sim \sigma} [x]$ is a pure Nash equilibrium and $x_e \stackrel{\text{a.s.}}{=} \bar{x}_e$. Then, as above, $\mathbb{E}_{x \sim \sigma} [c_i(x)] = c_i(\bar{x})$ and $c_i(y_i, \bar{x}_{-i}) = \mathbb{E}_{x \sim \sigma} [c_i(y_i, x_{-i})]$ for all players i and flows y_i . Also, $c_i(\bar{x}) \leq c_i(y_i, \bar{x}_{-i})$ because \bar{x} is a Nash flow. Consequently, $\mathbb{E}_{x \sim \sigma} [c_i(x)] \leq \mathbb{E}_{x \sim \sigma} [c_i(y_i, x_{-i})]$, that is, the distribution of x is a coarse correlated equilibrium. \square

5.4 General Latency

In this section, we investigate games with general latency functions that are non-decreasing, continuously differentiable, and semi-convex. We still assume a maximum slope a , that is, functions ℓ_e such that $|\ell_e(x_e) - \ell_e(y_e)| \leq a|x_e - y_e|$ for all flows x_e, y_e .

5.4.1 Structural Difference to Games with Affine Latency

In the previous section we found that for affine latency, coarse correlated equilibria essentially coincide with the pure Nash equilibrium of the game. Particularly, the worst-case price of anarchy cost guarantee for pure Nash equilibria also holds for all coarse correlated equilibria. The case of general latency functions is structurally different: We present a game with cubic cost functions that has a coarse correlated equilibrium which costs more than 5.29 times as much as an optimal flow. As Roughgarden and Schoppmann (2011) have shown that the price of anarchy for pure, mixed and correlated Nash equilibria with cubic cost functions is 5.0265, this implies that in general the price of anarchy bounds for coarse correlated equilibria do not coincide with the bounds for pure equilibria.

Lemma 5.13. *There is an atomic splittable congestion game with 3 players and cubic latency functions where a coarse correlated Nash equilibrium costs 5.29 times as much as the system optimum.*

Proof. Consider a game with three resources and three players. The cost of a resource is the cube of its load. Each player i allocates in total $\frac{1}{3}$ units of load. The strategy space is similar to the lower bound construction of Roughgarden and Schoppmann (2011): player i can choose either to use resource i or the other two resources.

In a system optimum, each player uses only her single resource strategy. The load on each resource is $\frac{1}{3}$ and the total cost is $\frac{1}{3}(\frac{1}{3})^3 + \frac{1}{3}(\frac{1}{3})^3 + \frac{1}{3}(\frac{1}{3})^3 = \frac{1}{27}$.

For the construction of a coarse correlated equilibrium $x \sim \sigma$, fix some $p \in [0, 1]$. Then, in x with probability p the optimal assignment is played by all players simultaneously, and with probability $1 - p$ all players assign their entire load to their two-resource strategy, such that the load on each resource is $\frac{2}{3}$. The expected cost of a player is

$$\begin{aligned}\mathbb{E}_{x \sim \sigma}[c_i(x)] &= p \frac{1}{3} \left(\frac{1}{3} \right)^3 + (1-p) \frac{1}{3} \left(\left(\frac{2}{3} \right)^3 + \left(\frac{2}{3} \right)^3 \right) \\ &= \left(\frac{1}{3} \right)^4 (p + 2^4(1-p)) = \frac{16 - 15p}{81}\end{aligned}$$

To show that $x \sim \sigma$ is a coarse correlated equilibrium with the right choice of p , we calculate the cost of a unilateral deviation. The pure strategies of any player can be parametrized by $q \in [0, 1]$, where $\frac{q}{3}$ load is put on the single resource strategy and $\frac{1-q}{3}$ load is put on the two-resource strategy. The expected cost of unilaterally deviating to such a strategy $y(q)$ is

$$\begin{aligned}\mathbb{E}_{x \sim \sigma}[c_i(y(q), x_{-i})] &= p \left(\frac{q}{3} \left(\frac{q}{3} \right)^3 + \frac{1-q}{3} \left(\left(\frac{2-q}{3} \right)^3 + \left(\frac{2-q}{3} \right)^3 \right) \right) \\ &\quad + (1-p) \left(\frac{q}{3} \left(\frac{2+q}{3} \right)^3 + \frac{1-q}{3} \left(\left(\frac{2-q}{3} \right)^3 + \left(\frac{2-q}{3} \right)^3 \right) \right) \\ &= \left(\frac{1}{3} \right)^4 (3q^4 - (6p+8)q^3 + (-12p+48)q^2 - (8p+32)q + 16)\end{aligned}$$

The profile x is a coarse correlated equilibrium if the expected cost in x is not greater than the expected cost of any deviation. We hence need to find $p \in [0, 1]$ such that for any $q \in [0, 1]$ the following expression is non-negative:

$$\begin{aligned}\mathbb{E}_{x \sim \sigma}[c_i(y(q), x_{-i})] - \mathbb{E}_{x \sim \sigma}[c_i(x)] \\ = \left(\frac{1}{3} \right)^4 (3q^4 - (6p+8)q^3 + (-12p+48)q^2 - (8p+32)q + 15p)\end{aligned}$$

This is a quartic function of q and, as shown by Lazard (1988), it is non-negative for all $q \in [0, 1]$ if its discriminant is non-negative and p satisfies the condition

$$8 \cdot 3 \cdot (48 - 12p) - 3 \cdot (-8 - 6p)^2 \geq 0.$$

While the latter is true for any $p \in [0, 1]$, the former is true for $p \in [0.714, 0.987]$. Hence, x is a coarse correlated equilibrium if we choose $p = 0.714$. Then, as there are three players,

$$\mathbb{E}_{x \sim \sigma}[C(x)] = 3 \cdot \mathbb{E}_{x \sim \sigma}[c_i(x)] = \frac{16 - 15p}{27}.$$

The ratio of the expected cost of x and the cost of a system optimum is $16 - 15p = 5.29$. \square

The above bound shows that price of anarchy bounds for pure Nash equilibria can in general not be applied to coarse correlated equilibria. However, this is only shown for cubic cost functions, for other classes of cost function (e.g. quadratic or higher degree polynomials) the question remains open. However, a similar lower bound may be possible for other cost function classes, possibly using an insight of the above proof: when an instance has two joint strategies x and y such that both $c_i(x_i, y_{-i}) \geq c_i(y)$ and $c_i(y_i, x_{-i}) \geq c_i(x)$, then a coarse correlated equilibrium may be found by having players alternate between strategies that are close to x and y . [Roughgarden and Schoppmann \(2011\)](#) give a general construction for lower bound instances for pure Nash equilibria in games with polynomial cost of a given degree, and their construction results in games that have the above property.

Remark 5.14. The above lower bound also holds for routing games, as the instance used in the proof can be expressed as a routing game as described in [Roughgarden and Schoppmann \(2011\)](#).

In a second example, we demonstrate a game where the cost difference between a pure Nash equilibrium and a coarse correlated equilibrium depends linearly on the size of the graph, the maximum slope and the number of players n , which corresponds to the style of the upper bound we give in the following sections.

Lemma 5.15. *For any a, k and n , there are atomic splittable routing games with n players, k nodes and latency with maximum slope a that have coarse correlated equilibria that cost $\Theta(\frac{ak}{n})$ more than the pure Nash equilibrium.*

Proof. Consider a graph consisting of two parallel paths of length k , where each edge e has latency $\ell_e(x_e) = a|x_e - \frac{1}{2}|_{\geq 0}$, that is, with zero latency for $x_e \leq \frac{1}{2}$. Clearly, the flow where each player splits her flow equally across both paths is a pure Nash equilibrium. This equilibrium has cost 0. A coarse correlated equilibrium can be described as follows: The players are uniformly at random divided into two sets N_1, N_2 of equal size. Each set is assigned to a path and the players in N_1 route all their flow along their path, while the players in N_2 each route $\frac{1}{2n}(\sqrt{n^2 + 6n + 1} - n - 1)$ flow along the path of N_1 and only the remainder along their own path. A simple calculation shows that the expected cost of a player in this setting is less than the expected cost of any pure strategy. The expected cost of a thus distributed flow x is

$$\mathbb{E}[C(x)] = ak \frac{3n + 1 - \sqrt{n^2 + 6n + 1}}{8n^2} \in \Theta\left(\frac{ak}{n}\right).$$

\square

Note how in the above example the edge flows are not equal to the pure Nash flow, contrasting the uniqueness statement of Theorem 5.12 for affine latency.

5.4.2 Convergence to an Equilibrium

As in the affine case, we investigate how no-regret sequences converge to an equilibrium. We start with a jitter bound.

Lemma 5.16. *Consider repeated play of an instance of atomic splittable routing. Given a no-regret sequence $(x^t)_{t=1}^T$ with average flow \hat{x} and time-average regret at most ε , the jitter of latency on the edges is bounded, that is, $\frac{1}{T} \sum_{t=1}^T |\ell_e(x_e^t) - \ell_e(\hat{x}_e)| \leq \sqrt{a\varepsilon_e}$ with $\sum_{e \in E} \varepsilon_e \leq \varepsilon + \frac{ak}{n}$.*

Proof. As in the previous section, we set $\varepsilon_e = \frac{1}{T} \sum_t \ell_e(x_e^t)(x_e^t - \hat{x}_e)$ and find again that on every edge e ,

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T |\ell_e(x_e^t) - \ell_e(\hat{x}_e)| &\leq \frac{1}{T} \sum_{t=1}^T \sqrt{a(x_e^t - \hat{x}_e)(\ell_e(x_e^t) - \ell_e(\hat{x}_e))} \\ &\leq \sqrt{a \frac{1}{T} \sum_{t=1}^T (x_e^t - \hat{x}_e) \ell_e(x_e^t) - a \underbrace{\frac{1}{T} \sum_{t=1}^T (x_e^t - \hat{x}_e) \ell_e(\hat{x}_e)}_{=0}} = \sqrt{a\varepsilon_e}, \end{aligned} \quad (5.15)$$

It remains to bound ε_e . We use the regret bound of the sequence $(x^t)_{t=1}^T$,

$$\begin{aligned} \sum_{e \in E} \varepsilon_e &= \sum_{i \in N} \sum_{e \in E} \frac{1}{T} \sum_{t=1}^T \ell_e(x_e^t) x_{i,e}^t - \ell_e(x_e^t) \hat{x}_{i,e} \\ &\leq \varepsilon + \sum_{i \in N} \sum_{e \in E} \frac{1}{T} \sum_{t=1}^T \ell_e(x_{-i,e}^t + \hat{x}_{i,e}) \hat{x}_{i,e} - \ell_e(x_e^t) \hat{x}_{i,e} \\ &\leq \varepsilon + \sum_{i \in N} \sum_{P \in \mathcal{P}_i} \hat{x}_{i,P} \sum_{e \in P} \frac{1}{T} \sum_{t=1}^T a |\hat{x}_{i,e} - x_{i,e}^t|, \end{aligned} \quad (5.16)$$

In inequality (5.16), we switch to summing up by path $P \in \mathcal{P}_i$, using $\hat{x}_{i,e} = \sum_{P \in \mathcal{P}_i; e \in P} \hat{x}_{i,P}$. By assumption also $\ell_e(x_{-i,e}^t + \hat{x}_{i,e}) - \ell_e(x_e^t) \leq a |\hat{x}_{i,e} - x_{i,e}^t|$.

Finally, with $|\hat{x}_{i,e} - x_{i,e}^t| \leq \frac{1}{n}$, we have $\sum_{e \in P} \frac{1}{T} \sum_{t=1}^T a |\hat{x}_{i,e} - x_{i,e}^t| \leq \frac{ak}{n}$ for any path P . Since we fixed the total flow to 1 unit, i.e. $\sum_{i \in N} \sum_{P \in \mathcal{P}_i} \hat{x}_{i,P} = 1$, we conclude

$$\sum_{e \in E} \varepsilon_e \leq \varepsilon + \frac{ak}{n}. \quad (5.17)$$

□

In the affine case, the latency on an edge in a coarse correlated equilibrium is almost surely equal to the average latency on that edge. Lemma 5.16 shows that in the general case, the latency in a coarse correlated equilibrium is within $\frac{ak}{n}$ of the average latency. This already hints at the main result of this section: A coarse correlated equilibrium may be more expensive than the pure equilibria, but only within a certain range. This jitter bound, showing that the latency observed on an edge over time is very similar to the latency of the average flow \hat{x} , will be essential for all upcoming statements. Next, we show that the average flow is an approximate Nash flow.

Lemma 5.17. *The average flow \hat{x} is a $\left(\varepsilon + \sqrt{ak(\varepsilon + \frac{ak}{n})} + 2\frac{ak}{n}\right)$ -approximate Nash flow.*

Proof. By assumption ℓ_e is semi-convex, i.e. $\ell_e(x_e)x_e$ is convex, and hence $C(\hat{x}) \leq \frac{1}{T} \sum_{t=1}^T C(x^t)$. We further estimate the average cost of $(x^t)_{t=1}^T$ with the regret bound,

$$\sum_{i \in N} \frac{1}{T} \sum_{t=1}^T c_i(x^t) \leq \varepsilon + \sum_{i \in N} \sum_{e \in E} \frac{1}{T} \sum_{t=1}^T \ell_e(x_{-i}^t + y_{i,e}) y_{i,e} \quad (5.18)$$

for all flows y , and we can estimate for every edge e

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \ell_e(x_{-i,e}^t + y_{i,e}) &\leq \frac{1}{T} \sum_{t=1}^T \ell_e(x^t) + a y_{i,e} \stackrel{(5.15)}{\leq} \ell_e(\hat{x}) + \sqrt{a\varepsilon_e} + \frac{a}{n} \\ &\leq \ell_e(\hat{x}_{-i,e} + y_{i,e}) + a \hat{x}_{i,e} + \sqrt{a\varepsilon_e} + \frac{a}{n} \\ &\leq \ell_e(\hat{x}_{-i,e} + y_{i,e}) + \sqrt{a\varepsilon_e} + 2\frac{a}{n} \end{aligned} \quad (5.19)$$

because the maximum slope is a . Note that for any path P , the Cauchy-Schwarz inequality with $|P| \leq k$ and the jitter bound give

$$\sum_{e \in P} \sqrt{a\varepsilon_e} \leq \sqrt{ak \sum_{e \in P} \varepsilon_e} \leq \sqrt{an(\varepsilon + \frac{ak}{n})} \quad (5.20)$$

Then, continuing where we left off with (5.18),

$$\begin{aligned} \sum_{i \in N} \sum_{e \in E} \frac{1}{T} \sum_{t=1}^T \ell_e(x_{-i}^t + y_{i,e}) y_{i,e} &\stackrel{(5.19)}{\leq} \sum_{i \in N} \sum_{P \in \mathcal{P}_i} y_{i,P} \sum_{e \in P} \left(\ell_e(\hat{x}_{-i,e} + y_{i,e}) + \sqrt{a\varepsilon_e} + 2\frac{a}{n} \right) \\ &\stackrel{(5.20)}{\leq} \sum_{i \in N} c_i(y_i, \hat{x}_{-i}) + \sqrt{an(\varepsilon + \frac{ak}{n})} + 2\frac{ak}{n}, \end{aligned} \quad (5.21) \quad (5.22)$$

Where in (5.21) we switch to summing up over paths and apply (5.19). Combining (5.18) and (5.22) gives the desired

$$C(\hat{x}) \leq \sum_{i \in N} \min_{y_i} c_i(y_i, \hat{x}_{-i}) + \varepsilon + \sqrt{ak(\varepsilon + \frac{ak}{n})} + 2\frac{ak}{n}.$$

□

5.4.3 Price of Total Anarchy

Although no-regret sequences do not converge to pure Nash equilibria in this general case, we are still able to relate their cost to the cost of pure Nash equilibria using the fact that their average flow is an approximate Nash flow. The goal of this section is to estimate the average cost of a no-regret sequence as a ratio to the cost of a socially optimal flow. We start out with the relation of the cost of the sequence to the cost of the average flow.

Lemma 5.18. *The average cost of a sequence $(x^t)_{t=1}^T$ with regret bound ε is at most $\sqrt{am(\varepsilon + \frac{ak}{n})}$ more than the cost of its average flow \hat{x} .*

Proof. The above jitter bound is used.

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T C(x^t) &= \frac{1}{T} \sum_{t=1}^T \sum_{e \in E} \ell_e(x_e^t) x_e^t = \frac{1}{T} \sum_{t=1}^T \sum_{e \in E} \ell(\hat{x}_e) x_e^t + (\ell_e(x_e^t) - \ell_e(\hat{x}_e)) \underbrace{x_e^t}_{\leq 1} \\ &\leq C(\hat{x}) + \sum_{e \in E} \sqrt{a\varepsilon_e} \end{aligned} \quad (5.23)$$

$$\leq C(\hat{x}) + \sqrt{am(\varepsilon + \frac{ak}{n})}. \quad (5.24)$$

In (5.23), we use $\frac{1}{T} \sum_t \ell_e(x_e^t) - \ell_e(\hat{x}_e) \leq \sqrt{a\varepsilon_e}$ from the jitter bound. In (5.24), we again use the Cauchy-Schwarz inequality and $\sum_{e \in E} \varepsilon_e \leq \varepsilon + \frac{ak}{n}$ from the jitter bound. \square

It remains to estimate the cost of ε -approximate Nash flows.

Lemma 5.19. *For an ε -approximate Nash flow x , $C(x) \leq \frac{\lambda}{1-\mu} \cdot C(y) + \frac{1}{1-\mu} \cdot (\varepsilon + 2\frac{ak}{n})$, if the game is locally (λ, μ) -smooth with respect to y .*

Proof. Recall that for games that are locally (λ, μ) -smooth with respect to y ,

$$C(x) \leq \frac{\lambda}{1-\mu} C(y) + \frac{1}{1-\mu} \sum_{i \in N} \sum_{e \in E} c'_{i,e}(x)(x_{i,e} - y_{i,e}),$$

where $c'_i(x)$ is the marginal cost as defined in (5.10). We use that

$$-\ell_e(x_e)y_{i,e} \leq -\ell_e(x_{-i,e} + y_{i,e})y_{i,e} + ay_{i,e}^2$$

and that

$$\sum_{e \in E} x_{i,e}^2 + y_{i,e}^2 \leq \sum_{P \in \mathcal{P}_i} (x_{i,P} + y_{i,P}) \sum_{e \in P} \frac{1}{n} \leq 2 \frac{k}{n^2}$$

because paths have maximum length k edges and $\sum_{P \in \mathcal{P}_i} (x_{i,P} + y_{i,P}) = \frac{2}{n}$ to estimate

$$\begin{aligned} \sum_{e \in E} c'_{i,e}(x)(x_{i,e} - y_{i,e}) &\leq \sum_{e \in E} (\ell_e(x) + ax_{i,e})(x_{i,e} - y_{i,e}) \\ &\leq \sum_{e \in E} (\ell_e(x_e)x_{i,e} + ax_{i,e}^2 - \ell_e(x_e)y_{i,e}) \leq c_i(x) - c_i(y_i, x_{-i}) + 2\frac{ak}{n^2} \end{aligned}$$

Summing this up across all players and using that x is an ε -approximate Nash flow gives

$$\sum_{i \in N} \sum_{e \in E} c'_{i,e}(x)(x_{i,e} - y_{i,e}) \leq \sum_{i \in N} \left(c_i(x) - c_i(y_i, x_{-i}) + 2\frac{ak}{n^2} \right) \leq \varepsilon + \frac{2ak}{n}.$$

Consequently,

$$C(x) \leq \frac{\lambda}{1-\mu} \cdot C(y) + \frac{1}{1-\mu} \left(\varepsilon + \frac{2ak}{n} \right).$$

\square

Corollary 5.20. *Given a class of latency functions \mathcal{L} with maximum slope a such that atomic splittable routing games with these latency functions are locally (λ, μ) smooth with respect to optimal flows y , for a no-regret sequence $(x^t)_{t=1}^T$ with time-average regret ε ,*

$$\frac{1}{T} \sum_{t=1}^T C(x^t) \leq \frac{\lambda}{1-\mu} C(y) + \frac{1}{1-\mu} \left(\varepsilon + \sqrt{ak(\varepsilon + \frac{ak}{n})} + \frac{4ak}{n} \right) + \sqrt{am(\varepsilon + \frac{ak}{n})}.$$

Although games with general latency can have significantly more expensive coarse correlated equilibria than games with affine latency, we have shown that at least for locally smooth games this cost increase is bounded. In Roughgarden and Schoppmann (2011), local smoothness bounds are shown for games with any kind of polynomial latency function.

5.5 Strict ρ -approximate Nash Equilibrium

In the previous sections, we showed the convergence of a no-regret sequence towards a ρ -approximate equilibrium instead of strict ρ -approximate equilibrium. This more ‘loose’ definition allowed to a stronger approximation factor that led to a better price of total anarchy result. However, one can also show that the average flow of a no-regret sequence is a strict approximate Nash equilibrium, that is, an approximate Nash equilibrium in the classic sense. Let us recall the definition of these strict approximate equilibria.

Definition 5.21 (strict ρ -approximate Nash flow)

In an atomic splittable routing game $(N, (V, E), (s_i, t_i)_{i \in N}, (\ell_e)_{e \in E})$, a flow x is a strict ρ -approximate Nash flow if for all players i and all strategies y_i , $c_i(x) \leq c_i(y_i, x_{-i}) + \rho$.

Even-Dar et al. (2009) have shown the average flow of a sequence $(x^t)_{t \in \mathbb{N}}$ with regret bound $\mathcal{R}_i(T)$ for each player $i \in N$ to be a strict $(\sum_{i \in N} \frac{1}{T} \mathcal{R}_i(T))$ -approximate Nash flow in games with affine latency functions. We state a different, potentially stronger bound for the affine case and a new bound for the case of general latency functions. To this end, denote by $\kappa^T = \frac{1}{T} \max_{i \in \mathbb{N}} \mathcal{R}_i(T)$ the maximum time-average regret across all players.

Lemma 5.22. *Consider repeated play of an instance of atomic splittable routing with affine latency. Then, the average flow \hat{x} of a no-regret sequence $(x^t)_{t=1}^T$ is a strict $(\kappa + \sqrt{\frac{am\kappa}{n}})$ -approximate Nash flow, where m is the number of edges and latencies have maximum slope a .*

Proof. We utilize the jitter bounds from the proof of Lemma 5.5. First, we relate the cost of a player i in \hat{x} to her average cost in the sequence $(x^t)_{t=1}^T$,

$$\begin{aligned} c_i(\hat{x}) &= \sum_{e \in E} \ell_e(\hat{x}_e) \hat{x}_{i,e} = \sum_{e \in E} \frac{1}{T} \sum_t \ell_e(\hat{x}_e) x_{i,e}^t \\ &\leq \sum_{e \in E} \frac{1}{T} \sum_{t=1}^T \ell_e(x_e^t) x_{i,e}^t + \sum_{e \in E} \frac{1}{T} \sum_t (\ell_e(\hat{x}_e) - \ell_e(x_e^t)) x_{i,e}^t \\ &\stackrel{(5.2)}{\leq} \sum_{e \in E} \frac{1}{T} \sum_{t=1}^T \ell_e(x_e^t) x_{i,e}^t + \sum_{e \in E} \sqrt{a\varepsilon_e} \frac{1}{n}. \end{aligned}$$

The average cost of player i in the sequence is compared to her cost in a unilateral deviation to some other strategy y_i using the regret bound,

$$\begin{aligned} \sum_{e \in E} \frac{1}{T} \sum_{t=1}^T \ell_e(x_e^t) x_{i,e}^t &\leq \min_{y_i} \sum_{e \in E} \frac{1}{T} \sum_{t=1}^T \ell_e(x_{-i,e}^t + y_{i,e}) y_{i,e} + \mathcal{R}_i(T) \\ &\leq \min_{y_i} \sum_{e \in E} \ell_e(\hat{x}_{-i,e} + y_{i,e}) y_{i,e} + \kappa, \end{aligned} \quad (5.25)$$

where the last inequality holds because the latency functions ℓ_e are affine. Using (5.3) from the jitter bound and the Cauchy-Schwarz inequality, we have $\sum_{e \in E} \sqrt{a\epsilon_e} \leq \sqrt{am\epsilon} \leq \sqrt{amn\kappa}$, so that we can combine the above to

$$c_i(\hat{x}) \leq \min_{y_i} \sum_{e \in E} \ell_e(\hat{x}_{-i,e} + y_{i,e}) y_{i,e} + \kappa + \sqrt{\frac{am\kappa}{n}}$$

that is, player i cannot improve more than $(\kappa + \sqrt{\frac{am\kappa}{n}})$ by deviating to y_i . \square

While the approximation factor of the above bound for strict approximate-equilibria is worse than the factor from Lemma 5.3, the lemma below shows a bound for general cost functions that is close to the bound from Lemma 5.17.

Lemma 5.23. *Consider repeated play of an instance of atomic splittable routing with general latency. The average flow \hat{x} of a no-regret sequence $(x^t)_{t=1}^T$ is a strict $\left(2\sqrt{\frac{am}{n}\left(\kappa + \frac{ak}{n^2}\right)} + 2\frac{am}{n^2} + \kappa\right)$ -approximate Nash flow, where m and k are the number of edges and nodes and the latency functions have maximum slope a .*

Proof. We fix a player i and a deviating strategy y_i , and denote the quantities

$$\alpha = \frac{1}{T} \sum_{t=1}^T c_i(x^t) - c_i(\hat{x}) \quad \text{and} \quad \beta = c_i(y_i, \hat{x}_{-i}) - \sum_{t=1}^T c_i(y_i, x_{-i}^t).$$

Then,

$$c_i(\hat{x}) = \frac{1}{T} \sum_{t=1}^T c_i(x^t) + \alpha \leq \sum_{t=1}^T c_i(y_i, x_{-i}^t) + \mathcal{R}_i(T) + \alpha = c_i(y_i, \hat{x}_{-i}) + \beta + \kappa + \alpha.$$

Before we bound α and β , note that with Cauchy-Schwarz and the jitter bound (5.17),

$$\sum_{e \in E} \sqrt{a\epsilon_e} \leq \sqrt{am \sum_{e \in E} \epsilon_e} \stackrel{(5.17)}{\leq} \sqrt{am \left(\epsilon + \frac{ak}{n}\right)}. \quad (5.26)$$

To bound α , we use $\hat{x}_{i,e} = \frac{1}{T} \sum_{t=1}^T x_{i,e}^t$ to regroup so that we can apply the jitter bound (5.15).

$$\begin{aligned}\alpha &= \frac{1}{T} \sum_{t=1}^T c_i(x^t) - c_i(\hat{x}) = \sum_{e \in E} \left(\frac{1}{T} \sum_{t=1}^T \ell_e(x_e^t) x_{i,e}^t - \ell_e(\hat{x}_e) \hat{x}_{i,e} \right) \\ &= \sum_{e \in E} \frac{1}{T} \sum_t (\ell_e(x_e^t) - \ell_e(\hat{x}_e)) x_{i,e}^t \\ &\stackrel{(5.15)}{\leq} \sum_{e \in E} \sqrt{a\varepsilon_e} \frac{1}{n}. \end{aligned} \quad (5.27)$$

$$\stackrel{(5.26)}{\leq} \sqrt{am \left(\varepsilon + \frac{ak}{n} \right)} \frac{1}{n} \quad (5.28)$$

In (5.27) we also estimated $x_{i,e}^t \leq \frac{1}{n}$ as player i has a total flow value of $\frac{1}{n}$.

For β , applying the jitter bound here is more tricky than in the affine case as, when player i deviates to y_i , the average latency of an edge is not equal to its latency in the average flow. To deal with this, we use the maximum slope assumption to estimate $\ell_e(\hat{x}_{-i} + y_{i,e}) \leq \ell_e(\hat{x}) + ay_{i,e}$ and $-\ell_e(x_{-i,e}^t + y_{i,e}) \leq -\ell_e(x^t) + ax_{i,e}^t$. Also, $y_{i,e} \leq \frac{1}{n}$.

$$\begin{aligned}\beta &= c_i(y_i, \hat{x}_{-i}) - \frac{1}{T} \sum_{t=1}^T c_i(y_i, x_{-i}^t) \\ &= \sum_{e \in E} y_{i,e} \left(\ell_e(\hat{x}_{-i} + y_{i,e}) - \frac{1}{T} \sum_{t=1}^T \ell_e(x_{-i,e}^t + y_{i,e}) \right) \\ &\leq \frac{1}{n} \sum_{e \in E} \left(\frac{1}{T} \sum_{t=1}^T (\ell_e(\hat{x}) - \ell_e(x^t)) + ay_{i,e} + \frac{1}{T} \sum_{t=1}^T ax_{i,e}^t \right) \\ &\stackrel{(5.15)}{\leq} \frac{1}{n} \sum_{e \in E} \left(\sqrt{a\varepsilon_e} + 2\frac{a}{n} \right) \\ &\stackrel{(5.26)}{\leq} \frac{1}{n} \left(\sqrt{am \left(\varepsilon + \frac{ak}{n} \right)} + 2\frac{am}{n} \right) \\ &= \frac{1}{n} \sqrt{am \left(\varepsilon + \frac{ak}{n} \right)} + 2\frac{am}{n^2} \end{aligned}$$

The lemma follows from

$$\begin{aligned}\alpha + \beta + \kappa &\leq 2\frac{1}{n} \sqrt{am \left(\varepsilon + \frac{ak}{n} \right)} + 2\frac{am}{n^2} + \kappa \\ &\leq 2\sqrt{\frac{am}{n} \left(\kappa + \frac{ak}{n^2} \right)} + 2\frac{am}{n^2} + \kappa. \end{aligned}$$

□

5.6 Conclusion

We considered the efficiency of no-regret sequences in atomic splittable routing. We showed that with affine latency functions, such sequences are guaranteed fast convergence to a pure Nash equilibrium. Our results not only improve the previously known bound, but also show that the flow ‘on most days’ of a no-regret sequence is an approximate equilibrium. Additionally, we showed that coarse correlated equilibria coincide with pure Nash equilibria in the sense that the flow on any edge is almost surely equal. These results imply that no-regret play behaves ‘nicely’ as hoped for, meaning that it gives a good approximation of the behavior modeled by the both more simple and more restrictive pure Nash equilibria, but requires less assumptions.

Such ‘nice’ behavior is not the case with general cost functions. Here, we gave lower bounds showing that already with cubic latency functions the worst-case price of anarchy bound derived for pure Nash equilibria does not hold for coarse correlated Nash equilibria. We then showed that the time-average cost difference between a no-regret sequence and an optimal flow is bounded by the price of anarchy for pure equilibria and an additive term based on the network size and the maximum slope of a latency function.

A tight upper bound on the price of anarchy for coarse correlated Nash equilibria in games with general latency functions remains open. A first step towards such a bound would be to construct, if possible, lower bound instances for all polynomial cost functions that exhibit coarse correlated equilibria beyond the worst-case price of anarchy bound for pure equilibria, possibly using the pattern of our lower bound construction for cubic cost.

Bibliography

- D. Acemoglu and M. K. Jensen. Aggregate comparative statics. *Games and Economic Behavior*, 81(0):27–49, 2013.
- H. Ackermann, H. Röglin, and B. Vöcking. Pure Nash equilibria in player-specific and weighted congestion games. *Theoretical Computer Science*, 410(17):1552–1563, 2009.
- S. Albers. On the value of coordination in network design. *SIAM Journal on Computing*, 38(6):2273–2302, 2009.
- M. Allais. Le comportement de l’homme rationnel devant le risque: critique des postulats et axiomes de l’école américaine. *Econometrica*, 21(4):503–546, 1953.
- R. Amir. Cournot oligopoly and the theory of supermodular games. *Games and Economic Behavior*, 15(2):132–148, 1996.
- S. P. Anderson and R. Renault. Efficiency and surplus bounds in cournot competition. *Journal of Economic Theory*, 113(2):253–264, 2003.
- E. Anshelevich, A. Dasgupta, J. 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.
- A. Archer, J. Feigenbaum, A. Krishnamurthy, R. Sami, and S. Shenker. Approximation and collusion in multicast cost sharing. *Games and Economic Behavior*, 47(1):36–71, 2004.
- S. Athey. Monotone comparative statics under uncertainty. *The Quarterly Journal of Economics*, 117(1):187–223, 2002.
- G. Ausiello, A. D’Atri, and M. Protasi. Structure preserving reductions among convex optimization problems. *Journal of Computer and System Sciences*, 21(1):136 – 153, 1980.
- B. Awerbuch and R. Kleinberg. Online linear optimization and adaptive routing. *Journal of Computer and System Sciences*, 74(1):97–114, 2008.
- B. Awerbuch, Y. Azar, A. Epstein, V. Mirrokni, and A. Skopalik. Fast convergence to nearly optimal solutions in potential games. In *Proceedings of the 9th ACM conference on Electronic Commerce*, pages 264–273. ACM, 2008.
- D. Bertsekas and R. Gallagher. *Data Networks*. Prentice-Hall, Upper Saddle River, NJ, USA, 1992.
- U. Bhaskar, L. Fleischer, D. Hoy, and C.-C. Huang. Equilibria of atomic flow games are not unique. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 748–757. SIAM, 2009.
- V. Bilò, A. Fanelli, M. Flammini, G. Melideo, and L. Moscardelli. Designing fast converging cost sharing methods for multicast transmissions. *Theory of Computing Systems*, 47(2):507–530, 2010.

- A. Blum and Y. Mansour. Learning, regret minimization, and equilibria. In [Nisan et al. \(2007\)](#), chapter 4, pages 79–101.
- A. Blum, E. Even-Dar, and K. Ligett. Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games. In *Proceedings of the 25th annual ACM Symposium on Principles of Distributed Computing*, pages 45–52. ACM, 2006.
- A. Blum, M. Hajiaghayi, K. Ligett, and A. Roth. Regret minimization and the price of total anarchy. In *Proceedings of the 40th annual ACM symposium on Theory of Computing*, pages 373–382. ACM, 2008.
- A. Bogomolnaia, R. Holzman, and H. Moulin. Sharing the cost of a capacity network. *Mathematics of Operations Research*, 35(1):173–192, 2010.
- A. Borodin and R. El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, Cambridge, UK, 1998.
- D. Braess. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung*, 12(1):258–268, 1968.
- J. A. Brander and B. J. Spencer. Export subsidies and international market share rivalry. *Journal of International Economics*, 18(1):83–100, 1985.
- J. Bulow, J. Geanakoplos, and P. Klempner. Multimarket oligopoly: Strategic substitutes and complements. *The Journal of Political Economy*, 93(3):488–511, 1985.
- I. Caragiannis. Efficient coordination mechanisms for unrelated machine scheduling. *Algorithmica*, 66(3):512–540, 2013.
- I. Caragiannis, A. Fanelli, N. Gravin, and A. Skopalik. Efficient computation of approximate pure Nash equilibria in congestion games. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science*, pages 532–541. IEEE, 2011.
- C. Chekuri, J. Chuzhoy, L. Lewin-Eytan, J. Naor, and A. Orda. Non-cooperative multicast and facility location games. *IEEE Journal on Selected Areas in Communications*, 25(6):1193–1206, 2007.
- H. Chen and T. Roughgarden. Network design with weighted players. *Theory of Computing Systems*, 45(2):302–324, 2009.
- Y. Chen and J. Zhang. Design of price mechanisms for network resource allocation via price of anarchy. *Mathematical Programming, Series A*, 131(1-2):333–364, 2012.
- X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3):14, 2009.
- H.-L. Chen, T. Roughgarden, and G. Valiant. Designing network protocols for good equilibria. *SIAM Journal on Computing*, 39(5):1799–1832, 2010.
- G. Christodoulou, E. Koutsoupias, and A. Nanavati. Coordination mechanisms. *Theoretical Computer Science*, 410(36):3327–3336, 2009.
- R. Cole, J. R. Correa, V. Gkatzelis, V. Mirrokni, and N. Olver. Inner product spaces for minsum coordination mechanisms. In *Proceedings of the 43rd annual ACM symposium on Theory of Computing*, pages 539–548. ACM, 2011.
- R. Cominetti, J. R. Correa, and N. E. Stier-Moses. The impact of oligopolistic competition in networks. *Operations Research*, 57(6):1421–1437, 2009.
- L. C. Corchón. Comparative statics for aggregative games. the strong concavity case. *Mathematical Social Sciences*, 28(3):151–165, 1994.
- J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.

- J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. On the inefficiency of equilibria in congestion games. In *Integer Programming and Combinatorial Optimization*, pages 167–181. Springer, 2005.
- J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. A geometric approach to the price of anarchy in nonatomic congestion games. *Games and Economic Behavior*, 64(2):457–469, 2008.
- A. A. Cournot. *Recherches Sur Les Principes Mathematiques De La Theorie De La Richesse*. Hachette, Paris, France, 1838.
- A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. *ACM Transactions on Algorithms*, 3(1):1–17, 2007.
- A. Czumaj, P. Krysta, and B. Vöcking. Selfish traffic allocation for server farms. *SIAM Journal on Computing*, 39(5):1957–1987, 2010.
- S. Dafermos. Traffic equilibrium and variational inequalities. *Transportation Science*, 14(1):42–54, 1980.
- C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- G. Debreu. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences of the United States of America*, 38(10):886–893, 1952.
- M. Desrochers, P. Marcotte, and M. Stan. The congested facility location problem. *Location Science*, 3(1):9–23, 1995.
- A. Dixit. Comparative statics for oligopoly. *International Economic Review*, 27(1):107–122, 1986.
- J. Eaton and G. M. Grossman. Optimal trade and industrial policy under oligopoly. *The Quarterly Journal of Economics*, 101(2):383–406, 1986.
- A. S. Edlin and C. Shannon. Strict monotonicity in comparative statics. *Journal of Economic Theory*, 81(1):201–219, 1998.
- J. Edmonds and R. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM*, 19(2):248–264, 1972.
- A. Epstein, M. Feldman, and Y. Mansour. Strong equilibrium in cost sharing connection games. *Games and Economic Behavior*, 67(1):51–68, 2009.
- E. Even-Dar, A. Kesselman, and Y. Mansour. Convergence time to Nash equilibrium in load balancing. *ACM Transactions on Algorithms*, 3(3):1–21, 2007.
- E. Even-Dar, Y. Mansour, and U. Nadav. On the convergence of regret minimization dynamics in concave games. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, pages 523–532. ACM, 2009.
- P. von Falkenhausen and T. Harks. Optimal cost sharing for resource selection games. *Mathematics of Operations Research*, 38(1):184–208, 2013a.
- P. von Falkenhausen and T. Harks. Optimal cost sharing protocols for scheduling games. In *Proceedings of the 12th ACM conference on Electronic Commerce*, pages 285–294. ACM, 2011.
- P. von Falkenhausen and T. Harks. Optimal cost sharing for capacitated facility location games. In *Proceedings of the 12th Cologne-Twente Workshop on Graphs and Combinatorial Optimization*, pages 99–102. CTIT University of Twente, 2013b.
- K. Fan. Fixed point and minmax theorems in locally convex topological linear spaces. *Proceedings of the National Academy of Sciences of the United States of America*, 38(2):121–126, 1952.
- P. Février and L. Linnemer. Idiosyncratic shocks in an asymmetric cournot oligopoly. *International Journal of Industrial Organization*, 22(6):835–848, 2004.
- A. Fiat and C. Papadimitriou. When the players are not expectation maximizers. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory*, volume 6386 of *LNCS*, pages 1–14. Springer, 2010.

- S. Fischer and B. Vöcking. Adaptive routing with stale information. *Theoretical Computer Science*, 410(36):3357–3371, 2009.
- S. Fischer, H. Räcke, and B. Vöcking. Fast convergence to wardrop equilibria by adaptive sampling methods. *SIAM Journal on Computing*, 39(8):3700–3735, 2010.
- D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. *Theoretical Computer Science*, 410(36):3305–3326, 2009.
- M. Gairing, B. Monien, and K. Tiemann. Routing (un-)splittable flow in games with player-specific affine latency functions. *ACM Transactions on Algorithms*, 7(3):31, 2011.
- G. Gaudet and S. W. Salant. Increasing the profits of a subset of firms in oligopoly models with strategic substitutes. *The American Economic Review*, 81(3):658–665, 1991.
- I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.
- I. L. Glicksberg. A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1):170–174, 1952.
- T. Harks and P. von Falkenhausen. Quantitative comparative statics for a multimarket paradox. In *Proceedings of the 9th International Conference on Web and Internet Economics*, volume 8289 of *LNCS*, pages 230–231. Springer, 2013.
- T. Harks. Stackelberg strategies and collusion in network games with splittable flow. *Theory of Computing Systems*, 48(4):781–802, 2011.
- T. Harks and M. Klimm. On the existence of pure Nash equilibria in weighted congestion games. *Mathematics of Operations Research*, 37(3):419–436, 2012.
- T. Harks and K. Miller. The worst-case efficiency of cost sharing methods in resource allocation games. *Operations Research*, 59(6):1491–1503, 2011.
- A. Haurie and P. Marcotte. On the relationship between Nash-Cournot and Wardrop equilibria. *Networks*, 15(3):295–308, 1985.
- E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.
- M. Hoefer. Competitive cost sharing with economies of scale. *Algorithmica*, 60(4):743–765, 2011.
- S. Ieong, R. McGrew, E. Nudelman, Y. Shoham, and Q. Sun. Fast and compact: A simple class of congestion games. In *Proceedings of the 20th National Conference on Artificial Intelligence*, pages 489–494. AAAI Press / MIT Press, 2005.
- N. Immorlica, L. Li, V. Mirrokni, and A. Schulz. Coordination mechanisms for selfish scheduling. *Theoretical Computer Science*, 410(17):1589–1598, 2009.
- R. Johari and J. N. Tsitsiklis. Efficiency loss in Cournot games. Technical report, LIDS-P-2639, Laboratory for Information and Decision Systems, MIT, 2005.
- R. Johari and J. N. Tsitsiklis. A scalable network resource allocation mechanism with bounded efficiency loss. *IEEE Journal on Selected Areas in Communications*, 24(5):992–999, 2006.
- R. Johari and J. N. Tsitsiklis. Efficiency of scalar-parameterized mechanisms. *Operations Research*, 57(4):823–839, 2009.
- S. Kakutani. A generalization of Brouwer’s fixed point theorem. *Duke Mathematics Journal*, 8(3):457–458, 1941.
- A. T. Kalai and S. Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.

- R. Kleinberg, G. Piliouras, and É. Tardos. Multiplicative updates outperform generic no-regret learning in congestion games. In *Proceedings of the 41st annual ACM Symposium on Theory of Computing*, pages 533–542. ACM, 2009.
- K. Kollia and T. Roughgarden. Restoring pure equilibria to weighted congestion games. In *Proceedings of the 38th International Colloquium on Automata, Languages and Programming*, volume 6756 of *LNCS*, pages 539–551. Springer, 2011.
- J. Könemann, S. Leonardi, G. Schäfer, and S. van Zwam. A group-strategyproof cost sharing mechanism for the Steiner forest game. *SIAM Journal on Computing*, 37(5):1319–1341, 2008.
- E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Symposium on Theoretical Aspects of Computer Science*, volume 1563 of *LNCS*, pages 404–413. Springer, 1999.
- P. Krugman. Scale economies, product differentiation, and the pattern of trade. *The American Economic Review*, 70(5):950–959, 1980.
- N. S. Kukushkin. A fixed-point theorem for decreasing mappings. *Economic Letters*, 46(1):23–26, 1994.
- D. Lazard. Quantifier elimination: optimal solution for two classical examples. *Journal of Symbolic Computation*, 5(1):261–266, 1988.
- M. J. Melitz. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725, 2003.
- G. B. Mertzios. Fast convergence of routing games with splittable flows. RWTH Aachen, Department of Computer Science, 2008.
- I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, 1996.
- P. Milgrom. Comparing optima: Do simplifying assumptions affect conclusions? *Journal of Political Economy*, 102(3):607–15, 1994.
- P. Milgrom and J. Roberts. Comparing equilibria. *American Economic Review*, 84(3):441–59, 1994.
- P. Milgrom and J. Roberts. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):1255–1277, 1990.
- P. Milgrom and C. Shannon. Monotone comparative statics. *Econometrica*, 62(1):157–80, 1994.
- V. S. Mirrokni and A. Vetta. Convergence issues in competitive games. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 3122 of *LNCS*, pages 183–194. Springer, 2004.
- H. Moulin. The price of anarchy of serial, average and incremental cost sharing. *Economic Theory*, 36(3):379–405, 2008.
- H. Moulin. An efficient and almost budget balanced cost sharing method. *Games and Economic Behavior*, 70(1):107–131, 2010.
- H. Moulin. Axiomatic cost and surplus sharing. In K. Arrow, A. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, chapter 6, pages 289–357. Elsevier, 2002.
- H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.
- J. Nash. *Non-cooperative games*. PhD thesis, Princeton, 1950.
- J. von Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320, 1928.
- J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, USA, 1944.
- A. Neyman. Correlated equilibrium and potential games. *International Journal of Game Theory*, 26(2):223–227, 1997.

- N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, Cambridge, UK, 2007.
- W. Novshek. On the existence of Cournot equilibrium. *The Review of Economic Studies*, 52(1):85–98, 1985.
- A. Orda, R. Rom, and N. Shimkin. Competitive routing in multi-user communication networks. *IEEE/ACM Transactions on Networking*, 1(5):510–521, 1993.
- M. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, MA, USA, 1994.
- R. Paes Leme, V. Syrgkanis, and É. Tardos. Sequential auctions and externalities. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 869–886. SIAM, 2012.
- M. Pál and É. Tardos. Group strategyproof mechanisms via primal-dual algorithms. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 584–593. IEEE, 2003.
- A. Pigou. *The Economics of Welfare*. Macmillan, London, UK, 1920.
- J. K.-H. Quah. The comparative statics of constrained optimization problems. *Econometrica*, 75(2): 401–431, 03 2007.
- H. C. Quirmbach. Comparative statics for oligopoly: Demand shift effects. *International Economic Review*, 29(3):451–459, 1988.
- R. Raz and S. Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proceedings of the 29th Annual ACM Symposium on Theory Computing*, pages 475–484. ACM, 1997.
- O. Richman and N. Shimkin. Topological uniqueness of the Nash equilibrium for selfish routing with atomic users. *Mathematics of Operations Research*, 32(1):215–232, 2007.
- J. Roberts and H. Sonnenschein. On the existence of Cournot equilibrium without concave profit functions. *Journal of Economic Theory*, 13(1):112–117, 1976.
- R. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- T. Roughgarden. Selfish routing with atomic players. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1184–1185. SIAM, 2005a.
- T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge, MA, USA, 2005b.
- T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2002.
- T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st annual ACM Symposium on Theory of Computing*, pages 513–522. ACM, 2009.
- T. Roughgarden and F. Schoppmann. Local smoothness and the price of anarchy in atomic splittable congestion games. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 255–267. SIAM, 2011.
- T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
- O. Rozenfeld and M. Tennenholtz. Strong and correlated strong equilibria in monotone congestion games. In *Proceedings of the 2nd International Workshop on Internet and Network Economics*, volume 4286 of *LNCS*, pages 74–86. Springer, 2006.
- A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, Berlin, Germany, 2003.
- C. Shannon. Weak and strong monotone comparative statics. *Journal of Economic Theory*, 5(2):209–27, 1995.
- É. Tardos and T. Wexler. Network formation games and the potential function method. In [Nisan et al. \(2007\)](#), chapter 19, pages 487–516.

- D. Topkis. Equilibrium points in nonzero n -person submodular games. *SIAM Journal on Control and Optimization*, 17(6):773–787, 1979.
- D. Topkis. *Supermodularity and Complementarity*. Princeton University Press, Princeton, NJ, USA, 1998.
- J. N. Tsitsiklis and Y. Xu. Profit loss in cournot oligopolies. *Operations Research Letters*, 41(4):415–420, 2013.
- Y. Ushio. Approximate efficiency of cournot equilibria in large markets. *Review of Economic Studies*, 52(4):547–56, 1985.
- A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Proceedings of the 43th Annual IEEE Symposium on Foundations of Computer Science*, pages 416–425. IEEE, 2002.
- X. Vives. Games with strategic complementarities: New applications to industrial organization. *International Journal of Industrial Organization*, 23(7-8):625–637, 2005.
- X. Vives. Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics*, 19(3):305–321, 1990.
- D. P. Williamson and D. B. Shmoys. *The design of approximation algorithms*. Cambridge University Press, Cambridge, UK, 2011.
- F. Yao, A. Demers, and S. Shenker. A scheduling model for reduced CPU energy. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 374–382. IEEE, 1995.
- M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning*, pages 928–936. AAAI Press, 2003.