Εθνικο Μετσοβιο Πολυτεχνειο
Σχολη Ηλεκτρολογων Μηχανικων και Μηχανικων Υπολογιστων
Τομεας Τεχνολογιας Πληροφορικης και Υπολογιστων

Ανάλυση Τιμήματος της Αναρχίας σε Συνδυαστικές Δημοπρασίες

Διπλωματικη εργασια

του

ΦΙΛΙΠΠΟΥ ΛΑΖΟΥ

Επιβλέπων: Δημήτριος Φωτάκης
Επίκουρος Καθηγητής Ε.Μ.Π.

Εργαστηριο Λογικης και Επιστημης Υπολογισμων
Αθήνα, Ιούλιος 2015

Εθνικό Μετσόβιο Πολυτεχνείο
Σχολή Ηλεκτρολόγων Μηχανικών και Μηχανικών Υπολογιστών
Τομέας Τεχνολογίας Πληροφορικής και Υπολογιστών
Εργαστήριο Λογικής και Επιστήμης Υπολογισμών

Ανάλυση Τιμήματος της Αναρχίας σε Συνδυαστικές Δημοπρασίες

Διπλωματική Εργασία

του

ΦΙΛΙΠΠΟΥ ΛΑΖΟΥ

Επιβλέπων: Δημήτριος Φωτάκης
Επίκουρος Καθηγητής Ε.Μ.Π.

Εγκρίθηκε από την τριμελή εξεταστική επιτροπή την 27η Ιουλίου 2015.

(Υπογραφή)	*(Υπογραφή)*	*(Υπογραφή)*
....	....	....
Δημήτριος Φωτάκης	Νικόλαος Παπασπύρου	Αριστείδης Παγουρτζής
Επίκουρος Καθηγητής	Αναπληρωτής Καθηγητής	Αναπληρωτής Καθηγητής
Ε.Μ.Π.	Ε.Μ.Π.	Ε.Μ.Π.

Αθήνα, Ιούλιος 2015

(Υπογραφή)

....

Φιλιππος Λαζος
Διπλωματούχος Ηλεκτρολόγος Μηχανικός και Μηχανικός Υπολογιστών Ε.Μ.Π.

Ευχαριστίες

Με την κατάθεση της Διπλωματικής μου εργασίας ολοκληρώνεται ο κύκλος των προπτυχιακών μου σπουδών στη Σχολή Ηλεκτρολόγων Μηχανικών και Μηχανικών Ηλεκτρονικών Υπολογιστών του Εθνικού Μετσόβιου Πολυτεχνείου. Στη διάρκεια αυτών των σπουδών απέκτησα πολύτιμη εμπειρία και γνώση σε πολλά επιστημονικά πεδία και παρακολούθησα τα μαθήματα σημαντικών καθηγητών. Προς όλους αυτούς θα ήθελα να εκφράσω και τις ευχαριστίες και την ευγνωμοσύνη μου.

Ιδιαίτερες ευχαριστίες, ωστόσο, οφείλω σε όλους όσοι στάθηκαν αρρωγοί και συμπαραστάτες στη διαδικασία εκπόνησης και παρουσίασης της Διπλωματικής μου εργασίας. Θερμά ευχαριστώ τον επόπτη μου καθηγητή κ. Δημήτρη Φωτάκη για την επιστημονική του καθοδήγηση και την αμέριστη βοήθειά του σε όλα τα στάδια της μελέτης μου. Τους καθηγητές κ. Άρη Παγουρτζή και κ. Νίκο Παπασπύρου για όσα μου έμαθαν κατά τη διάρκεια των σπουδών μου και για την ουσιαστική τους υποστήριξη στη διαδικασία ολοκλήρωσης αυτής της προσπάθειας. Τον καθηγητή κ. Ηλία Κουτσουπιά για τις καίριες υποδείξεις του και την εμπιστοσύνη με την οποία με περιβάλλει. Τους φίλους και συναδέλφους Αλέξανδρο Κάββο, υποψήφιο διδάκτορα του Oxford University και Μικελή Κιτρομηλίδη, υποψήφιο διδάκτορα του Imperial College για την ηθική τους υποστήριξη. Τους γονείς μου για την υλική και συναισθηματική τους συμπαράσταση σε όλη τη διάρκεια των σπουδών μου.

Περίληψη

Το Τίμημα της Αναρχίας ορίζεται ως ο λόγος του κόστους της χειρότερης ισορροπίας Nash προς την καλύτερη δυνατή έκβαση του στρατηγικού παιγνίου και χρησιμοποιείται για την μοντελοποίηση της απόδοσης ενός παιγνίου. Σε πολλά περιβάλλοντα έχουν βρεθεί άνω και κάτω φράγματα για το Τίμημα της Αναρχίας. Θα παρουσιάσουμε μια θεώρια η οποία χαρακτηρίζει παίγνια με κάποια ευρωστία στη δομή τους, η οποία επιτρέπει τον υπολογισμό τέτοιων φραγμάτων (κάποιες φορές με ακρίβεια) με πιο χαλαρές προυποθέσεις ως προς τον ορισμό της ισορροπίας στην οποία θα φτάσουν οι παίχτες . Στη συνέχεια, θα δείξουμε πως εφαρμόζεται αυτή η θεωρία στις Δημοπρασίες, επιτρέποντας μας να κάνουμε προβλέψεις άνω φραγμάτων σε καταστάσεις που οι ίδιοι παίχτες συμμετέχουν σε πολλές δημοπρασίες ταυτόχρονα. Τέλος θα παρουσιάσουμε και μια θεωρία η οποία χρησιμοποιεί αποτελέσματα από Computational Complexity και Communication Complexity για να δώσει κάτω φράγματα στο Τίμημα της Αναρχίας, συμπληρώνοντας την ανάλυση που έγινε για τα άνω φράγματα.

Λέξεις Κλειδιά

Τίμημα της Αναρχίας, Στρατηγικά Παίγνια, Υπολογιστική Πολυπλοκότητα, Πολυπλοκότητα Επικοινωνίας, Δημοπρασίες, Ισορροπία Nash , Σχεδιασμός Μηχανισμών

Abstract

The price of anarchy (POA), defined as the ratio of the worst-case objective function value of a Nash equilibrium of a game and that of an optimal outcome, quantifies the inefficiency of selfish behavior. In many cases both upper and lower bounds for the POA have been found. We will present an approach that defines some games as 'robust', allowing bounds of the POA (often tight) to automatically extend to weaker definitions of equilibria reached by the players. Next we will show how a variation of this approach applies to auctions, giving us the tools to find how upper bounds on the POA for isolated auctions extend when players take part in many such auctions at the same time or sequentially. Lastly, we will exploit results from Computational Complexity and Communication Complexity to show a way to find lower bounds for auctions and other games, complementing our previous upper bound analysis.

Keywords

Price of Anarchy, Strategic Games, Computational Complexity, Communications Complexity, Auctions, Nash Equilibrium, Mechanism Design, Congestion Games

Contents

Chapter 1

Introduction

1.1 Auctions and Strategic Games

The problem of allocating discrete, sparse resources among players has long been studied in Economics and lately is getting a lot of attention in Computer Science as well. One way to approach this problem is to model it through auctions, strategic situations in which the strategic players (which can be any number of entities) try to obtain certain outcomes by using money to leverage their preferences. As an auction we can classify an vast amount of situations where players pay (or lose some fractional resource) in order to gain something more valuable to them.

The auction setting is interesting for a variety of reasons. First, auctions are widely used in real life in consumer, corporate or government settings. One could argue that today they are more relevant than ever, since the emergence of the World Wide Web gave rise to many more opportunities for allocation of resources like bandwidth or space for ads on a webpage and the number of strategic players has increased as well because over the internet humans and computers interact in the same way. Companies like Google, and many more, rely on running an inconceivable amount of auctions in order to properly outsource their resources and generate a large part of their income. Governments use auctions not so much for their own gain, but as a fair way to distribute sparse resources to those who need them the most.

There are many auction types, with some types being used in practice and others formulated from theory. One of most common auctions used in practice are the English auction, instantly recognizable to anyone interested in buying art. It consist of the auctioneer increasing prices in a stepwise fashion, until only one bidder remains. Another type of auction, also for one item is the first price auction, in which bidders submit one closed bid and the auctioneer gives the item to the highest of them. The idea behind all these auctions is that whoever wants the item most will be most willing to pay.

When the goal is maximizing social welfare, disregarding the payments and focusing on making players happy, an early theoretical breakthrough lead to an optimal auction: the Vickrey-Clarke-Groves mechanism (or VCG) [Vic61, Cla71, Gro73]. This auction (which can also be used as a recipe to create optimal mechanisms in a variety of settings) achieves its performance by internalizing externalities. In other words, ensuring the winning players pay proportionately to the loss they caused to the rest of society. However, even though optimal, the VCG is rarely used in practice because it is quite complicated to explain to players and largely inefficient in a variety of settings. Revenue maximization has proven to be far more elusive with initial results from Myerson for single item auction [Mye81] and recent developments for multiple items by Giannakopoulos [GK14].

However, many of the auction derived from theory are impractical to implement. Consequently, a reasonable alternative would be to analyze preexisting auctions and maybe fine-tune them. Especially interesting is the case of auctioning multiple items. Trying to use the VCG for this case is almost impossible, since it would require an exponential amount of bids from each participant. Not only that, but bids would have to be submitted secretly and at the same time. A line of research from Christodoulou et al [CKS08], Bhawalkar and Roughgarden [BR11], Paes Leme et al [LST12], Caragiannis et al [CKKK11] and others has focused on quantifying the efficiency of simple mechanisms (like the first price auction) when running simultaneously and sequentially. Finally, there was a significant advancement by Syrgkanis and Tardos [ST13] which unified the previous approaches by finding their underlying characteristics. Through their research, we have a deeper understanding of simple mechanisms, as well as concrete proof of their efficiency. As a reference point, an optimal auction (say VCG) is only $\frac{e}{e-1}$ times better than the much simpler parallel first price auction.

Auctions have been studied from both the Economic standpoint, in terms of generated revenue or welfare to society, and from Computer Science, to make sure they are easy to use and tractable to implement. In this thesis we showcase the use techniques from the modern algorithmic game theory toolbox, to show that simple auctions have very good efficiency. The goal is double: to formally define auctions and prove their efficiency and show the intricacies of these techniques and how they can be used in different settings. We will focus on showing good qualities in terms of maximizing the benefit to society, disregarding payments and revenue.

1.2 Outline of this thesis

This thesis is divided into several chapters and is written in a way to be useful for readers with a modest mathematical background, as it assumes very little in the way of Economics or CS. The second chapter will lay all the groundwork required in order to model situations of strategic interaction between selfish players. The basics of Game Theory will be presented, with a focus on the efficiency of different equilibria. The basics of Mechanism Design will be shown as well.

In Chapter 2 we will present the first major new technique of algorithmic game theory, Smooth Games. Although not directly useful for auctions, it is still relevant for a deeper understanding of the techniques to follow. In Chapter 3 an immediate follow-up to Smooth Games will be discussed, which is specifically tailored for measuring efficiency in parallel auctions, as well as helping us design them. Both these chapters will provide tools to get upper bound proofs for games, which will be complemented by another technique from Chapter 5 that produces very robust lower bounds.

In Chapters 3 and 5, there will be a treatment of congestion games using similar techniques, to be juxtaposed with auctions and to provide a gentler introduction to some more complex proofs.

Chapter 2

Preliminaries

2.1 Strategic Games

Game Theory is the study of strategic decision making. It is a branch of mathematics that models and studies strategic situations where players choose different actions towards maximizing some goal. The nature of Game Theory is very general and encompasses a large variety of situations across different sciences.

In this chapter we will lay the necessary groundwork onto which auctions, among other settings, will be analysed. As expected, the main mathematical object studied in Game Theory is the (strategic) *game*, comprised by of following sets:

- A set N of n players
- For each player i, a set S_i of his available strategies. An instance of the game where every player has chosen a strategy is called an outcome and defined by a vector $s \in S_1 \times \ldots \times S_n = S$.
- Each player also has a utility $u_i : S \to \mathbb{R}$ which is a representation of his 'gain' or 'payoff' from each outcome.

Players are assumed to be strategic and will chose strategies in order to maximize their payoff. Also, everything is assumed to be public information and players choose their action at the same time. For this reason, these games are also called *one-shot* games.

Just to get things started, imagine the following game theoretic scenario, where two prisoners (the players) can either cooperate or defect. Defecting will reduce their sentence, but if they both defect they will serve for longer than if they both cooperated.

	Cooperate	Defect
Cooperate	-2,-2	-5,0
Defect	0,-5	-4,-4

Since we have only two players, we can represent the whole game concisely by a matrix. One player chooses the row and the other the column. Each cell of the matrix has two values, which are the payoffs for each player. The numbers can be thought of as years in prison and are negative since players want to maximize their gain, or in this case minimize their loss.

This simple example shows how the incentives of each player are influenced by his valuation. For example, if the row player chooses to cooperate, the other player would prefer to defect. However, we need some way to define which outcome (among the 4 available) players will choose, since right now we can calculate each player's gain for any outcome but we aren't sure what constitutes a 'good' move.

2.1.1 Equilibria

Equilibria are a solution concept for strategic games. They represent a stable situations reached by selfish, strategic players who try to maximize their own gain, after perfect play. The pivotal assumption of Game Theory is that players will be able to reach an equilibrium, thus our goal is to find them and quantify their properties.

2.1.1.1 Pure Nash Equilibria

Out of many different possible 'solutions' to a game, the most important is the pure Nash Equilibrium. It was devised by John Nash in 1950 and was (among other theoretical breakthroughs) instrumental in revitalizing Game Theory and helping it transcend to other disciplines.

Definition 1 (Pure Nash Equilibrium (PNE))**.** *An outcome* $\mathbf{s} \in S$ *is a PNE if for each player* i *we have*

$$u_i(s'_i, \mathbf{s}_{-i}) \leq u_i(\mathbf{s})$$

for any $s'_i \in S_i$*.*

By the way, the notation s_{-i} means every element of the vector but s_i, so $(s'_i, \mathbf{s}_{-i})$ replaces s_i for s'_i in $\mathbf{s}$. The PNE describes a stable outcome where no player can increase his gain through unilateral change of strategy. That's why it's reasonable to assume that whenever possible players will reach a PNE, since in other outcomes there would be players who have chosen the wrong action. Another way of looking at the PNE is as an outcome where every player plays his best response against the others.

Let's try and find a PNE of the previous game. The only outcome that is a PNE is $s =$ (defect,defect). The gain of each player is $u_1(s) = u_2(s) = -4$ and their only alternative strategy is to cooperate which leads to u_1(cooperate,defect) $=$ u_2(defect,cooperate) $= -5$ which is worse. So, s is a PNE, *even* though players could have both cooperated and gotten u_1(cooperate,cooperate) $= u_2$(cooperate,cooperate) $=$ -2 which is preferable for both. The power of Game Theory lies in it's ability to justify poor outcomes caused by perfect strategic, selfish play.

In general, games could have more than one PNE or none at all. Also, there is good reason to suspect that PNE's may not be tractable [DGP09], making the assumption that players will reach them less reasonable and motivating the formulation of pore permissive solution concepts.

2.1.1.2 Mixed Nash Equilibria

Consider the classic 2 player game of rock-paper-scissors. Two players chose simultaneously either rock, papers or scissors and the winner is declared according to the following rules:

- Rock beats Scissors
- Scissors beats Paper
- Paper beats Rock

This game can be modeled as follows:

	Rock	Paper	Scissors
Rock	0,0	0,1	1,0
Paper	1,0	0,0	0,1
Scissors	0,1	1,0	0,0

This game has no PNE because for any outcome s either the row or column player does not win but can change his action and beat his opponent.

Mathematically, the problem is that players do not have fine control over the outcome of the game and changing their actions causes them to 'overshoot' the equilibrium. A solution is to convexify their set of actions using probability. In this setting, every player i chooses probability distribution $\boldsymbol{\sigma_i}$ over his strategies S_i. This leads to the following equilibrium definition.

Definition 2 (Mixed Nash Equilibrium (MNE)). *A product distribution over outcomes $\boldsymbol{\sigma} = \times_i \sigma_i$ is a mixed Nash equilibrium (MNE) if for each player i:*

$$\mathbb{E}_{\mathbf{s}_{-i} \sim \boldsymbol{\sigma}_{-i}}[u_i(s'_i, \mathbf{s}_{-i})] \leq \mathbb{E}_{\mathbf{s} \sim \boldsymbol{\sigma}}[u_i(\mathbf{s})]$$

for any $s'_i \in S_i$

In this setting, rock-paper-scissors does have an equilibrium $\boldsymbol{\sigma}_i = (1/3, 1/3, 1/3)$ for both players. Knowing the opponent will choose an action uniformly at random, any strategy chosen has expected utility equal to 1/3.

Fortunately, MNE's are guaranteed to exist, as famously proven by Nash in 1951 [Nas51].

Theorem 1. *Any game with a finite number of players choosing among a finite number of strategies has at least one MNE*

2.1.1.3 Correlated and Coarse Correlated Nash Equilibria

Despite the generality of the MNE, there still exist certain situations where we would like more flexibility, as is apparent from the following game.

	Cross	Stop
Cross	-100,-100	+1,0
Stop	0,+1	0,0

This game corresponds to two vehicles at a crossroad, where only one can cross at a time. This game has two PNE's, (Cross,Stop) and (Stop,Cross). In both of them however, one player knows his payoff will be poor, even though the game is completely symmetrical. The MNE, where each player crosses with probability $\frac{1}{101}$ doesn't fare much better because most of the time no car will cross the road, meaning very small expected utility for both players.

This situation could be easily rectified by having a coordinator guide the players, or in game theoretic terms, have a *correlated equilibrium (CE)*, introduced in [Aum74]. In a correlated equilibrium, we can image the coordinator picking randomly strategies and assigning them to the players.

Definition 3 (CE). *A probability distribution $\boldsymbol{\sigma}$ over outcomes is a CE if for any player i:*

$$\mathbb{E}_{\mathbf{s}_{-i}\sim\boldsymbol{\sigma}_{-i}}[u_i(s'_i, \mathbf{s}_{-i})|(s_i)] \leq \mathbb{E}_{\mathbf{s}\sim\boldsymbol{\sigma}}[u_i(\mathbf{s})]$$

for any $s'_i \in S_i$. Note that $\boldsymbol{\sigma}$ needs not be *a product distribution.*

In the previous game, there exists a correlated equilibrium where the coordinator chooses strategies (Cross,Stop) and (Stop,Cross) uniformly at random. The expected utility of each player is 0.5 (significantly better that the MNE) and the game is still completely symmetric.

It is important to note the following relation:

$$\text{PNE} \subset \text{MNE} \subset \text{CE}$$

which is fairly obvious, considering each time we relax the constraints over which players are assigned strategies.

There is one last interesting equilibrium concept we need to note, even though we will only briefly use it. It is the coarse correlated equilibrium (CCE) and it is similar to CE, where players are not signaled about their strategy. Intuitively, one can think of the difference between CE and CCE as:

- In a CE you agree to play the game, then you are assigned a strategy and then you decide what to play knowing what the coordinator proposed

- In a CCE you can either deviate up front or participate in the game and immediately play whatever the coordinator proposes to you

Formally, the definition is:

Definition 4 (CCE). *A probability distribution $\boldsymbol{\sigma}$ over outcomes is a CCE if for any player i:*

$$\mathbb{E}_{\mathbf{s}_{-i}\sim\boldsymbol{\sigma}_{-i}}[u_i(s'_i, \mathbf{s}_{-i})] \leq \mathbb{E}_{\mathbf{s}\sim\boldsymbol{\sigma}}[u_i(\mathbf{s})]$$

for any $s'_i \in S_i$.

We also have that CE $\subset$ CCE.

2.1.2 Price of Anarchy

Up to now we have defined some solution concepts and agreed that strategic players will be able to reach them. However, we do not have any property of these equilibria in order to differentiate between them. Intuitively, we can sense that some equilibria are worse than others, but it is time to formalize this notion.

As every player has his utility function, we define a 'global' utility function of the game, called either the welfare or the cost, depending on whether we are trying to maximize or minimize it, defined by $W : S \to \mathbb{R}$. Usual candidates for the welfare maximization objective are:

- The social surplus objective: $\sum_{i=1}^{n} u_i(s)$, where we sum each players utility
- The egalitarian objective: $\min_{i=1}^{n} u_i(s)$, where we try to make sure all players gain at least some utility.

A cost minimization objective might look like $\sum_{i=1}^{n} c_i(s)$ which is similar to the previous expression, only we are trying to minimize it (and the players want to minimize their cost c_i).

Of course we are mostly interested in the welfare of either some equilibrium or the optimal outcome. The most common measure to compare the efficiency of different equilibria is the Price of Anarchy (or POA) defined by Koutsoupias and Papadimitriou in [KP09].

Definition 5. *For a welfare maximization game $G = (N, S, u)$, a welfare function $W : S \to R$ and equilibrium concept E the POA is:*

$$POA = \frac{max_{\mathbf{s} \in S} W(\mathbf{s})}{min_{\mathbf{s} \in E} W(\mathbf{s})}$$

Equivalently, for cost minimization we have:

$$POA = \frac{max_{\mathbf{s} \in E} C(\mathbf{s})}{min_{\mathbf{s} \in S} C(\mathbf{s})}$$

To get a feel for the POA (and inefficiency in general), we present the following scenario, known as the pollution game.

- We have n players
- Each player has two actions: pollute or be environmentally friendly
- The cost of each player i is $c_i(s) = \sum_{j \neq i : s_j = \text{pollute}} 1 + 3e_i$ where $e_i = 1$ if he is eco-friendly, else $e_i = 0$.

In other words, if a player pollutes he increases everyone else's cost by 1 but if he is eco-friendly he only increases his own by 3.

The optimal outcome is for every player to remain eco-friendly, incurring a welfare of $\sum_i u_i(s) = 3n$. However, the only PNE of this game is for everyone to pollute, since changing to pollution increases utility by 3 (decreasing the other player's utilities does not matter to selfish players). The welfare of this PNE is $\sum_i u_i(s') = n(n-1)$ leading to a POA of $\frac{3n}{n(n-1)} = \frac{3}{n-1}$

Armed with these tools, we are ready to tackle various strategic scenarios where selfish players compete for the better outcome.

2.1.3 Extensive From Games

This is a minor detour from the flow of this chapter, as we will only use extensive form games for one theorem later on.

In the games defined previously, everything happens at once, which is in contrast to most of the 'games' played which usually involve more than one step. Think for example of playing chess or buying stocks one day after another for some time. In both these situations strategies are dynamic and depend on the flow of the game.

Extensive form games capture this element of having multiple rounds. Informally, think of an extensive form game as a tree, where players take turn and decide which node to descend to. This is represented by:

- A finite set of players $\{1, \ldots, n\}$
- A finite set of nodes X which form the game tree. A set $Z \subseteq X$ contains the terminal nodes.
- A set of functions:
  - $i(x)$: determines which player chooses an action at node x
  - $A(x)$: the set of possible actions at x
  - $n(x, a)$: the node to move to having chosen action a at node x
- Utility functions for each player, as before $u_i : Z \to \mathbb{R}$
- The information partition $h(x)$. For each node x this denotes the set of nodes that are player $i(x)$ could be on, based on the information available to him. It will become clearer with an example.

We also define the set of available information for player i:

$$H_i = \{h(x) : \forall x \in X \text{ with } i(x) = i\}$$

Having that, a player's pure strategy is a function $s_i : H_i \to A_i$ with $\forall x : s_i(h(x)) \in A(h(x))$, which maps his available information to an action. Note that this isn't just one move, but a whole contingency plan for all different paths the game might follow.

Just to clarify, we will present rock-paper-scissors as a turn based game. The nodes are $\{\text{start}, R, P, S, RR, RP, RS, PR, PP, PS, SR, SP, SS\}$ denoting the move of the first and second player. The first player begins and has $h(\text{start}) = \text{start}$ as it's the beginning of the game. The second player however, might have $h(R) = h(P) = h(S) = \{R, P, S\}$ or $h(R) = R, h(P) = P$ and $h(S) = S$. In the first case he does not know what player 1 has chosen since his information set shows he could be on either node. In the second case however, he knows exactly where he is on the game tree and can beat player 1.

The PNE (and MNE) is described in exactly the same way, only instead of single strategies we have these strategy functions. The only issue with the PNE is that if a player decides to deviate, the rest of the game is up for grabs, since the PNE only describes one path of the game tree.

The other important solution concept of extensive form games is the *Subgame Perfect Equilibrium*, first described as an analog to the PNE by nobel laureate Reinhard Selten. Image being on node y. From now till the end of the game, what happened before does not matter. It's as if there is a new extensive form game starting from that node. We call this an *induced subgame*.

Definition 6. *Let G be an extensive form game. A subgame G' of G consists of a subset Y of the nodes of X created by a non-terminal node y and all of it's successors which has the property that if $q \in Y$, $q' \in h(q)$ then $q' in Y$. The information sets, actions, terminal nodes and payoffs remain the same.*

A subgame perfect equilibrium is a subset of the PNE that guarantees that every possible subgame is also in a PNE.

Definition 7 (Subgame Perfect Equilibrium (SPE))**.** *A strategy profile s is in a subgame perfect equilibrium if it induces a PNE in every subgame G'.*

Since a game is a subgame of itself, every SPE is also a PNE.

The POA for games in extensive form is exactly the same as in one-shot games. Also, although extensive form games are an obvious superset of normal form games (just include every node in each information set, thus players do not know anything when choosing their action), we can also convert them to one shot games, albeit with an enormous and complicated set of strategies. However, knowing the terminology of extensive form games gives us the ability to construct more detailed arguments, especially when discussing subgame perfect equilibria.

2.2 Introduction to Mechanism Design

The goal of this thesis is twofold. Besides showing how certain techniques from the modern algorithmic toolbox are used for analyzing games, we are also interested in designing games with desirable properties ourselves.

As Tim Roughgarden often says, Mechanism Design is the science of rule making. One can think of it as reverse Game Theory. Instead of having a game found in the wild handed to us for studying, we design our own one on top of some strategic foundation, according to our goals.

2.2.1 Mechanism Design Setting

A mechanism design setting is a set $(N, \mathcal{X}, \mathcal{V})$ where:

- $N = \{1, \ldots, n\}$ is the set of players
- $\mathcal{X} \subseteq \times_i \mathcal{X}_i$ is the set of outcomes
- $\mathcal{V} = \times_i \mathcal{V}_i$ is the set of valuations $v_i : \mathcal{X}_i \to \mathbb{R} \in \mathcal{V}_i$ for each player.

This setting just makes a connection between players and desirable outcomes without mentioning the actions through which the players will reach them. The valuation of an outcome is similar to the utility function and signifies how much a player wants an outcome. For example, we could model an auction with one item a by taking n players, each of them with outcomes $\mathcal{X}_i = \emptyset, a$ and valuations $v_i(\emptyset) = 0, v_i(a) = v_i$ where the overall outcome $\mathcal{X}$ only contains outcomes where at most one player gets the item. How the players might obtain the item or how much they will pay for it is irrelevant to the setting and can be imposed by the mechanism designed later. The setting only describes a social scenario.

2.2.2 Mechanism Design

In order to make a strategic game out of the previous setting we need to complete our model with some interaction from the players. We complement the previous setting with a mechanism $M = (\mathcal{A}, X, P)$ where:

- $\mathcal{A} = \times_i \mathcal{A}_i$ is the set of actions of each player
- $X : \mathcal{A} \to \mathcal{X}$ is the *allocation* function. X_i is the allocation for each player.
- $P : \mathcal{X} \to \mathbb{R}^n$ is the payment function, with P_i for each player.

We assume that after payments, each player has utility

$$u_i^{v_i}(x_i, p_i) = v_i(x_i) - p_i$$

This preference is called *quasilinear*, since it's first term is a function of the outcome (but the actual value is determined by the player) and the second term produces a linear change in utility through payments, which we can control. This gives the designed great influence over the outcome choses by the players.

As an example, consider the auction setting described in the previous section. One possible mechanism is the first price auction, which we will study in detail in the following chapters. In a first price auction, every player submits a hidden bid and the auctioneer awards the top bidder the item and charges him his bid. Formally:

- Players can bid any nonnegative value $b_i \in A_i = \mathbb{R}^+$.
- Only the highest bidder gets the item:

$$x_i = \begin{cases} a & b_i = \max_i b_i \\ 0 & \text{otherwise} \end{cases}$$

- Only the winner pays his bid

$$p_i = \begin{cases} b_i & i = j \\ 0 & \text{otherwise} \end{cases}$$

A different formulation could be what is called the *second* price auction. Here the allocation function remains the same, but winner pays the second highest bid $p_j = \max_{i \neq j} b_i$.

Both these mechanisms imply strategic games as defined previously. This is because the utility of the players is implicitly a function of their actions: $u_i(x) = u_i(X(a)) = u'_i(a)$, thus an n player mechanism induces the game $G = (N, \mathcal{A}, u')$ which accepts equilibria as normal.

For example, in the second price auction one PNE is when everyone bids his actual valuation v_i.

Theorem 2. *The outcome $a = (v_1, \ldots, v_n)$ is a PNE of the second price auction*

Proof. Let i be any player and b_1, b_2 the bids above and below his respectively. The only way he can influence the outcome is by increasing or decreasing his bid enough to change the relative order of bids. If he bids $b'_i > b_2$, he will either still lose the auction and leave his utility unchanged or get $u_i(b'_i, b_2) = v_1 - b_2 < 0$. If he bids $b'_i < b_1$ then he will lose for utility 0. In any case, he cannot gain positive utility by changing his bid, thus bidding $v_i = b_i$ is a PNE. □

The first price auction does not have a PNE like that, since if the top two bids are $b_1 > b_2$, the winner can bid $b'_1 = \frac{b_1+b_2}{2}$ to pay less and increase his utility. Actually, to be absolutely precise, there could be PNE if for example two players have the same valuation $v_i = v_j$ and bid $b_i = b_j = v_i$. In this case, as long as the winning player is picked consistently no player can raise his utility by deviating. As a side note, in the first price auction you cannot gain utility by bidding your actual valuation.

With mechanisms we can still use the POA as an inefficiency metric. In auctions our goal is to maximize the social welfare which is usually defined as $SW(a) = \sum_i v_i(X(a))$. Sometimes it is also defined as the sum of valuations plus the payments which cancel out, since we can consider payments as 'gains' of the auctioneer. As we can see, the second price auction is has a POA of 1, since the PNE is optimal as the highest bidder also has the highest valuation.

2.2.3 Mechanism Design and Incomplete Information

The mechanisms we have described until now are not particularly realistic, especially in the auction setting, because the players' valuations is public information. Without getting into to much detail, one way to solve this is to use probability distributions over valuations. In this setting each player has a distribution $\mathcal{F}_i \sim \mathcal{V}_i$ over his valuations. This distribution is public information, instead of the actual valuation. The mechanism definition is left unchanged, but since players don't know their valuation beforehand their actions are encoded in the form of functions $s_i : \mathcal{V}_i \to \mathcal{A}_i$. The standard solution concept for these games is the Bayes-Nash equilibrium (BNE), defined as follows:

Definition 8 (BNE)**.**

$$\forall a_i \in \mathcal{A}_i : \mathbb{E}_v[u_i(s(v))] \geq \mathbb{E}_v[u_i(a_i, s_{-i}(v)_{-i})]$$

As expected, the social welfare of an outcome s is defined in expectation:

$$SW(s) = \sum_i \mathbb{E}_v u_i(s(v))$$

with the POA being similar.

In this setting, the second price auction retains it's equilibria of bidding your own valuation and the POA is still 1. However, the first price auction has a completely non trivial POA of $\frac{e}{e-1}$ [CKST13].

2.2.4 Truthfullness

The last part of mechanism design we will lightly touch upon is truthfulness. We call a mechanism truthful when the players maximize their gain by submitting their true value, instead of strategizing. This is a particularly valuable property because it means that as designers we can be certain about the players actions and how to assign the optimal outcome and the players have a very easy optimal strategy.

Before formally defining truthfulness, we give the definition of an 'optimal' (or *dominant* as it's most commonly called) strategy.

Definition 9. *A strategy s_i is dominant for player i if for every s_i', s_{-i} we have*

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$$

This means that a dominant strategy always guarantees maximum utility, no matter what the other players are doing. If a game has an optimal strategy for every player, the resulting strategy vector is obviously a PNE called a dominant strategy equilibrium.

Technically, we can only define truthfulness for mechanisms where $\mathcal{A}_i = \mathcal{V}_i$, also called *direct revelation* mechanisms. A direct revelation mechanism is truthful if choosing action v_i(the actual valuation) is a dominant strategy. In mechanisms where these sets are different the player can only implicitly be truthful. The second price auction is a truthful game. This is easy to prove, just take a second look at the PNE proof and observe that we did not care about the remaining players' bids. We proved that bidding truthfully was optimal irrespective of the other players actions.

Not every game with a dominant strategy however is truthful. As a toy example, consider the second price auction where the payment is **twice** the second highest bid. Here the dominant strategy is to bid $\frac{v_i}{2}$, which is not truthful. The following important theorem, due to Allan Gibbard [Gib73], clarifies this relationship.

Theorem 3 (Revelation Principle). *For every mechanism M in which every participant has a dominant strategy there is an equivalent direct revelation mechanism M'*

Proof. Since mechanism M has a dominant strategy equilibrium, assume $s_i : \mathcal{V}_i \rightarrow \mathcal{A}_i$ to be each players strategy (we kept the function notation because this applies to the incomplete information setting as well).

Mechanism M' is direct revelation and works in the following way. It accepts players reported valuations (which may be untrue) and then feeds actions $s_i(v_i)$ to mechanism M and forwards the resulting allocation and payments to the players. Thus, reporting the true valuation is still a dominant strategy and M' is truthful. □

In other words, we can delegate choosing the dominant strategy to a direct revelation mechanism.

We will use dominant strategies in more complicated settings such as the combinatorial auction. Here we have n players, a set U of m items and each player has a valuation $v_i : 2^U \to \mathbb{R}$ over all possible allocations. A truthful mechanism would require exponential communication from the players which is just not feasible in practice. An alternative would be to run independent single item auctions for each item, which as we will see remains quite efficient.

Chapter 3

Smooth Games

3.1 Overview

In many situations POA bounds are not only exceedingly difficult to calculate but might not offer a great deal of insight. This could be because of the absence of specific equilibria or because the assumption that the players will be able to reach them may not always be reasonable. As such, alternative tools that provide expressiveness (in terms of bounds) in these conditions may be required. Smooth Games, introduced and proven to be practical by Tim Roughgarden in [Rou09], are such a toolset.

Before getting into formal definitions, Smooth Games utilize the age-old mathematical technique of restricting the scope in favor of structure. Smooth Games are a natural subset of cost minimization (or maximization) games which entail bounds to configurations much more expressive than the Nash equilibrium.

Essentially, Smooth Games capitalize on the robust nature many important games exhibit. By calling a game robust, we mean that certain worst-case bounds similar to the POA hold even when players haven't reached a Nash equilibrium. Moreover, Smooth Games are very useful in a more practical sense: bounds acquired via smoothness techniques extend to various common equilibrium concepts, as we will soon see. In some cases smoothness bounds bind the POA tightly.

3.2 Defining Smooth Games

Just to clarify, we will be using cost minimization games in order to showcase smoothness arguments. By cost minimization game we mean a game where the goal of each player is to minimize his own cost while the goal of the designer is to minimize the joint cost function $C(s) = \sum_{i=1}^{k} C_i(s)$.

Definition 10. *A cost minimization game is* $(\lambda\text{-}\mu)$*-smooth if for every two outcomes* $\mathbf{s}$ *and* $\mathbf{s}^*$ *we have:*

$$\sum_{i=1}^{k} C_i(s_{i*}, \mathbf{s_{-i}}) \leq \lambda \cdot C(\mathbf{s}^*) + \mu \cdot C(\mathbf{s}) \tag{3.1}$$

What we have achieved here is quantifying the difference between any two outcomes by using the cost of unilateral deviation. In contrast, most equilibrium concepts give us a weaker amount of knowledge: the Nash equilibrium is only an inequality between one specific outcome and all the others. The term 'smooth' now becomes apparent. In smooth games, the social cost of unilateral deviation is nicely bounded and represents a 'smooth' transition between two outcomes.

The cornerstone of this chapter is following short but deep proof that uses smoothness to give an upper bound on the POA.

Theorem 4. *If a game is $(\lambda$-$\mu)$-smooth with $\lambda > 0$ and $\mu < 1$ then each of it's pure Nash equilibria $\mathbf{s}$ has cost at most $\lambda/(1-\mu)$ times that of an optimal solution $\mathbf{s}^*$*

Proof.

$$C(\mathbf{s}) = \sum_{i=1}^{k} C_i(\mathbf{s}) \tag{3.2}$$

$$\leq \sum_{i=1}^{k} C_i(s_i^*, \mathbf{s}_{-i}) \tag{3.3}$$

$$\leq \lambda \cdot C(\mathbf{s}^*) + \mu \cdot C(\mathbf{s}) \tag{3.4}$$

Inequality (3.2) follows from the definition of social cost; inequality (3.3) comes from applying the Nash equilibrium hypothesis once for each player for deviation s_i^* and inequality (3.4) comes from the definition of smoothness. Rearranging terms gives us the claimed inequality □

During this proof we did not exploit our definition to the fullest. For example, we did not use any combination of outcomes and we used equality for our objective function. However, giving a more general definition for smoothness is what allows us to use this bound beyond pure Nash equilibria.

Note that even though a POA bound derived by this proof would give us an inequality between the pure Nash equilibrium and any other outcome, we 'proved' something more. Using any two outcomes and joining them by the cost of unilateral deviation is what makes this proof reveal interesting structural properties of the game, beyond pure Nash equilibria.

For completeness the definition of smoothness for payoff-maximization is:

$$\sum_{i=1}^{k} \pi_i(s_{i*}, \mathbf{s_{-i}}) \geq \lambda \cdot V(\mathbf{s}^*) - \mu \cdot V(\mathbf{s}) \tag{3.5}$$

Here $V(\cdot)$ is an objective function that satisfies $V(\mathbf{s}) \geq \sum_{i=1}^{k} \pi_i(\mathbf{s})$, or in other words it is payoff dominated. Similarly, the POA for such games is $\frac{\lambda}{1+\mu}$.

We formally define a lower bound on what can be proven by smoothness arguments.

Definition 11 (Robust POA)**.** *The robust price of anarchy of cost-minimization game is*

$$\inf \left\{ \frac{\lambda}{1-\mu} : (\lambda, \mu) \text{ such that the game is } (\lambda, \mu)\text{-smooth} \right\}$$

This definition allows us to express the quality of smoothness argument bounds with a single number and compare it to the regular POA.

Before moving on to specific examples, a remark may be in order. First of all, there are classes of games, like valid utility games, which are smooth, even though they may have no pure Nash equilibria. Essentially, the smoothness argument provides us with POA bounds *if* a pure Nash equilibrium exists. Combined with the extensibility of bounds proven by smoothness arguments, we see that Smooth Games are not about just PNE's but give insights that are used alongside the standard tools for evaluating inefficiency of equilibria.

3.2.1 Examples

One class of games that we will be using throughout this chapter is the Congestion Game, introduced by Rosenthal in [Ros73] and at first studied due to the fact that they always had a PNE, despite their generality at modelling various strategic sharing scenarios in society. The Congestion Game is a cost-minimization game defined by a set E of resources, a set of k players with strategy sets $S_1, \dots S_k \subseteq 2^E$ and a cost function $c_e : \mathbb{Z}^+ \to \mathbb{R}$. We will assume these functions are nonnegative and nondecreasing. Given a strategy profile $\mathbf{s} = (s_1, \dots s_k)$ with $s_i \in S_i$ for each i we define the load of each resource as the number of players using it: $x_e = |\{i : e \in s_i\}|$. The cost for each player i is $C_i(\mathbf{s}) = \sum_{e \in s_i} c_e(x_e)$. The total cost is:

$$C(\mathbf{s}) = \sum_{i=1}^{k} C_i(\mathbf{s}) = \sum_{e \in E} c_e(x_e) \cdot x_e$$

Intuitively, congestions games model situations where players must share a set of resources which in this case get more expensive the more they are used.

Example 3.2.1 (Congestion Game with Affine Cost Functions)**.** We will use the smoothness framework to derive POA bounds for the case where cost functions are affine: $c_e(x) = a_e \cdot x + b_e$. This case has been studied by Koutsoupias and Christodoulou in [CK05]. Using the lemma:

$$x(y+1) \leq \frac{5}{3}x^2 + \frac{1}{3}y^2$$

which is true for all nonnegative integers x, y we will show that congestion games with affine cost functions are $\frac{5}{3}, \frac{1}{3}$ smooth.

Using the lemma, we have:

$$\begin{aligned} ax(y+1) + bx &\leq \frac{5}{3}ax^2 + \frac{1}{3}ay^2 + bx \\ &\leq \frac{5}{3}ax^2 + \frac{1}{3}ay^2 + \frac{5}{3}bx + \frac{1}{3}by \\ &\leq \frac{5}{3}(ax^2 + bx) + \frac{1}{3}(ay^2 + by) \end{aligned}$$

for all $a, b \geq 0$.

Consider two outcomes $\mathbf{s}, \mathbf{s}^*$ with induced loads $\mathbf{x}, \mathbf{x}^*$. For outcome $(s_i^*, \mathbf{s_{-i}})$ the number of players using each resource can increase by at most one. Adding everything together:

$$\sum_{i=1}^{k} C_i(s_i^*, \mathbf{s_{-i}}) \leq \sum_{i=1}^{k} \sum_{e \in s_i^*} c_e(x_e + 1) \tag{3.6}$$

$$\leq \sum_{e \in E} (a_e(x_e + 1) + b_e)x_e^* \tag{3.7}$$

$$\leq \sum_{e \in E} \frac{5}{3}(a_e x_e^* + b_e)x_e^* + \sum_{e \in E} \frac{1}{3}(a_e x_e + b_e)x_e \tag{3.8}$$

$$\leq \frac{5}{3}C(\mathbf{s}^*) + \frac{1}{3}C(\mathbf{s}) \tag{3.9}$$

Where (3.6) and (3.7) are just writing out the objective function and doing a reversal and regrouping of sums and (3.8) is utilizing the lemma to split the two outcomes into smooth game form. Thus, by smoothness the POA of affine congestion games is at most $\frac{\lambda}{1-\mu} = \frac{5}{2}$.

Example 3.2.2 (Location Game)**.** To show a different approach, our next example is one where the goal is payoff maximization. The Location Game was designed by Tim Roughgarden and is based on the *valid utility game*[Vet02].

The Location game is defined by:

- A set F of possible locations
- A set of k players, where player i chooses only one among $F_i \subset F$ locations.
- A set M of markets. Each market has a value v_j which is known to all the players. Every players 'sells' the same goods to each market. The market only buys from the player who sells the cheapest according to market equilibrium. More on that later.
- For each location $i \in F$ and market $j \in M$ there is a 'distance' cost c_{ij}.

Before defining the payoffs we will create a toy example to clarify how the prices are agreed.

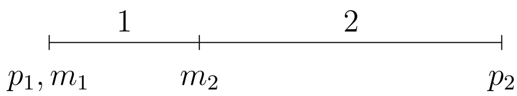

Here we have 2 players, p_1,p_2 and markets m_1,m_2 located on this line segment with their respective distances. Assuming all markets buy at a cost of 3, at which price would the players sell?

At m_1 player 1 has zero distance while p_2 has $c_{21} = 3$. Player 2 cannot sell to m_1 at any positive price, so p_1 sells for the maximum he can, which is 3. At m_2 however, things are more interesting. Player 1 has distance 1 and player 2 has distance 2.

This means that in order to be profitable every player must sell his goods for more that the cost of distance. Which means that p_1 sells for 1 and p_2 sells for 2. Player 1 of course maximizes his profits by selling for $2-\epsilon$ and the market buys from him. When designing the actual payoff function we will disregard this ϵ because it does not add any meaningful information.What we can see is that the only player who sells at a market is the one located closest to it. The location of the other players defines his price.

In general, in a strategy profile $\mathbf{s}$ of a location game, the payoff of player i is defined as:

$$\pi_i(\mathbf{s}) = \sum_{j \in M} \pi_{ij}(\mathbf{s}) \tag{3.10}$$

where, assuming C is the set of chosen locations (by all the players) and i chose $j \in C$ we have:

$$\pi_{ij}(\mathbf{s}) = \begin{cases} 0 & if c_{lj} > v_j \text{ or } l \text{ is not the closest location of } C \text{ to } i \\ \delta_j^{(2)}(\mathbf{s}) - c_{ij} & otherwise \end{cases} \tag{3.11}$$

Where $\delta_j^{(2)}(\mathbf{s})$ is equal to the minimum of v_J and the second smallest distance between a location in C and market j. Following our example, since p_2 has a distance of 2 to m_2, p_1 can sell his goods for 2 and beat him. There is a subtlety here when players both are at an equal minimal distance from a market. In this case both players get *nothing*. If this sounds unnatural, it can be considered a subgame perfect equilibrium were prices are sequentially undercut until they reach 0.

The maximization objective in this setting is the social surplus, which we define:

$$V(\mathbf{s}) = \sum_{j \in M} \max\{v_j - d_j(\mathbf{s}), 0\} \tag{3.12}$$

where $d_j(\mathbf{s})$ is the smallest distance. The *max* is somewhat cumbersome, especially when all we are trying to write is that every market is served by the closest location unless the distance is higher than v_j. To slightly abuse the notation for convenience we will use $V(T)$ for the surplus when players occupy locations $T \subset F$. Each player is interchangeable and $V(\mathbf{s})$ only uses the strategy profile to find these positions.

We start by proving some simple properties of this game.

P1 For any strategy profile $\mathbf{s}$ we have

$$\sum_{i=1}^{k} \pi_i(\mathbf{s}) \leq V(\mathbf{s}) \tag{3.13}$$

This follows because every market $j \in M$ serviced by player i gives him payoff $d_j^{(2)}(\mathbf{s}) - d_j(\mathbf{s})$ but adds $v_j - d_j(\mathbf{s})$ to the social surplus. By definition, $v_j \geq d_j^{(2)}(\mathbf{s})$

P2 For any strategy profile $\mathbf{s}$ we have

$$\pi_i(\mathbf{s}) = V(\mathbf{s}) - V(\mathbf{s}_{-i}) \tag{3.14}$$

There are two cases. If market j is not used by any player in $\mathbf{s}_{-i}$ but the new player uses it, then its contribution to both sides of the equation is $v_j - d_j(\mathbf{s})$.

If market j changes user when the player is added, the payoff for the player is $d_j^{(2)}(\mathbf{s}) - d_j(\mathbf{s})$ since the new player is closer. The difference in social surplus is $v_j - d_j(\mathbf{s}) - (v_j - d_j^{(2)}(\mathbf{s}))$ which is exactly the same. Summing over all markets gives us the desired equation.

P3 The function $V(\cdot)$ is monotonic and submodular. Monotonicity means that $V(T_1) \leq V(T_2)$ for $T_1 \subset T_2$. This follows immediately from (3.12).

Submodularity models diminishing returns and is defined as

$$V(T_2 \cup \{l\}) - V(T_2) \leq V(T_1 \cup \{l\}) - V(T_1)$$

for all $l \in C$ and $T_1 \subset T_2$. The proof is similar to the previous property.

Theorem 5. *The location game is* $(1,1)$*-smooth.*

Proof. Our proof will follow the standard smoothness argument procedure, which is to write the PNE hypothesis once for each player, not using again and disentangle the payoffs for two outcomes using the previous properties.

We begin by the letting $\mathbf{s}$ be the PNE and $\mathbf{s}^*$ any other outcome. We have

$$\pi_i(\mathbf{s}) \geq \pi_i(s_i^*, \mathbf{s}_{-i}) \tag{3.15}$$

Summing over all players and using (3.13) we have:

$$\begin{aligned}
V(\mathbf{s}) &\geq \sum_{i=1}^{k} \pi_i(\mathbf{s}) && (3.16)\\
&\geq \sum_{i=1}^{k} \pi_i(s_i^*, \mathbf{s}_{-i}) && (3.17)\\
&= \sum_{i=1}^{k} [V(s_i^*, \mathbf{s}_{-i}) - V(\mathbf{s}_{-i})] && (3.18)\\
&\geq \sum_{i=1}^{k} [V(s_1^*, \dots, s_i^*, \mathbf{s}_{-i}) - V(s_1^*, \dots, s_{i-1}^*, \mathbf{s}_{-i})] && (3.19)\\
&= V(s_1^*, \dots, s_k^*, s_1, \dots, s_k) - V(\mathbf{s}) && (3.20)\\
&\geq V(\mathbf{s}^*) - V(\mathbf{s}) && (3.21)
\end{aligned}$$

Where (3.18) from (3.14), (3.19) from submodularity, (3.20) from rolling out the telescoping sum and (3.21) from monotonicity.

Solving for the POA, we get $\frac{1}{2}$. □

3.3 Extension Theorems

One extra reason to use smoothness techniques when studying games is that bounds proven by the standardized smoothness argument automatically extend to more general equilibrium concepts. This is quite useful, since analysis of a game usually starts by pure Nash equilibra which often are specific and easier to think about, with the logical next step being moving on to mixed Nash equilibria and so forth.

Game Theory is quite intricate and we know of many games where the mixed Nash equilibrium is worse than the pure. This means that there can never be an extension theorem using just the POA without other properties of the game. Smooth Games automatically have a POA bound that does not depend on the game itself directly, but on the two smoothness parameters (λ, μ). This gives us an extra layer of abstraction that makes POA proofs reusable.

The intuition behind the extension theorems lies inside the canonical smoothness bound proof, which can be succinctly described in list fashion:

- Choose an outcome $\mathbf{s}$ of the game. It doesn't need to be the optimal.
- Invoke the Nash equilibrium hypothesis once for each player (for the previous outcome) without bringing extra inequalities into the proof that might be specific to the equilibrium or the outcome.
- Combine the previous inequalities, along with the definition of smoothness to come up with a bound

Bounds proven this way *do not* make use of any specific equilibria properties or configurations and are very robust to extensions and alterations.

3.3.1 One-Shot Games

The first extension theorem concerns randomized equilibrium concepts. Just a very quick reminder mixed Nash,correlated and coarse correlated equilibria to show how they generalize:

$$\mathbb{E}_{\mathbf{s}\sim\sigma}\left[C_i(\mathbf{s})\right] \leq \mathbb{E}_{\mathbf{s}_{-i}\sim\sigma_{-i}}\left[C_i(s_i', \mathbf{s})\right] \tag{3.22}$$

$$\mathbb{E}_{\mathbf{s}\sim\sigma}\left[C_i(\mathbf{s})|s_i\right] \leq \mathbb{E}_{\mathbf{s}_{-i}\sim\sigma_{-i}}\left[C_i(s_i', \mathbf{s})|s_i\right] \tag{3.23}$$

$$\mathbb{E}_{\mathbf{s}\sim\sigma}\left[C_i(\mathbf{s})\right] \leq \mathbb{E}_{\mathbf{s}_{-i}\sim\sigma_{-i}}\left[C_i(s_i', \mathbf{s})\right] \tag{3.24}$$

Where in (3.22) $\sigma = (\sigma_1, \ldots \sigma_k)$ is a product distribution. We only need to show the extension theorem for coarse correlated equilibria, since the same bound applies for the rest.

Theorem 6 (Extension Theorem - One-shot)**.** *For every cost minimization game G with robust POA $\rho(G)$, every coarse correlated equilibrium σ of G and every outcome $\mathbf{s}^*$ of G we have*

$$\mathbb{E}_{\mathbf{s}\sim\sigma}[C(\mathbf{s})] \leq \rho(G) \cdot C(\mathbf{s}^*)$$

Proof. The proof is quite straightforward, similar to the first theorem:

$$\mathbb{E}_{\mathbf{s}\sim\sigma}[C(\mathbf{s})] = \mathbb{E}_{\mathbf{s}\sim\sigma}\left[\sum_{i=1}^{k} C_i(\mathbf{s})\right] \tag{3.25}$$

$$= \sum_{i=1}^{k} \left[\mathbb{E}_{\mathbf{s}\sim\sigma} C_i(\mathbf{s})\right] \tag{3.26}$$

$$\leq \sum_{i=1}^{k} \left[\mathbb{E}_{\mathbf{s}\sim\sigma} C_i(s_i^*, \mathbf{s_{-i}})\right] \tag{3.27}$$

$$= \mathbb{E}_{\mathbf{s}\sim\sigma}\left[\sum_{i=1}^{k} C_i(s_i^*, \mathbf{s})\right] \tag{3.28}$$

$$\leq \mathbb{E}_{\mathbf{s}\sim\sigma}\left[\lambda \cdot C(\mathbf{s}^*) + \mu \cdot C(\mathbf{s})\right] \tag{3.29}$$

$$\leq \lambda \cdot C(\mathbf{s}^*) + \mathbb{E}_{\mathbf{s}\sim\sigma}\left[\mu \cdot C(\mathbf{s})\right] \tag{3.30}$$

Where (3.26) and (3.28) by linearity of expectation, (3.27) by the definition of the coarse correlated equilibrium and the inequality (3.29) by (λ, μ)-smoothness.

Inequality (3.30) holds for every (λ, μ) for which the game is smooth. We can therefore rearrange terms and reach (or get arbitrarily close to) the robust POA, thus proving the theorem. □

3.3.2 Repeated Play and No-Regret Sequences

The previous extension theorem can be easily proven for sequences of outcomes with certain properties. Consider sequence of outcomes $\mathbf{s}^1, \ldots, \mathbf{s}^t$ of an (λ, μ) smooth game and $\mathbf{s}^*$ the optimal outcome. For every i, t we can define:

$$\delta_i(\mathbf{s}^t) = C_i(\mathbf{s}^t) - C_i(s_i^*, \mathbf{s}_{-i}^t) \tag{3.31}$$

This is player i's 'improvement' for deviating to $\mathbf{s}^*$. Summing over all players and deploying the definition of smoothness, we get

$$C(\mathbf{s}^t) \leq \frac{\lambda}{1-\mu} \cdot C(\mathbf{s}^*) + \frac{\sum_{i=1}^{k} \delta_i(\mathbf{s}^t)}{1-\mu} \tag{3.32}$$

We are only interested in sequences of outcomes when every player experiences *vanishing average regret*:

$$\frac{1}{T}\sum_{i=1}^{T} C_i(\mathbf{s}^t) \leq \frac{1}{T}\left[\min_{s_i'} \sum_{t=1}^{T} C_i(s_i', \mathbf{s}_{-i})\right] + o(1) \tag{3.33}$$

where $o(1)$ is some function which goes to 0 as $T \to \infty$. Regret is defined as the difference in cost for a certain player had he chosen the best fixed response for $\mathbf{s}^1, \ldots, \mathbf{s}^t$. The previous equation shows vanishing average regret, since playing s^t is asymptotically competitive against any time invariant best response strategy. This can be considered a generalisation of the Nash equilibrium for learning scenarios.

Averaging (3.33) over T steps we get

$$\frac{1}{T}\sum_{i=1}^{T} C_i(\mathbf{s}^t) \leq \frac{\lambda}{1-\mu}\cdot C(\mathbf{s}^*) + \frac{1}{1-\mu}\sum_{i=1}^{k}\left(\frac{1}{T}\sum_{t=1}^{T}\delta_i(\mathbf{s}^t)\right) \tag{3.34}$$

By the no regret guarantee, we have that

$$\frac{1}{T}\sum_{t=1}^{T}\delta_i(\mathbf{s}^t) \to 0$$

as $T \to \infty$. Since this holds for every player, we can restate in more formal fashion.

Theorem 7. *For every cost minimization game G with robust POA $\rho(G)$, every outcome sequence $s^1, \ldots, s^T$ that satisfies* (3.33) *for every player and every outcome $\mathbf{s}^*$ of G,*

$$\frac{1}{T}\sum_{i=1}^{T} C_i(\mathbf{s}^t) \leq [\rho(G) + o(1)]\cdot C(\mathbf{s})$$

as $T \to \infty$

Blum et al. investigated bounds of this type in [BHLR08].

This theorem applies to a much broader scale than it's equivalent on PNE's because there exist online learning algorithms which require much fewer restrictions on the game and it's players than the existence of a Nash equilibrium. Even more, online learning algorithms are often quite simple where a PNE may be intractable. A thorough investigation of online learning algorithms in game theoretic settings can be found in [CBL06].

Even more general than our previous theorem, we can extend smoothness bounds to mixed-strategy no regret sequences which in turn cover an even larger number of settings. In this case we would have

$$\frac{1}{T}\sum_{i=1}^{T} \mathbb{E}_{\mathbf{s}^t\sim\sigma^t} C_i(\mathbf{s}^t) \leq [\rho(G) + o(1)]\cdot C(\mathbf{s})$$

3.4 Tight Classes of Games

A natural question is wether POA bounds that come from the smoothness arguments, a restricted class of proofs, are tight. Before showing certain classes of games for which it is impossible to obtain tight bounds through smoothness, we need to formally define what tight means.

Introducing some new notation, let $\mathcal{G}$ define the set of cost-minimization games with a nonnegative cost function. Let $\mathcal{A}(\mathcal{G})$ denote the set of (λ, μ) values such that every game in $\mathcal{G}$ is (λ, μ)-smooth. Let $\hat{\mathcal{G}} \in \mathcal{G}$ denote the games with at least one PNE and $\rho_{pure}(G)$ the POA of such equilibria in $\mathcal{G} \in \mathcal{G}$. The 3 line smoothness argument (3.2)-(3.4) shows that for every $(\lambda, \mu) \in \mathcal{A}(\mathcal{G})$ and every $G \in \hat{\mathcal{G}}$ we have that $\rho_{pure}(G) \leq \frac{\lambda}{1-\mu}$

Definition 12 (Tight Class of Games). *A set $\mathcal{G}$ of cost-minimization games is tight if:*

$$\sup_{G\in\hat{\mathcal{G}}} \rho_{pure}(G) = \inf_{(\lambda,\mu)\in\mathcal{A}(\mathcal{G})} \frac{\lambda}{1-\mu} \tag{3.35}$$

The right hand side is the best upper bound that can be proven by smoothness argument while the left hand side is the actual worse POA amongst all the games with at least one PNE. Obviously, the right hand side is at least as large as the left, so the condition we have hear is reminiscent of a satisfied min-max. Although there might be specific games which have specific equilibria with better bounds, any proof that is general enough to encompass all instances of games in the class $\hat{\mathcal{G}}$ is at most as good as the best smoothness bound.

Of course, showing that a class of games is tight for PNE's also extends to equilibria we discussed in the previous section. The proof is very simple.

Corollary 7.1. *A tight class of games is also tight for MN, correlated and coarse correlated equilibria*

Proof.

$$\sup_{G\in\hat{\mathcal{G}}} \rho_{pure}(G) \leq sup_{G\in\hat{\mathcal{G}}}\rho_{mixed}(G) \tag{3.36}$$

$$\leq sup_{G\in\hat{\mathcal{G}}}\rho_{correlated}(G) \tag{3.37}$$

$$\leq sup_{G\in\hat{\mathcal{G}}}\rho_{coarse}(G) \tag{3.38}$$

$$\leq \inf_{(\lambda,\mu)\in\mathcal{A}(\mathcal{G})} \frac{\lambda}{1-\mu} \tag{3.39}$$

Where inequalities (3.36)-(3.38) hold because each equilibrium concept is a superset of the previous one.

But, because $\sup_{G\in\hat{\mathcal{G}}} \rho_{pure}(G) = \inf_{(\lambda,\mu)\in\mathcal{A}(\mathcal{G})} \frac{\lambda}{1-\mu}$ by definition of a tight class, we have that all the bounds are equal and tight. □

3.4.1 Congestion Games are Tight

In **Example 3.2.2** we used the smoothness framework to showcase how a POA proof would go. The resulting bound was $\frac{5}{2}$ and although we did not investigate the tightness of this result, it turns out it actually is tight, in terms of **Definition 12**.

Congestion Games can be parametrized in two ways:

- Imposing limitations on the structure of the shared resources
- Restricting the set of allowable cost functions

Previous research in this area by [ADG+06, AAE05, CK05] investigated tight bounds for variants of the original definition of Congestion Games. In all these cases, the bound was defined in relation to the cost functions. For example for polynomial cost functions of maximum degree d and nonnegative coefficients it was found that the worst-case POA bound was exponential in d but independent of the number of players or the way resources were shared. This is not to say that the sets $S_1, \ldots S_k \in 2^E$ that players choose their strategy from have no impact on the POA. They allowed strategies for each player does have an effect, but there is already some structure since all strategies are subsets of the available resources E. However, leaving the class of cost functions $\mathcal{C}$ unrestricted imposes much higher difficulty and it might even unreasonable to expect a worst-case bound for each set $\mathcal{C}$ to be meaningful, let alone expressible in closed form.

We will now show that for any set $\mathcal{C}$, whose functions are nonnegative and nondecreasing, the induced set of congestion games $\mathcal{G}(\mathcal{C})$ forms a tight class. Restated without jargon, knowing that every congestion game has at least one PNE[Ros73]:

$$\sup_{G \in \mathcal{G}(\mathcal{C})} \rho_{pure}(G) = \inf_{(\lambda,\mu) \in \mathcal{A}(\mathcal{G}(\mathcal{C}))} \frac{\lambda}{1-\mu} \tag{3.40}$$

By **Corollary 7.1** this bound will remain tight for those equilibrium concepts as well.

3.4.1.1 Simplifying the Smoothness Constraints

We have already shown that Congestion Games are smooth for affine cost function, now we need to do the same for the general case. The proof presented in this chapter is basically the same as presented in [Rou09] with more details when needed.

We will be using the same properties we used in **Example 3.2.2**, namely:

- The objective cost function and every player's cost function is additive over the resources.
- If one player deviates from the current strategy profile he can at most increase the load of each resource by 1.

This means that, as we did before, we can simplify the range where (λ, μ) parameters of interest are found by considering what happens to one resource.

Let $\mathcal{C}$ be a nonempty set of nondecreasing, nonnegative cost functions. We will disallow the all zero cost function, meaning that $c(x) \geq 0$ for $x \geq 0$. This is done only for convenience, since we can simulate the zero cost function by $c(x) = \epsilon$ without affecting the equilibria or the optimal cost more by than a bounded multiple of ϵ. Let $\mathcal{A}(\mathcal{C})$ (*not* $\mathcal{A}(\mathcal{G}(\mathcal{C}))$, we are trying to find bounds independent of structure) be the set of parameters (λ, μ) with $\mu < 1$ that satisfy:

$$c(x+1)x^* \leq \lambda \cdot c(x^*)x^* + \mu \cdot c(x)x \tag{3.41}$$

for every cost function $c \in \mathcal{C}$ and every nonnegative integer x and every positive integer x^*. The reason for unbounded integers is that we want to be independent of the number of players as well. Only x^* needs to be positive since it represents that one player deviated to it. The induced load of the first outcome, x, can be 0.

Inequality (3.41) is a worst case on the cost change of the unilateral deviation. Since we are hoping our bounds are tight, we expect limiting our (λ, μ) to the worst case won't be a problem.

We define $\gamma(\mathcal{C})$ to be the best POA that can be proved by smoothness for games $\mathcal{G}(\mathcal{C}))$ by the condition (3.41). That is:

$$\gamma(\mathcal{C}) = \inf \left\{ \frac{\lambda}{1-\mu} : (\lambda, \mu) \in \mathcal{A}(\mathcal{C}) \right\} \tag{3.42}$$

For completeness, we define $\gamma(\mathcal{C}) = \infty$ for $\mathcal{A}(\mathcal{C}) = \emptyset$.

We will prove a useful lemma:

Lemma 3.4.1 (Nonnegativity of μ). *For every nonempty set $\mathcal{C}$ of strictly positive functions and every $(\lambda, \mu) \in \mathcal{A}(\mathcal{C}), \mu > 0$*

Proof. Taking $x = n$ and $x^* = 1$ in (3.41) with any cost function $\text{į}n\mathcal{C}$ we get:

$$\mu \leq \frac{c(n+1) - \lambda \cdot c(1)}{c(n)n} \leq \frac{1-\lambda}{n} \tag{3.43}$$

Where we have that c is nondecreasing. Since n can become arbitrarily large we can only assume that $\mu \geq 0$ and nothing more. □

We now need to show that only considering parameters that satisfy (3.41) does not affect the tightness of our bound. More concretely, we will show that any game in $\mathcal{G}(\mathcal{C})$ is smooth for all parameters in $\mathcal{A}(\mathcal{C})$ which in turn implies:

Lemma 3.4.2. *For every set of cost function $\mathcal{C} \neq \emptyset$, the robust POA of every game in $\mathcal{G}(\mathcal{C})$ is at most $\gamma(\mathcal{C})$.*

Proof. If $\gamma(\mathcal{C})$ is not finite, we are done.

For $\gamma(\mathcal{C}) < \infty$, assume $c \in \mathcal{C}$ is not strictly positive. Hence there exists $z \geq 0$ such that $c(z) = 0$ and $c(z+1) > 0$. Substituting $x = x^* = z$ into (3.41) we get $c(z+1) \leq 0$ thus $\mathcal{A}(\mathcal{C}) = \emptyset$ and $\gamma(\mathcal{C}) = \infty$ by definition leading to a contradiction. We can then assume that every $c \in \mathcal{C}$ is strictly positive.

Let $G \in \mathcal{G}(\mathcal{C})$ and $(\lambda, \mu) \in \mathcal{A}(\mathcal{C})$. By the previous assumption and **Lemma 3.4.1** $\mu \geq 0$. For every outcome pair $\mathbf{s}$ and $\mathbf{s}^*$ of G with induced loads $\mathbf{x}$ and $\mathbf{x}^*$ we have:

$$\sum_{i=1}^{k} C_i(s_i^*, \mathbf{s}_{-i}) \leq \sum_{e \in E: x_e^* > 0} c_e(x_e + 1)x_e^* \tag{3.44}$$

$$\leq \sum_{e \in E: x_e^* > 0} [\lambda \cdot c_e(x_e^*)x_e^* + \mu \cdot c_e(x_e)x_e] \tag{3.45}$$

$$\leq \sum_{e \in E} [\lambda \cdot c_e(x_e^*)x_e^* + \mu \cdot c_e(x_e)x_e] \tag{3.46}$$

$$= \lambda \cdot C(\mathbf{s}^*) + \mu \cdot C(\mathbf{s}) \tag{3.47}$$

In inequality (3.44) we are just summing the costs going through the resources instead of the players. From the perspective of the deviating player, when he moves to resource e the cost can at most be $c_e(x_e + 1)$, because the other players stayed put. If x_e^* players deviate to e we get the first inequality.

One may ponder why we sum over $e \in E : x_e^* > 0$ since obviously

$$\sum_{e \in E: x_e^* > 0} c_e(x_e + 1)x_e^* = \sum_{e \in E} c_e(x_e + 1)x_e^*$$

The reason is that we need $x_e^* > 0$ in order to expand to inequality (3.45) by (3.41) and the definition of $\mathcal{A}(\mathcal{C})$. The we use that $\mu \geq 0$ to get to the standard definition of smoothness with (3.47). □

We have now reduced our search for parameters (and the accompanying upper bound) to $\mathcal{A}(\mathcal{C})$ which is only dependent on the cost functions and easier to calculate as an optimization problem for specific sets $\mathcal{C}$.

3.4.1.2 Characterization of the Optimal Smoothness Parameters

After carefully considering (3.41) we conclude that we are dealing with an optimization problem with two parameters, λ and μ. We want to minimize the function $\frac{\lambda}{1-\mu}$ over the feasible region $\mathcal{A}(\mathcal{C})$.

We can see that as λ and μ increase (3.41) continue to hold while our objective function $\frac{\lambda}{1-\mu}$ is increasing. Broadly, the feasible region is similar to the intersection of two halfplanes with equality in (3.41) being satisfied somewhere along the 'southwestern' boundary' of $\mathcal{A}(\mathcal{C})$ for specific x_e and x_e^*.

Right now, we will assume that the set $\mathcal{C}$ is finite, that every cost function is strictly positive and that there is an upper bound on the load of each resource. Formally, let $\mathcal{A}(\mathcal{C}, n)$ be the set of (λ, μ) parameters with $\mu < 1$ that satisfy (3.41) for every $c \in \mathcal{C}$, $x \in \{0, 1 \ldots n\}$ and $x^* \in \{1, 2 \ldots n\}$. As before, we define

$$\gamma(\mathcal{C}, n) = \inf \left\{ \frac{\lambda}{1 - \mu} : (\lambda, \mu) \in \mathcal{A}(\mathcal{C}, n) \right\} \tag{3.48}$$

The set $\mathcal{A}(\mathcal{C}, n)$ is not empty because it contains $\max_{c \in \mathcal{C}}(\frac{c(n+1)}{c(1)}, 0)$.

What we have achieved by bounding and 'discretizing' our problem (but keeping the feasible region continuous) is ability to prove in a mostly hassle free way the existence of points that transform (3.41) into an equality. The next seemingly complicated lemma proves what we argued about, that the bound $\gamma(\mathcal{C}$,$n)$equality usually appears at the intersection of two lines for which inequality (3.41) holds true.

Lemma 3.4.3. *Let $\mathcal{C}$ be a finite set of strictly positive cost functions and n a positive integer. Suppose there exists $(\hat{\lambda}, \hat{\mu}) \in \mathcal{A}(\mathcal{C}, n)$ such that:*

$$\frac{\hat{\lambda}}{1 - \hat{\mu}} = \gamma(\mathcal{C}, n) \tag{3.49}$$

Then there exist $c_1, c_2 \in \mathcal{C}$, $x_1, x_2 \in \{0, 1 \ldots n\}$, $x_1^*, x_2^* \in \{1, 2 \ldots n\}$ and $\eta \in [0, 1]$ such that

$$c_j(x_j + 1)x_j^* = \hat{\lambda} \cdot c_j(x_j^*)x^* + \hat{\mu} \cdot c_j(x_j)x \tag{3.50}$$

for $j = 1, 2$ and

$$\eta \cdot c_1(x_1 + 1)x_1^* + (1 - \eta)c_2(x_2 + 1)x_2^* = \eta c_1(x_1)x_1 + (1 - \eta)c_2(x_2)x_2 \tag{3.51}$$

Proof. Let

$$\mathcal{H}_{c,x,x^*} = \{(\lambda, \mu) : c(x+1)x^* \leq \lambda \cdot c(x^*)x^* + \mu \cdot c(x)x\} \tag{3.52}$$

This set is the feasible region for specific configurations. The union of all these sets needs not be $\mathcal{A}(\mathcal{C})$. We use these sets to allow greater freedom when analyzing sections of the feasible region. Write $\partial\mathcal{H}_{c,x,x^*}$ for it's boundary which is comprised by the (λ, μ) which that satisfy equality. Define

$$\beta_{c,x,x^*} = \frac{c(x)x}{c(x+1)x^*} \tag{3.53}$$

which is well defined since $x^* > 0$ and c strictly positive. If $x \geq 1$ then we can uniquely express λ in terms of μ on the curve $\theta\mathcal{H}_{c,x,x^*}$. Solving for $\frac{\lambda}{1-\mu}$ in the equality from the definition of $\mathcal{H}_{c,x,x^*}$ we get

$$\frac{\lambda}{1-\mu} = \frac{c(x+1)}{c(x^*)} \frac{1 - \beta_{c,x,x^*}\mu}{1-\mu} \tag{3.54}$$

If $x = 0$, the only interesting case is

$$\mathcal{H}_{c,0,1} = \{(\lambda, \mu) : \lambda \geq 1\}$$

because if $x^* > 1$ the inside inequality becomes $c(1) \leq \lambda c(x^*)$ which is superseded by the previous case. In that case we have $\beta_{c,0,1} = 0$ and $\frac{\lambda}{1-\mu} = \frac{1}{1-\mu}$ for points in $\partial\mathcal{H}_{c,x,x^*}$.

In any case, what we need to remember is that λ and μ are defined by each other on the boundary, μ is a decreasing function of λ and as λ increases along the line $\mathcal{H}_{c,x,x^*}$ the value $\frac{\lambda}{1-\mu}$ is:

- strictly increasing if $\beta_{c,x,x^*} > 0$
- constant if $\beta_{c,x,x^*} = 0$
- strictly decreasing if $\beta_{c,x,x^*} < 0$

We use this knowledge to pinpoint $(\hat{\lambda}, \hat{\mu}) \in \mathcal{A}(\mathcal{C}, n)$. Since $\frac{\lambda}{1-\mu}$ is strictly increasing in both λ and μ $(\hat{\lambda}, \hat{\mu})$ must be on the boundary of $\mathcal{A}(\mathcal{C})$. Because neither λ or μ can move freely, this point must be in $\mathcal{A}(\mathcal{C}) \cap \mathcal{H}_(c, x, x^*)$ for some c, x, x^* that satisfy (3.50).

We define two cases, according to β_{c,x,x^*}. If $\beta_{c,x,x^*} = 1$ we have that $c(x+1)x^* = c(x)x$ from the definition of β (3.53). If we take $x_1 = x_1 = x$, $x_1^* = x_2^* = x^*$, $c_1 = c_2 = c$ and an arbitrary value of η we can plug everything in (4.47) and we are done.

If $\beta_{c,x,x^*} > 1$, without loss of generality, then $\hat{\lambda}, \hat{\mu}$ is an endpoint of the line segment $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c,x,x^*}$. Also, it is the endpoint of another line segment $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c',y,y^*}$ which has $\beta_{c,x,x^*} < 1$. The reason for both line segments is that $(\hat{\lambda}, \hat{\mu})$ is a boundary point and we cannot decrease both λ and μ.

Putting everything together, both pairs c, x, x^* and c', y, y^* satisfy equality (3.50) for $\hat{\lambda}, \hat{\mu}$. Moreover, by definition (3.53) we have that:

$$\beta_{c,x,x^*} > 1 \Rightarrow c(x+1)x^* > c(x)x$$

$$\beta_{c,y,y^*} > 1 \Rightarrow c'(y+1)y^* < c'(y)y$$

Relabeling $c_1 = c$,$c_2 = c'$ and $x_1, x_2 \to x, y$, plugging into (4.47) and considering that for $\eta = 0$ the left hand side is larger but for $\eta = 1$ the right hand we conclude that there exists $\eta \in [0, 1]$ which gives us the desired equality. □

For this lemma we assumed that the infimum in (3.48) is attained for some point in $\mathcal{A}(\mathcal{C},n)$. The next lemma treats the remaining case.

Lemma 3.4.4. *Let $\mathcal{C}$ be a finite set of strictly positive cost functions and n a positive integer. Suppose no point $(\lambda, \mu) \in \mathcal{A}(\mathcal{C},n)$ satisfies $\frac{\lambda}{1-\mu} = \gamma(\mathcal{C}, n)$. Then, there exists c such that*

$$\gamma(\mathcal{C}, n) = \frac{c(n)n}{c(1)} \tag{3.55}$$

and

$$c(n+1) \geq c(n)n \tag{3.56}$$

Proof. The idea behind the proof is that the infimum of (3.48) is not attained is if the set $\mathcal{A}(\mathcal{C}, n)$ has an unbounded boundary face $\mathcal{A}(\mathcal{C}, n) \cap \mathcal{H}_{c,x,x^*}$ with $\beta_{c,x,x^*} < 1$.

Since the proof is mathematically demanding but does not offer any new insights, the curious reader is pointed towards [Rou09]. □

3.4.1.3 Lower Bound Construction: The Finite Case

We will now present a lower bound construction. As before, we are still working under the assumption that $\mathcal{C}$ is finite, contains strictly positive, nondecreasing cost functions and there is an upper bound on the load of each resource.

We need to construct a game for which the inequalities (3.44)-(3.47) from **Lemma 3.4.2** are replaced by equalities for some outcomes. The construction follows a theme that appears often in congestion games. We devise a setting when each player has only two strategies: one which he uses a small number of resources and one which uses many. Those strategies are carefully crafted so that we only have 2 pure Nash equilibria. In each of them, all players all use they same strategy, either few or many resources. We try to make the outcome were everyone is using a few strategies the optimal one and then measure the other PNE's inefficiency.

More concretely, image a game with k players and resources. The cost function of each resource is linear, unless all players use a resource, when it becomes extremely large. Player i has 2 strategies: use resource i or use all the others. If every player uses his own resource, we have a PNE, because deviating and using all other resources has greater cost. However, every player using all resources but his own is also a PNE. This time, deviating to using only one resource will cause that resource to incur the large cost.

Doing something as simple for every $\mathcal{C}$ is not that easy, but it's possible using the properties of $(\hat{\lambda}, \hat{\mu})$ that we discovered in the previous section.

Theorem 8 (Main Construction). *Let $\mathcal{C}$ be a non-empty finite set of strictly positive cost functions and n a positive integer. There exist congestion games with cost functions in $\mathcal{C}$ and (pure) POA arbitrarily close to $\gamma(\mathcal{C}, n)$*

Proof. As before, we deal with two cases separately. We begin by analyzing the case were the value $\gamma(\mathcal{C}, n)$ is not attained by any $(\lambda, \mu) \in \mathcal{A}(\mathcal{C}, n)$. Let $c \in \mathcal{C}$ be the cost function satisfying the properties of **Lemma 3.4.4**. We define a congestion game similar to the above example. Let $E = \{e_1, e_2 \dots e_n\}$ and $n+1$ players, each of them having strategies e_i and $E \setminus \{e_i\}$ where i is the player. If all players choose their first strategy, the total cost of the outcome is $(n+1)c(1)$. If they choose their alternative strategy, the cost is $(n+1)c(n)n$.

Every player using his many-resource strategy is a PNE, by (3.56), because his cost before deviating is $c(n)n$. Dividing the two costs, we get a POA of $\frac{c(n)n}{c(1)}$ as (3.55) proving our first case.

The construction when $\gamma(\mathcal{C}, n)$ is attained within the region $\mathcal{A}(\mathcal{C}, n)$ is more complicated. Let $\frac{\hat{\lambda}}{1-\hat{\mu}} = \gamma(\mathcal{C},n)$. By **Lemma 3.4.3** we have $c_1, c_2, x_1, x_1^*, x_2, x_2^*$ with properties (3.50) and (4.47). We define a congestion game with $k = max x_1 + x_1^*, x_2 + x_2^*$ players and strategies $E_1 \cup E_2$, where E_1, E_2 are disjoint. Each set contains k resources, labelled 1 to k which are arranged in a cycle. The cost of resources is $\eta \cdot c_1(x)$ from E_1 and $(1-\eta) \cdot c_2(x)$ from E_2 .

Each of the players has two strategies. Player i's first strategy is P_i uses x_j resources from E_j starting from the i-th resource of each cycle and wrapping around if necessary, for $j = 1, 2$. His second strategy, Q_i is similar. In Q_i he uses x_j^* resources from E_j ending in $i-1$, wrapping around again if needed. We have choses k so it is large enough that the cycle parts P_i and Q_i do not overlap.

Let $\mathbf{y}$ and $\mathbf{y}^*$ denote the two outcomes where each players chooses strategy P_i and Q_i respectively. Every players chooses the same number of resources from each set E_i, just rotated by 1. Taking into account that the number of resources of E_i is equal to the number of players, we get that if players choose P_i then $y_e = x_1$ for $e \in E_1$ and $y_e = x_2$ for $e \in E_2$. If they choose Q_i, we get about the same: $y_e * = x_1^*$ for $e \in E_1$ and $y_e^* = x_2^*$ for $e \in E_2$. Thus $x_i(^*)$ show both the number of resources used and their load and our setting is completely symmetric.

We need to show that $\mathbf{y}$ is a PNE. For player i we have

$$
\begin{aligned}
C_i(\mathbf{y}) &= \sum_{e \in P_i \cap E_1} \eta \cdot c_1(y_e) + \sum_{e \in P_i \cap E_2} (1-\eta) c_2(y_e) && (3.57)\\
&= \eta c_1(x_1+1)x_1^* + (1-\eta)c_2(x_2+1)x_2^* && (3.58)\\
&= \sum_{e \in Q_i \cap E_1} \eta \cdot c_1(y_e+1) + \sum_{e \in Q_i \cap E_2} (1-\eta) c_2(y_e+1) && (3.59)\\
&= C_i(y_i^*, \mathbf{y}_{-i}) && (3.60)
\end{aligned}
$$

Where (3.58) from **Lemma 3.4.3** and (3.59) because P_i and Q_i are disjoint, meaning deviating to Q_i from P_i the player i only uses resources which are already used by x_1 or x_2 players. We have showed that deviation does not lead to any decrease in cost, thus $\mathbf{y}$ is a PNE.

We now need to compare the costs of the outcomes. In similar fashion to (3.59)

$$C(\mathbf{y}) = \sum_{i=1}^{k} C_i(\mathbf{y}) \tag{3.61}$$
$$= k \cdot [\eta \cdot c_1(x_1+1)x_1^* + (1-\eta) \cdot c_2(x_2+1)x_2^*] \tag{3.62}$$
$$= k\eta[\hat{\lambda}c_1(x_1^*)x_1^* + \hat{\mu}c_2(x_1)x_1] + k(1-\eta)[\hat{\lambda}c_1(x_2^*)x_2^* + \hat{\mu}c_2(x_2)x_2] \tag{3.63}$$
$$= \hat{\lambda} \cdot C(\mathbf{y}) + \hat{\mu} \cdot C(\mathbf{y}^*) \tag{3.64}$$

Rearranging the terms we get the desired lower bound. □

The attentive reader may have spotted that we used functions $\eta \cdot c_1(x)$ and $(1-\eta) \cdot c_2(x)$ which may not be in $\mathcal{C}$. However, we using standard techniques such as converting to rationals and scaling we can get arbitrarily close to this bound. In broad strokes, we can convert η to rationals which will not impact the PNE and only change the POA by an arbitrarily small amount. Then we scale these functions up so they are integer multiples. If we needed to scale up by m, instead of actually scaling the costs we just multiply the number of resources my m. Thus we are still using original cost functions, but the compound cost for each player is the same as the scaled up version. PNE's and POA's are not changed. Similar techniques are explained in more detail in [Rou03].

3.4.1.4 Lower Bound Construction: The General Case

The only thing left is to consider how the previous proof extends to the general case, when we $\mathcal{C}$ is infinite and contains just nondecreasing nonnegative functions.

Theorem 9. *For every nonempty set $\mathcal{C}$ of cost functions, the set of congestion games with cost functions in $\mathcal{C}$ is tight.*

Proof. As usual, we first treat the case where $\mathcal{C}$ contains functions that are not strictly positive. We will show that in this case, there exists a congestion game with infinite POA. Suppose $c \in \mathcal{C}$ satisfies $c(z) = 0$ and $c(z+1) > 0$ for some $z \geq 1$. Perform the main construction of **Theorem 8** using $c_1 = c_2 = c, \eta = \frac{1}{2}$, $x_1 = x_1^* = x_2^* = z$ and $x_2 = z+1$. The outcome $\mathbf{y}$ has cost $kc(z)z = 0$. The other outcome has positive cost. Precisely, it's cost is $C(\mathbf{y}) = \frac{k}{2}c(z)z + \frac{k}{2}c(z+1)(z+1) > 0$. As before, we need to prove that the first outcome is a PNE

$$\begin{aligned} C_i(y_i^*, \mathbf{y}_{-i}) &= \frac{1}{2}(zc(z+1) + zc(z+2)) \\ &\geq \frac{1}{2}(z+1)c(z+1) \\ &= C_i(\mathbf{y}) \end{aligned}$$

Deviating increases the cost, thus outcome $\mathbf{y}$ is a PNE and has infinite POA.

We can now assume that $\mathcal{C}$ only contains strictly positive cost functions. We will also assume that $\mathcal{C}$ is countable, which means that we can order the cost functions and use $\mathcal{C}_n$ to denote first n of them. We will use **Theorem 8** which applies to every $\mathcal{C}_n$.

Broadly, our argument starts by an assumption and shows that all paths we can go from there lead to the construction of game with the desired POA. We start assuming that there are infinitely many games which are not tight, because we have proven that finite games are tight for finite $\mathcal{C}$. Then we do case analysis on all the possible values of (λ, μ), since even thought we assumed that there infinite subsets of $\mathcal{C}$ which induce are games with loose POA bounds, any finite subset $\mathcal{C}_n$ is tight. Every item of the list will represent one analyzed case.

- Assume that for infinitely many n, $\gamma(\mathcal{C}_n)$ is not attained by any $(\lambda, \mu) \in \mathcal{A}(\mathcal{C}_n, n)$. By **Lemma 3.4.4** there are congestion games in $\mathcal{C}$ with arbitrarily high POA. Thus by the previous construction we have $(\lambda_n, \mu_n) \in \mathcal{A}(\mathcal{C}_n, n)$ with $\frac{\lambda_n}{1-\mu_n} = \gamma(\mathcal{C}_n, n)$ for all sufficiently large n.

- By using the definition of (λ, μ), rearranging constraint (3.41) and choosing arbitrary $c \in \mathcal{C}$,$x = n$ and $x^* = 1$ we get that

$$\lambda \geq \frac{c(n+1) - \mu c(n)n}{c(1)} \geq 1 - \mu n \tag{3.65}$$

  For all $(\lambda, \mu) \in \mathcal{A}(\mathcal{C}_n, n)$. If $\mu_n < -1$, then dividing the previous inequality by $1 - \mu_n$ we get

$$\frac{\lambda_n}{1 - \mu_n} > \frac{1 - \mu_n n}{1 - \mu} > \frac{n}{2}$$

  By **Theorem 8** we can construct games with arbitrarily high POA.

- We now assume that $\mu_n \geq -1$ for all sufficiently large n. Let's now assume that λ_n grows unbounded as n increases. Again, by **Theorem 8** we are done.

- Assume that $\lambda_n < M$ for some M for all sufficiently large n. Knowing that $\lambda \geq 1$ (easily proven by plugging $x = 0$,$x^* = 1$ in (3.41)) we have that $(\lambda_n, \mu_n) \in [1, M] \times [-1, 1]$ which is closed. By the Bolzanno-Weierstrass theorem, there exists a subsequence of (λ_n, μ_n) which converges to (λ^*, μ^*) since the domain is contains all it's limit points. If $\mu^* = 1$ then there are infinite pairs (λ_n, μ_n) with $\mu \to 1$ and μ increasing in n. Again, by **Theorem 8** we can construct games with arbitrarily high POA by picking a correct pair (λ_n, μ_n) from the subsequence.

- This time assume the subsequence converges on (λ^*, μ^*) with $\mu^* < 1$. Because the function $\frac{\lambda}{1-\mu}$ is continuous the subsequence $\frac{\lambda_n}{1-\mu_n}$ also converges on $\frac{\lambda^*}{1-\mu^*}$, so by the known theorem we can construct games with POA arbitrarily close to $\frac{\lambda^*}{1-\mu^*}$ We are not done yet thought, we need to show that $(\lambda^*, \mu^*) \in \mathcal{A}(\mathcal{C}_n, n)$.

- We claim that $(\lambda^*, \mu^*) \notin \mathcal{A}(\mathcal{C}_n, n)$. This means that there exist $x \geq 0, x^* > 0, c \in \mathcal{C}$ for which

$$c(x+1)x^* > \lambda^* c(x^*)x^* + \mu^* c(x)x$$

  . However, moving everything to the left hand side yields a function which is continuous in (λ, μ). Since for all (λ_n, μ_n) we have that

$$c(x+1)x^* - \lambda_n c(x^*)x^* - \mu_n c(x)x \leq 0$$

it cannot converge to a negative value. Thus $(\lambda^*, \mu^*) \notin \mathcal{A}(\mathcal{C}_n, n)$ and by definition $\gamma(\mathcal{C}_n, n) \leq \frac{\lambda^*}{1-\mu^*}$.

After all this, we have proven that under any circumstance we can either make a game with arbitrarily high POA or find a tight bound by smoothness argument.

The proof is not over yet. It remains to show that the same results hold for uncountable $\mathcal{C}$. This is done by a standard technique known as a density argument, which proves that we can approximate any result in the real number domain by a sufficiently close rational analogue (since the rationals are 'dense' within the reals). A rough description of the rest of the proof would be that using the countability of rational numbers, we construct rational approximations of cost functions and show that the equilibria and POA change by an arbitrarily small amount. □

Chapter 4

Composable and Efficient Mechanisms

4.1 Introduction

In the previous chapter we discussed a method of proving POA bounds which has great theoretical applications for studying games, is often natural to use and extends to various more permissive equilibrium concepts among other properties it also has. In this chapter we will follow a different approach along the same path. We will build upon our definition of smoothness and refine it, yielding the notion of *smooth mechanisms*.

Our goal this time is not to study the properties of games, but to design them. Smooth Mechanisms, introduced and proven to be useful by Vasilis Syrgkanis and Eva Tardos in [ST13] are a class of mechanisms with similar properties as smooth games but which are also quantifiably efficient when run in parallel or in sequence. This is of great importance to the effective implementation of mechanism in society, since the same mechanisms are used by players at the same time in a variety of settings. Imagine having many different buyers and sellers on eBay competing for the best prices or having multiple Search Engines wanting to post the most profitable adwords.

While each of these settings when studied in isolation has many well studied mechanisms with reasonable performance guarantees, the efficiency of the overall market is not entirely clear. One solution, which has its merits from a theoretical standpoint, would be to implement a centralized mechanism that coordinates users. In practice however, this is impossible and we are better off designing simple auctions users will be able to undestand.

To state the goal of this chapter very succinctly, we will try to define local properties of mechanisms which guarantee efficiency in a market setting where the same mechanism is used by the same agents for a variety of different purposes.

4.2 Smooth Mechanism Design Setting

4.2.1 Mechanism Design Setting

Before we begin, it will be useful to formalize the notion of mechanism for our setting. We will try to adapt traditional mechanism design notation conventions to facilitate writing expressions which include multiple mechanisms with the same participants. Moreover, we will assume that our mechanisms are quasilinear in money and players are risk neutral.

We begin by restating the *mechanism design setting*. A mechanism design setting consists of a set of players n, a set of outcomes $\mathcal{X} \subseteq \times \mathcal{X}_i$ where $\mathcal{X}_i$ is the set of allocations for player i. To complete the picture, each player also has a *valuation* $v_i : \mathcal{X}_i \to \mathbb{R}_+$. We use $\mathcal{V}_i$ to denote the set of possible valuations for player i. Since we will be discussing auctions, we will consider users with quasilinear preferences with payments. For player i, given an allocation $x_i \in \mathcal{X}_i$ and a payment p_i we have utility:

$$u_i^{v_i}(x_i, p_i) = v_i(x_i) - p_i \tag{4.1}$$

You may have noticed that allocation space can be any subset of the product space of individual allocations. This may seem counterintuitive at first, but is actually quite general, in a similar way that the correlated equilibrium is a subset of the product space but superset of the mixed Nash equilibrium. As such, this formulation can handle a wide variety of auction settings which may contain externalities or cooperation between players. Just to give a few examples, this framework can easily handle the *combinatorial auction* where $\mathcal{X}_i$ is the power set of items sold and $\mathcal{X}$ does not contain allocations were an item is sold more than once. *Combinatorial Public Projects* were $\mathcal{X}_\rangle$ is the power set of projects and $\mathcal{X}$ is the subset of the product space where each coordinate is the same. *Position Auctions* where $\mathcal{X}_i$ is the set of positions each player can attain and $\mathcal{X}$ is a subset where no two players occupy the same position. Using this subset of product structure for our allocation space also gives us an easy way to argue about the induced mechanism and valuation for player i, when many mechanisms run in parallel. As the proofs in the following sections will show, this is the reason we can exploit locally good properties of mechanisms to infer good global performance.

There is no need to use the same construction of the valuation space, which is $\mathcal{V} = \times \mathcal{V}_i$. If we want to show externalities or shared outcomes between players, we can choose specific v_i's and exploit the structure of the allocation space. Having both is just redundant.

Given an allocation $\mathcal{X}$ and a valuation space $\mathcal{V} = \times \mathcal{V}_i$ we have mechanism $\mathcal{M} = (\mathcal{A}, X, P)$. Here $A = \times A_i$ is the set of actions player i has (actions are performed independently and simultaneously), $X : A \to \mathcal{X}$ is the allocation function and $P : \mathcal{X} \to \mathbb{R}^n$ is the payment function. Each player gains some items according to his action and the allocation rule and pays a certain price.

Mechanisms we will study will also contain an extra (mostly technical) caveat. We will give players the opportunity to opt out of the mechanisms and gain zero utility. This will only be useful when for proving certain extension theorems were we will need to guarantee that no player will get negative utility because of incomplete information or randomness.

4.2.2 The Composition Framework

We are finally getting into the more interesting part of this chapter. As we have already discussed, we need a framework which allows us to study mechanisms that run simultaneously or sequentially as is more common in practice, instead of in isolation.

To be consistent we will assume n players and m mechanisms and we will use superscripts to show the mechanisms and subscripts to show the player. We will also favor using j when indexing mechanisms and i for players. As before, each mechanism is $M_j = (\mathcal{A}^j, X^j, P^j)$ where $\mathcal{A}^j = \times\mathcal{A}_i^j$, $X^j : \mathcal{A}^j \to \mathcal{X}^j$ and $P^j : \mathcal{A}^j \to \mathbb{R}_+^n$

In order to somehow link the mechanisms together, the valuation space $\mathcal{V}_i = \times_j \mathcal{V}_i^j$ of each player contains functions of the form $v_i : \times_j \mathcal{X}_i^j \to \mathbb{R}^+$. The individual allocation across all mechanisms is still denoted by $\mathcal{X}_i$. The players continue to have quasilinear utilities in this setting. The generalization is quite natural: using vectors to group allocations and payments from all mechanisms we have $x_i = (x_i^1, \dots x_i^m)$ and $p_i = (p_i^1, \dots, p_i^m)$:

$$u_i^{v_i}(x_i) = v_i(x_i) - \sum_{j=1}^{m} p_i i^j \tag{4.2}$$

We now need to show how the player's actions are composed. We will study two types of composition: simultaneous and sequential. In the case of simultaneous composition, each player i plays on all mechanisms at the same time, using action a_i^j on each mechanism j. In the case of sequential compositions we still use a_i^j for actions but if index j signifies the sequence of the mechanisms we also have that $a_i^j = a_i^j(h^j)$ where h^j is the history of observed actions (and possibly valuations) that have been revealed until mechanism j.

Putting everything together, simultaneous composition can be viewed as one global mechanism $\mathcal{M} = (\mathcal{A}, X, P)$ where $\mathcal{A}_i = \times_j \mathcal{A}) i^j$, $X(a) = (X^j(a^j))_j$ and $P(a) = \sum_{j=1}^m P^j(a^j)$. Sequential mechanisms are somewhat more complicated and can be represented as a game in extensive form or as a complicated mechanism were actions are functions of the previously observed history of play. We will not delve too deep into this definition, as it is only marginally used in one proof which is not too complicated anyway.

As always, our efficiency measure of an action profile a will be the social welfare

$$SW^v(a) = \sum_i v_i(X_i(a) \tag{4.3}$$

For every valuation $v \in \times\mathcal{V}_i$ there exists an optimal allocation $x^*(v)$ (but not necessarily an accompanying action profile) that maximizes the social welfare over all allocations $x \in \mathcal{X}$. For this allocation we have

$$OPT(v) = \sum_i v_i(x_i^*(v)) \tag{4.4}$$

As usual, we efficiency (or lack thereof) of a mechanism will be measured using the Price of Anarchy.

4.2.3 Valuations

One important but complicated and somewhat 'dry' topic concerning mechanisms is the user's valuations. As is common when studying any type of auction (the mechanisms we will focus on), we consider valuations that are complement free across mechanisms. The reason being that having complements amongst allocation or externalities between users makes inferring good global behavior from local properties impossible. Since most of our individual mechanisms will usually be single item auctions complement free can be easily captured by additive, subadditive or submodular valuations. In order to generalize our setting, we will extend these notion to classes of valuations where individual mechanisms have arbitrary allocation spaces and no assumptions will be made on the per mechanism valuations of each player.

Each player (the index i will be dropped because we will only argue about one arbitrary player) has a valuation of the form $v : \mathcal{X} \to \mathbb{R}_+$ where $\mathcal{X} = \times \mathcal{X}^j$. Notice this is not a subset of a product space, since mechanisms are completely independent and linked together only by the players' valuations.

Every valuations considered will be *monotone*. This means $v(S) \leq v(T)$ for any $S \subseteq T$, or in other words, the more players get the happier they will be. Going back to having only single item auctions as mechanisms, the holy grail of complement free valuations is the subadditive valuation:

$$v(S_1 \cup S_2) \leq v(S_1) + v(S_2) \tag{4.5}$$

For sets of items S_1 and S_2 acquired by the player across all mechanisms. Unfortunately, not all extendability results hold under this condition.

In order to get around this issue we will design some valuations which are general enough, are a subset of subadditive valuations across mechanisms and leave few assumptions about the individual per mechanism valuations. The first such valuation is the *fractionally subadditive*, which is a variation on the original fractionally subadditive valuation introduced by Feige in [Fei09]

Definition 13 (Fractionally Subadditive)**.** *A valuation is fractionally subadditive across mechanisms if*

$$v(x) \leq \sum_l a_l v(y^l) \tag{4.6}$$

whenever each coordinate x_j is covered in the set of solutions y^l, that is

$$\sum_{l : x_j = y^l_j} a_l \geq 1$$

At first this definition may seem rather arcane. In essence, the solutions y^l are the backbone of our valuation. Using those we try to make a fractional cover of any other outcome and then show that the valuation of that outcome is less than a weighted additive valuation of the covering solution. Again, considering only single item auctions, this definition is a subset of subadditive valuations.

A second important and much more useful set of valuation is the *XOS* valuation introduced in [LLN01] and proven to be equivalent to *fractionally subadditive* by

Feige in [Fei09]. Here we will use a modification of *XOS* because we have a vector of outcomes across mechanisms. The equivalence between the two valuations still holds however.

Definition 14 (XOS)**.** *A valuation is XOS if there exists a set* $\mathcal{L}$ *of additive valuations* $v_j^l(x_j)$ *such that*

$$v(x) = max_{l \in \mathcal{L}} \sum_{j=1}^{m} v_j^l(x_j) \tag{4.7}$$

This time instead of choosing from a fractional set cover, we are picking the additive valuation that maximizes the players gain for a given allocation. Again, *XOS* valuations are a subset of subadditive.

Theorem 10. *XOS* $\subset$ *Subadditive*

Proof. Let x be an allocation given to a player. Split this allocation into $x^1 + x^2 = x$ arbitrarily. From the definition of the *XOS* valuation, we have $l \in \mathcal{L}$ be the maximizing additive valuation. Let $l_1, l_2 \in \mathcal{L}$ be the maximizing valuations for x^1 and x^2 accordingly. Now we have:

$$v(x) = max_{l \in \mathcal{L}} \sum_{j=1}^{m} v_j^l(x_j) \tag{4.8}$$

$$\leq max_{l \in \mathcal{L}} \sum_{j=1}^{m} v_j^l(x_j^1) + max_{l \in \mathcal{L}} \sum_{j=1}^{m} v_j^l(x_j^2) \tag{4.9}$$

$$= v(x^1) + v(x^2) \tag{4.10}$$

Where the inequality holds by definition, because instead of choosing l_1 and l_2 we could have chosen l and gotten an equality. □

Again, *XOS* valuations are equivalent to *fractionally subadditive* but in most cases in this chapter they will be much easier to use.

Now that we have the necessary valuation classes to show how players may participate in multiple mechanisms, it's time to move on to the main definition of this chapter.

4.3 Smooth Mechanisms

We finally introduce the notion of a smooth mechanism, taking inspiration from the work done by Roughgarden in the previous chapter and extending it to better support players with quasilinear preferences especially in auction settings.

Definition 15 (Smooth Mechanism)**.** *A mechanism* $\mathcal{M}$ *is* (λ, μ)*-smooth if for any valuation profile* $v \in \times \mathcal{V}_i$ *and for any action profile* a *there exists a randomized action* $\mathbf{a}_i^*(v, a_i)$ *for each player such that:*

$$\sum_i u_i^{v_i}(\mathbf{a}_i^*, a_{-i}) \geq \lambda OPT(v) - \mu \sum_i P_i(a) \tag{4.11}$$

for some $\lambda, \mu \geq 0$. *We denote by* $u_i^{v_i}(\mathbf{a})$ *the expected utility of a player if* $\mathbf{a}$ *is a vector of randomized strategies. In general, boldface action will denote expectation.*

One natural way of interpreting the smooth mechanism condition is as guaranteeing an approximate analog of market clearing prices.

Although appearing similar, this notion of smoothness differs *significantly* from the Smooth Games studied in the previous chapter. There are two main differences. First of all, there is no direct comparison between outcomes, as with the previous chapter. Vaguely, $OPT(v)$ could play the role of the optimal outcome and $\sum_i P_i(a)$ could be interpreted as having similar value to the players gain for action profile a. To make this more concrete, we can transform our mechanism into a game by adding one more player with no actions and utility $\sum_i P_i(a)$. Now, an (λ, μ)-smooth mechanism could be seen as a $(\lambda, \mu - 1)$-smooth game where the term $-(\mu - 1)\sum_i u_i^{v_i}(a)$ is dropped in order to make the expression easier to use for quasilinear preferences.

The biggest difference however, is that the players know the others valuation and their own action when deviating. So, while Smooth Games required for us to connect either two arbitrary outcomes or one arbitrary with the optimal, here we only need to provide a randomized deviation that already knows the valuation and previous action. This makes smooth mechanism quite a lot easier to use. Since we are summing everything in the left hand side of the inequality (without affecting $OPT(v)$ or $P_i(a)$) we can focus on trying to maximize the gain of each player when he already knows the others' valuation. Then we can optimize in (λ, μ) to achieve the best POA bound.

However, relying on the knowledge of one's action for the randomized deviation also limits the extension of POA to correlated equilibria, instead of coarse correlated as normal smooth games. Intuitively, this is because we cannot use the smooth mechanism hypothesis in the coarse correlated equilibrium setting, since players are not signaled about their action. On the other hand, relying on your own previous actions allows bounds to extend to sequential composition of mechanisms.

4.3.1 Price of Anarchy and Extension Theorems

4.3.2 Extension to general equilibria

Before moving on to composition, we need to show that smooth mechanisms have low price of anarchy for a variety of equilibria when run in isolation. The proofs are fairly similar to the ones in the previous chapter.

Theorem 11. *If a mechanism is* (λ, μ)*-smooth and the players have the option to withdraw (gaining 0 utility) then the expected cost at any correlated equilibrium of the game is at least* $\frac{\lambda}{max\{1,\mu\}}$ *of the optimal social welfare*

Proof. Let $\mathbf{a}$ be a correlated equilibrium over action profiles $a \in \mathcal{A}$ such that for any player i and any strategy a_i

$$\mathbf{E}_{a_{-i}}[u_i^{v_i}(a_i, \mathbf{a}_{-i})] \geq \mathbb{E}_{a_{-i}}[u_i^{v_i}(a_i', \mathbf{a}_{-i})|a_i] \tag{4.12}$$

Plugging $\mathbf{a}^*(v, a_i)$ for a'_i we get:

$$\mathbf{E}_{a_{-i}}[u_i^{v_i}(a_i, \mathbf{a}_{-i})] \geq \mathbb{E}_{a_{-i}}[u_i^{v_i}(\mathbf{a}_i^*(v, a_i), \mathbf{a}_{-i})|a_i] \tag{4.13}$$

Adding for all players, taking advantage of the smooth mechanism property and taking expectations over $\mathbf{a}_i$ we get:

$$\mathbf{E}_{\mathbf{a}}[\sum_i u_i^{v_i}(\mathbf{a})] \geq \mathbb{E}_{\mathbf{a}}[\sum_i u_i^{v_i}(\mathbf{a}_i^*(v, a_i), \mathbf{a}_{-i})|a_i] \tag{4.14}$$

$$\geq \lambda OPT(v) - \mu \mathbf{E}_{\mathbf{a}} \sum_i P_i(\mathbf{a}) \tag{4.15}$$

Using $u_i^{v_i}(a) = v_i(X_i(a)) - P_i(a)$ we get:

$$\mathbf{E}_{\mathbf{a}}[\sum_i v_i(a)] \geq \lambda OPT(v) - (1 - \mu) \sum_i \mathbb{E}_{\mathbf{a}} P_i(\mathbf{a}) \tag{4.16}$$

If $\mu \leq 1$ then $\mathbf{E}_{\mathbf{a}}[\sum_i v_i(X_I(a))] \geq OPT(v)$ and we are done. If $\mu > 1$ and knowing that players can withdraw, thus always have non negative utility $v_i(X_i(a)) \geq P_i(a)$:

$$\mu \mathbf{E}_{\mathbf{a}}[\sum_i v_i(X_i(a))] \geq \mathbf{E}_{\mathbf{a}}[\sum_i v_i(a)] + (\mu - 1) \sum_i \mathbb{E}_{\mathbf{a}} P_i(\mathbf{a}) \tag{4.17}$$

$$\geq \lambda OPT(v) \tag{4.18}$$

Thus, we indeed have that the POA is at least $\frac{\lambda}{max\{1,\mu\}}$ □

It is important to take note that this proof would not have been enough for coarse correlated equilibria, however it obviously is enough for mixed Nash equilibria. In a coarse correlated equilibrium, player i is not signaled about his strategy so he is unable to deviate to $\mathbf{a}_i^*$. In an MNE he isn't signaled either though. However, this is not an issue because in an MNE the distribution of the other players' actions is independent of his own. Thus, our player can calculate every $\mathbf{a}_i^*(v, a_i)$ corresponding to any of his actions, and then sample these deviations.

4.3.3 Extension to Incomplete Information

We now turn to the incomplete information setting. As we have showed before, in the incomplete information setting instead of having the actual valuations be common knowledge we have distributions over valuations. Each player has a distribution $\mathcal{F}_i \sim \mathcal{V}_i$ over his valuations which is independent and known to the other players.

Just to recap, in the incomplete information setting a mechanism is still represented as $\mathcal{M} = (\mathcal{A}, X, P)$. The difference is the way actions are picked. Since valuations are not known a priori, every player i has a function $s_i : \mathcal{V}_i \to \mathcal{A}_i$. We will calculate the POA of smooth mechanisms around the Bayes-Nash equilibrium:

$$\forall a_i \in \mathcal{A}_i : \mathbb{E}_v[u_i^{v_i}(s(v))] \geq \mathbb{E}_v[u_i^{v_i}(a_i, s_{-i}(v)_{-i})] \tag{4.19}$$

Given a strategy $s : \times\mathcal{V}_i \to \times\mathcal{A}_i$, as before we will compare the social welfare

$$\mathbb{E}_v[SW^v(s(v))] \tag{4.20}$$

with the optimal welfare

$$\mathbb{E}_v[OPT(v)] \tag{4.21}$$

We continue to prove that even in the incomplete information setting, smooth mechanisms defined as in **Definition 15** achieve (in expectation) the same fraction of the optimal welfare as in the complete information setting, irrespective of the distribution of valuations. We shall prove this result for pure Bayes-Nash equilibria, but the proof can be generalized in a straightforward way up to mixed Bayes-Nash equilibria.

Before we begin proving, we need to address one issue about the deviating strategy of any player i: $\mathbf{a}_i^*(v, a_i)$. It depends on his action and the valuation of the other players, which is not public knowledge in the incomplete information setting. To surpass this difficulty, each player i will use random sampling to substitute knowing valuations v_{-i} and then come up with good deviations.

Theorem 12. *If a mechanism is (λ, μ)-smooth and players have the possibility to withdraw, then for any set of independent distribution F_i over valuations, every mixed Bayes-Nash equilibrium s of the game has expected social welfare at least $\frac{\lambda}{max\{1,\mu\}}$ of the expected optimal social welfare.*

Proof. We will prove this result around the pure Bayes-Nash equilibrium $s(v)$. In order to prove this, we need to discover a way for our players to use the randomized deviation given in the definition of smooth mechanisms, plug those deviations into the Bayes-Nash equilibrium hypothesis and try to rearrange terms to suit our purpose.

We will focus on the randomized deviation of arbitrary player i. When player i needs to submit his action, he only knows about his own valuation v_i, his own action $s_i(v_i)$, the distribution of the other players' valuations and the Bayes-Nash equilibrium $s(\cdot)$. To get around the lack of knowledge of the other players' valuations he performs random sampling $w \sim \times_i\mathcal{F}_i$ on all valuation profiles, including his 'own'. Then he devises a randomized action

$$\mathbf{a}_i^*((v_i, w_{-i}), s_i(w_i)) \tag{4.22}$$

Essentially the player uses his best guess of the other players valuation combined with his own actual valuation, but instead of using his actual action he deviates around his action in the sampled equilibrium $s(w)$. This method is reminiscent of a bluffing technique introduced in [Syr12] for sequential auctions. It is important to note that there is no actual bluffing going on here as players only submit one action. However, the randomized deviation does depend on w_i which is speculative information.

From the definition of the Bayes-Nash equilibrium deviating instead of playing $s_i(v_i)$ is not profitable:

$$\mathbb{E}_v[u_i^{v_i}(s(v))] \geq \mathbb{E}_{v,w}[u_i^{v_i}(\mathbf{a}_i^*((v_i, w_{-i}), s_i(w_i)), s_{-i}(v_{-i}))] \quad (4.23)$$
$$= \mathbb{E}_{v,w}[u_i^{w_i}(\mathbf{a}_i^*((w_i, w_{-i}), s_i(v_i)), s_{-i}(v_{-i}))] \quad (4.24)$$
$$= \mathbb{E}_{v,w}[u_i^{w_i}(\mathbf{a}_i^*(w, s_i(v_i)), s_{-i}(v_{-i}))] \quad (4.25)$$

The first inequality comes from the definition of the Bayes-Nash equilibrium. The second and third equalities hold because the distributions of v_i and w_i are the same, thus we can substitute them in this case.

Summing over all players:

$$\mathbb{E}_v[\sum_i u_i^{v_i}(s(v))] \geq \mathbb{E}_{v,w}[\sum_i u_i^{w_i}(\mathbf{a}_i^*(w, s_i(v_i)), s_{-i}(v_{-i}))] \quad (4.26)$$
$$\geq \mathbb{E}_{v,w}[\lambda OPT(w) - \mu \sum_i P_i(s(v))] \quad (4.27)$$
$$\geq \lambda \mathbb{E}_w[OPT(w)] - \mu \mathbb{E}_v[\sum_i P_i(s(v))] \quad (4.28)$$
$$\geq \lambda \mathbb{E}_v[OPT(v)] - \mu \mathbb{E}_v[\sum_i P_i(s(v))] \quad (4.29)$$

In (4.27) we have used the definition of the smooth mechanism. We get $OPT(w)$ instead of $OPT(v)$ since we are calculating the utilities around w. Finally, in (4.29) we use linearity of expectation to separate terms in a more useful form. The rest of the proof goes exactly as in the complete information case: if $\mu \leq 1$ we are finished, if $\mu > 1$ we use that players can withdraw and get the desired fraction of the expected optimal social welfare. □

Having established good POA bounds for both the complete and incomplete information setting, without assuming anything about our valuations (we *might* need a specific valuation to prove smoothness but these result do not impose new restrictions) we are ready to move on to quantifying the efficiency of such mechanisms when run in parallel or in sequence.

4.4 Compositionality of Smooth Mechanisms

4.4.1 Simultaneous Composition

As we briefly discussed during our revision of valuations, we do need specific valuation profiles across mechanism in order to guarantee good behaviour under composition. We will begin by proving POA bounds of simultaneous composition of mechanisms. The following proof also holds for *fractionally subadditive* valuations, but will be proven for XOS because it is more natural. It is important to note that the following proof is not by itself enough to prove POA bounds for general subadditive valuations. There may be some connection between simultaneous composition and subadditivity, but if there is it will have different bounds that those proven here.

Theorem 13. *Consider the simultaneous composition of m mechanisms. Suppose that each mechanism M^j is (λ, μ)-smooth when the mechanism restricted valuations come from a class $(\mathcal{V}_i^j)_{i\in[n]}$. If the valuation $v_i : \mathcal{X}_i \to \mathbb{R}^+$ of each player across mechanisms is fractionally subadditive, and can be expressed as an XOS valuation by component valuations $v_{ij}^k \in V_i^j$ then the global mechanism is also (λ, μ)-smooth.*

Proof. The main argument behind this proof is best described in the following steps:

- Consider the optimal outcome for any given XOS valuation
- Using the XOS valuation definition, obtain purely additive valuations that are easier to use. Have one such valuation for the optimal outcome and one for the deviation.
- Rewrite these 2 additive valuations as the induced per mechanism valuations, fixing one mechanism and considering all the others 'constant'.
- Use the smooth mechanism definition to connect per mechanism deviations to the optimal social welfare.
- Add over all players and all mechanism and perturb the sum to reach the desired inequality

More specifically, assume a valuation profile v and let $x^* \in \mathcal{X}$ be its accompanying allocation vector which optimizes social welfare. From the XOS definition, let v_{ij}^* be the additive valuation chosen for x^*, indexed by player and mechanism. By definition, we have that for any $x_i \in \mathcal{X}_i$ $v_i(x_i) \geq \sum_j v_{ij}^*(x_{ij}^*)$. Note we are summing over mechanisms. To complete our setting, assume a is the action profile submitted by the players.

Using this information, we need to show how players will devise a randomized deviation $\mathbf{a}_i(v, a_i)$ such that:

$$\sum_i u_i^{v_i}(\mathbf{a}_i^*, a_{-i}) \geq \lambda \sum_i v_i(x_i^*) - \mu \sum_i P_i(a) \tag{4.30}$$

To define such a deviation, we focus on arbitrary player i and mechanism M^j and devise the best randomized deviation locally using the smoothness property. For mechanism M^j each player has valuation v_{ij}^* on $\mathcal{X}_i^j$ and let v_j^* be the valuation profile for this mechanism. We have assumed that mechanism M^j is smooth for valuations coming from $\mathcal{V}_i^j$. This is not too strong of an assumption, since $\mathcal{V}_i^j$ only contains additive valuations. For most purposes (regarding auctions at least) mechanisms are usually smooth in this case. As a result of this, players can deviate to $\mathbf{a}_{ij}^* = \mathbf{a}_{ij}^*(v_j^*, a_i^j)$ used by the smoothness of M^j. By the smoothness property we have:

$$\sum_i u_{ij}^{v_{ij}^*} \geq \lambda OPT(v_j^*) - \mu \sum_i P_i(a_i^j) \tag{4.31}$$

$$= \lambda \sum_i (v_{ij}^*(x_{ij}^*)) - \mu \sum_i P_i(a_i^j) \tag{4.32}$$

Where the last equality comes from the definition of the XOS valuation. This is one instance where the proof about subadditive valuations would differ, since we cannot make this claim in the general case.

To produce the deviation necessary for the global mechanism, we bundle together all the individual, mechanism-wise deviations. Thus, for player i his global randomized deviation is $\mathbf{a}_i^* \mathbf{a}_i^*(v, a_i)$ where

$$(\mathbf{a}_i^*)_j = \mathbf{a}_{ij}^*(v_j^*, a_i^j) \tag{4.33}$$

consisting of the independent randomized deviation around the induced 'optimal' XOS valuation from each mechanism j as previously described. For each action a_i^* in the support of $\mathbf{a}_i^*$ we denote $X_i(a_i^*, a_{-i})$ the allocation vector of this particular deviation and action profile. By the XOS property and the induced optimal additive valuation v_i^* we have that:

$$v_i(X_i(a_i^*, a_{-i})) \geq \sum_j v_{ij}^*(X_i^j(a_i^*, a_{-i})) \tag{4.34}$$

Summing over mechanisms, the expected utility of player i performing this randomized deviation is:

$$u_i^{v_i}(\mathbf{a}_i^*, a_{-i}) \geq \mathbb{E}_{\mathbf{a}_i^*} \sum_j [v_{ij}^*(X_i^j(\mathbf{a}_{ij}^*, a_{-i}^j) - P_i^j(\mathbf{a}_{ij}^*, a_{-i}^j)] \tag{4.35}$$

Adding over all players, we have:

$$\sum_i u_i^{v_i}(\mathbf{a}_i^*, a_{-i}) \geq \sum_{i,j} \mathbb{E}_{\mathbf{a}_i^*}[v_{ij}^*(X_i^j(\mathbf{a}_{ij}^*, a_{-i}^j) - P_i^j(\mathbf{a}_{ij}^*, a_{-i}^j)] \tag{4.36}$$

Here we notice that the right hand side sum (over players, fixing one arbitrary mechanism) can be rewritten as:

$$\sum_i \mathbb{E}_{\mathbf{a}_i^*}[v_{ij}^*(X_i^j(\mathbf{a}_{ij}^*, a_{-i}^j) - P_i^j(\mathbf{a}_{ij}^*, a_{-i}^j)] = \sum_i \mathbb{E}_{\mathbf{a}_i^*}[u_{ij}^{v_j}(\mathbf{a}_{ij}^*, a_{-i}^j)] \tag{4.37}$$

We have used quasi-linearity and grouped terms to get the utility of all players participating in mechanism j and deviating from action profile a^j according to valuation v_{ij}^*. Combining (4.36) and (4.37) and using the smoothness of each mechanism $\mathcal{M}_j$:

$$\sum_i u_i^{v_i}(\mathbf{a}_i^*, a_{-i}) \geq \sum_{i,j} \mathbb{E}_{\mathbf{a}_i^*}[u_{ij}^{v_j}(\mathbf{a}_{ij}^*, a_{-i}^j)] \tag{4.38}$$

$$\geq \sum_j (\lambda OPT(v^j) - \mu \sum_i P_i^j(a^j)) \tag{4.39}$$

$$= \lambda \sum_{ij} v_{ij}^*(x_{ij}^*) - \mu \sum_{ij} P_i^j(a^j) \tag{4.40}$$

$$= \lambda \sum_i v_i^*(x_i^*) - \mu \sum_i P_i(a) \tag{4.41}$$

Where we used the definition of the XOS valuation twice. Once in the first inequality using that the starred versions of allocations and component valuations form the optimal social welfare and once in the last equality: $v_i^*(x_i^*) = \sum_j v_{ij}^*(x_{ij}^*)$ by definition of component valuations.

We have finally produced a randomized deviation and proven that the overall mechanism is also (λ, μ)-smooth □

This is a significant result because traditionally, even for restricted valuations, producing good mechanisms when multiple goods where involved is exceedingly tricky. As a side note, truthfulness is not preserved under simultaneous composition (we will present one example shortly). However, having such an elegant way of extending POA bounds, especially when truthful mechanisms perform poorly at the same tasks, surpasses this limitation.

4.4.2 Sequential Composition

There are many scenarios when mechanisms may not be run in parallel but in sequence. Actually, most of them are run this way. Imagine a network trying to allocate resources on the fly or a company trying to bid on public projects while securing a loan. Also when participating in any high risk situation (buying and selling stocks for example) players do may not want to simultaneously wager on outcomes. They might prefer to play slower and have more control over their actions.

Unfortunately, since we required to define a specific class of complement free valuation in order to guarantee decent behavior in simultaneous composition, one can image how difficult it is to do the same for sequential. In the sequential case we need to account for signaling and leaking of information as time goes on and more mechanisms are played. In this regard, going from local smoothness to smoothness in parallel does not seem that complicated, since players actually want to do as good as they can in each mechanism. Because there is no sharing of information, locally good behavior leads to global smoothness and i'ts a happy coincidence that the (λ, μ) remain the same.

However, as [LST12, Syr12] have shown, sequential composition leads to terrible POA bounds even for auctions that are tame and truthful in the isolated case. To counter this, we will be able to prove sequential composition only under an *extremely* limiting valuation set.

Definition 16 (Unit-demand)**.** *A valuation v_i is unit demand across m mechanisms if for $x_i \in \times_j \mathcal{X}_i^j$ we have:*

$$v_i(x_i) = \max_{j \in [m]} v_i^j(x_i^j) \tag{4.42}$$

Where v_i^j are per mechanism valuations.

Or in other words, the overall valuation of an outcome is equal to it's most valuable component, as dictated by mechanism-wise valuations. Notice that as in the case of simultaneous composition, we have not made any assumptions about the valuation of each mechanism. Moreover, unit-demand valuations are a subset of XOS. The proof is very simple, and relies on viewing unit-demand valuations as XOS where all elements except one of each additive component valuation is zero.

The reason why this valuation was chosen is because essentially players only need to perform well once, at their 'best' mechanism. This makes smoothness arguments where players only make one good deviation work.

Something to take note is that we in proving sequential composition we cannot take independent randomized deviations as before. As the mechanism unfolds, each new action and deviation will depend on the previously observer history (even though our valuation is restrictive). The statement of the theorem and it's proof rely to a certain extent on extensive form games, however our limited explanation provided in the introduction will suffice to graps the essence of the proof, which is the leaking of information and choosing the right point to deviate. As before, truthfulness is sacrificed.

Theorem 14. *Consider the sequential composition of* m*,* (λ, μ)*-smooth mechanisms defined on valuation spaces* $\mathcal{V}_i^j$*. If each valuation* $v_i : \mathcal{X}_i \to \mathbb{R}^+$ *is of the form* $v_i(x_i) = \max_{j \in [m]} v_i^j(x_i^j)$*, with* $v_i^j \in \mathcal{V}_i^j$ *then the global mechanism is also* $(\lambda, \mu + 1)$*-smooth, regardless of what information was released to the players during the sequential rounds*

Proof. Consider a valuation profile v and an action profile a of the (complete) sequential composition. Our goal is to design randomized deviations for each player to prove overall smoothness. Keep in mind that this time we do not have just $\mathbf{a}_{ij}^* = \mathbf{a}_{ij}^*(v, a_i^j)$, since this formulation misses the information released throughout the mechanisms. Instead, we have $\mathbf{a}_{ij}^* = \mathbf{a}_{ij}^*(h_i^j)$ where h_i^j encapsulates the history of play up to mechanism $\mathcal{M}_j$.

As we did with simultaneous composition, we start arguing about the optimal outcome. Let x^* be the optimal allocation for valuation profile v. Given unit demand valuations for each player i of the form:

$$v_i(x_i) = \max_{j \in [m]} v_i^j(x_i^j) \tag{4.43}$$

where $v_i^j \in \mathcal{V}_i^j$. To denote the actual maximizing valuation and mechanism pair we will use $j_i^* = \text{argmax}_{j \in [m]} v_i^j(x_{ij}^*)$, in other words $v_i(x_i) = v_i^{j_i^*}(x_{ij_i^*})$. So for each player, j_i^* represents his most valued individual outcome, disregarding payments.

To prove the theorem, we will construct a randomized deviation $\mathbf{a}_i^*(v, a_i)$ for each player such that:

$$\sum_i u_i^{v_i}(\mathbf{a}_i^*(v, a_i), a_{-i}) \geq \lambda \sum_i v_i(x_i^*) - (1 + \mu) \sum_i P_i(a) \tag{4.44}$$

We will focus our attention on player i. His randomized deviation will be $\mathbf{a}_i^* = \mathbf{a}_i^*(v, a_i) = (\mathbf{a}_i^*(h_i^j, v, a_i))$ where he will play exactly as a_i until mechanism $j = j_i^*$ where he will play a randomized strategy $\mathbf{a}_{ij}^*$ depending on the mechanism $\mathcal{M}_{j_i^*}$ and the observer actions. As we see, the deviations is a function of the observed history, which goes against the rules of smooth mechanisms where they only depend on v and a_i. Knowing that we can convert a sequential mechanism from an extensive from game to a normal form game with much more complicated actions, in order to keep the notation simple and expose the spirit of the proof we will slightly abuse our rules.

Continuing, the utility of player i is at least:

$$u_i^{v_i}(\mathbf{a}_i, a_{-i}) \geq \mathbb{E}_{a_{ij}^*}[v_i^j(X_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))) - P_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))] - P_i^{j-}(a) \tag{4.45}$$

$$\geq \mathbb{E}_{a_{ij}^*}[v_i^j(X_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))) - P_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))] - P_i(a) \tag{4.46}$$

where $P_i^{j-}(a)$ is the price paid by the player to mechanisms before $\mathcal{M}_j$ and $a_{-i}^j(h_{-i}^j)$ is the action profile submitted by the rest of the players at mechanism j when they played action profile a up to that mechanism, producing history h_i^j. Although the previous inequality seems confusing, it holds due to the following argument: the utility of player i deviating to $\mathbf{a}_{ij}^*$ is whatever he gains up to mechanism j plus whatever he gains after that. Since the second part can be made equal to 0, obviously his gain up to j is at least his expected gain at j minus the total payment.

Summing over all players and mechanisms, we get:

$$\sum_i u_i^{v_i}(\mathbf{a}_i^*, a_{-i}) \geq \sum_j \sum_{i:j_i^*=j} \mathbb{E}_{a_{ij}^*}[v_i^j(X_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))) - P_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))] - \sum_i P_i(a) \tag{4.47}$$

Note that:

$$\sum_{i:j_i^*=j} \mathbb{E}_{a_{ij}^*}[v_i^j(X_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))) - P_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))] \tag{4.48}$$

is exactly the utility gained by each player that deviated to $\mathbf{a}_{ij}^*$ at M_j, while the remaining players play actions $a_{-i}(h_{-i}^j)$ and all players with $j = j_i^*$ have valuations $v_i^j : \mathcal{X}_i^j \to \mathbb{R}^+$ and the rest have 0 valuation for any outcome. Also note that the history of play remains h_i^j, caused by the original action profile a.

By the smoothness of mechanism M_j, for the induced valuation profile (where only players with $j_i^* = j$ have nonzero value) there must exist a strategy $\mathbf{a}_{ij}^* = \mathbf{a}_{ij}^*(v, a_i^j(h_i^j))$ such that:

$$\sum_i u_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j)) = \sum_{i:j_i^*=j} \mathbb{E}_{a_{ij}^*}[v_i^j(X_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))) - P_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))] \tag{4.49}$$

$$\geq \lambda OPT(v^j) - \mu \sum_i P_i^j(a_{-i}^j(h_{-i}^j)) \tag{4.50}$$

$$\geq \lambda \sum_{i:j_i^*=j} v_i^j(x_{ij}^*) - \mu \sum_i P_i^j(a_{-i}^j(h_{-i}^j)) \tag{4.51}$$

since the value of the optimal outcome is at least the value of outcome x_j^*. Again, the payment is of the correct form, since it only depends on history h_i^j caused by strategies a, before the deviation. Plugging this result back to (4.47), we get:

$$\sum_i u_i^{v_i}(\mathbf{a}_i^*, a_{-i}) \geq \sum_j [\lambda \sum_{i:j_i^*=j} v_i^j(x_{ij}^*) - \mu \sum_i P_i^j(a_{-i}^j(h_{-i}^j))] - P_i(a) \tag{4.52}$$

$$\geq \lambda OPT(v) - (1+\mu) \sum_i P_i(a) \tag{4.53}$$

And we finally reach the desired result. □

The main point of this proof is how the specific qualities of the unit-demand valuation are put into play. Basically, the two most interesting techniques are only analyzing the utility of each player up until he deviates (without affecting the rest of the game) and still obtaining a significant portion of the optimal outcome, both of which would have been impossible with an XOS valuation.

4.5 Case Study: Auctions

4.5.1 First Price Auction

We will now attempt to use the tools devised in this chapter to study our first type of mechanism, the first price auction.

Consider a first price auction with n players with arbitrary valuation profiles $v \in \times\mathcal{V}_i$, submitting bids b_i. Of course, the optimal social welfare is $OPT(v) = \max_i v_i$ and total price paid by the players $\max_i b_i$. Given bids (actions) b, we would need to produce unilateral deviation $\mathbf{b}_i^* = \mathbf{b}_i^*(v, b_i)$ to optimize the following:

$$\sum_i u_i^{v_i}(\mathbf{b}_i^*, b_{-i}) = \lambda \max_i v_i - \mu \max_i b_i \tag{4.54}$$

Although the valuations are public information, the deviating players only know their own bid. As such, we cannot aggressively optimize because we need to take into account scenarios where players may bid above their valuation.

Having all players deviate less that their value guarantees the left hand side being non negative. This covers most paradoxical bids, but is not enough when the players actually bid intelligently i.e. $\max_i v_i \geq s \max_i b_i$. Moreover, we can craft bids where any player but $argmax_i v_i$ cannot have positive utility. This can be done by having the second highest player bid above his valuation and the top player bidding just over that. Thus we can only rely on the top valued player to maximize the sum of utilities in the worst case. So we can safely assume all other players deviate to 0.

Our first candidate deviation will be having the top player (let's call him h) bid half his valuation $\mathbf{b}_h = \frac{v_h}{2}$. Now, if $v_h \geq 2b_h$ he will get $u_h^{v_h} = \frac{v_h}{2}$. If not, he will get 0. By setting $\lambda = \frac{1}{2}$ and $\mu = 1$, if $v_h \geq 2b_h$ we have $\lambda v_h - \mu b_h \leq \frac{v_h}{2}$ else $\lambda v_h - \mu b_h \leq 0$. Thus the first price auction is $(\frac{1}{2}, 1)$-smooth, having a POA of 2.

However, we can do better. To maximize his utility, player h can submit a randomized bid, drawn from the distribution

$$f(x) = \frac{1}{v_h - x} \tag{4.55}$$

and support $[0, (1-\frac{1}{e})v_h]$. In this case, his utility is

$$u_i^{v_i}(\mathbf{b}_i, b_{-i}) \geq \int_{\max_{i\neq h} b_i}^{(1-\frac{1}{e})v_h} (v_h - x)\frac{1}{v_h - x}dx \tag{4.56}$$

$$\geq \int_{\max_{i\neq h} b_i}^{(1-\frac{1}{e})v_h} dx \tag{4.57}$$

$$\geq (1-\frac{1}{e})v_h - \max_{i\neq h} b_i \tag{4.58}$$

showing that the first price auction is indeed $(1-\frac{1}{e}, 1)$-smooth with a POA of $\frac{e}{e-1}$ which is a tighter upper bound than 2. This is not that impressive on its own, but it's extension to m simultaneous first price auctions with XOS valuations is. Although this bound had been already proven in [Syr12], this construction provides more direct insight into the role of the valuation and player strategies.

4.5.2 Second Price Auction

Despite having several very nice properties independently, like being truthful and optimal, the second price auction is not smooth. The main culprit for this is the loose connection between the payments and value of each outcome.

To show that a mechanism is not smooth, we need to find valuations $v \in \times\mathcal{V}_i$ and bid profile b such that the smoothness condition is not achieved for any randomized bid $\mathbf{b}_i^* = \mathbf{b}_i^*(v, b_i)$.

Assume a 2 player second price auction and take any valuation v. Without loss of generality, assume $v_1 > v_2$. As our bid profile, we take $b_1 = v_2$ and $b_2 = v_1$. The utilities of the players are $u_1 = 0$ and $u_2 = v_2 - v_2 = 0$. As a result, the right hand side of the smoothness condition is:

$$\lambda v_1 - \mu v_2 \tag{4.59}$$

Let's now try to maximize each players utility through deviation. Since $b_2 = v_1$, player 1 cannot achieve positive utility. Also, because $b_1 = v_2$, player 2 neither player 2 can. So we have:

$$0 \geq \lambda v_1 - \mu v_2 \Leftrightarrow \tag{4.60}$$

$$\frac{v_2}{v_1} \geq \frac{\lambda}{\mu} \tag{4.61}$$

This cannot hold for all v_1, v_2 and a specific pair of λ, μ. So, the second price auction is not a smooth game.

It is also somewhat interesting to see that the restriction of the randomized deviation not to depend on b was key in this proof, although it did not matter at all in proving the first price auction were smooth. It did however indirectly affect the smoothness constrains (λ, μ). However, having deviations independent of b is what makes these POA bounds extend, by forbiding from exploiting the full information PNE structure too much.

The importance of this restriction becomes readily apparent when considering sequential composition of second price auction with unit-demand valuations. Remember that sequential first price auctions with the same valuations are $(1-\frac{1}{\epsilon}, 2)$-smooth by **theorem 14**. The following construction is from [LST12].

4.5.2.1 Arbitrarily High POA in Sequential Second Price Auctions with Unit-Demand Valuations

Assume a sequential second price auction with unit-demand valuations of 4 items and 4 players. We will denote items by A_1, B_1, A_*, B_* auctioned in that order and players by valuations v_1, v_a, v_b, v_c. Their valuations are:

- $v_1(A_1) = 1 - \epsilon$ and $v_1(B_1) = \delta$
- $v_a(A_*) = 1$
- $v_b(A_*) = v_b(B_*) = 2$
- $v_c(A_*) = v_c(B_*) = 2$

Any value not accounted for is 0. We begin by focusing our attention to the last two items, A_* and B_*. Here we have 2 subgame perfect equilibria. In the first, denoted Spe_1, player b bids $1+\epsilon$ and 0 and c bids 0 and ϵ, gaining utilities 1 and 2 each. The second, Spe_2, is symmetrical by swapping b for c, gaining 2 and 1 this time. Remember that v_b prefers Spe_2 and v_c prefers Spe_1.

We will construct an Spe which exhibits high POA.

- If players b or c win auction A_1 at 0 price, then Spe_2 will be used.
- If player 1 wins auction A_1 then Spe_1 will be implemented.
- If player 1 loses auction A_1 but sets a positive price then if either b or c win B_1 Spe_2 is implemented else Spe_1

To sum up, player b wants to win A_1 at 0 price, or let player c win the first auction and then win B_1 to secure outcome Spe_2. On the other hand, player c wants either player 1 to win A_1 or player 1 to win B_1. Thus, player c does not have any incentive to bid on either A_1 or B_1.

We are left with players 1 and b. If player b wins A_1, then at B_1 he has utility 2 if he wins and 1 if he loses. Thus, he bids 1 at B_1 and beats player 1 who bids at most δ. Knowing he has utility $1-\epsilon$ at B_1, player b can bid 1 at A_1. However, player 1 knows he can't win B_1, thus bids nothing on A_1.

So we are end up with the following equilibrium: only player b bids 1 on A_1, player 1 takes B_1 for free and then Spe_2 is played. This might be obtuse, but no player can win by deviating. Keep in mind the only payments happen at auctions A_* and B_*.

Now suppose that instead of having just A_1 and B_1 we have n players and corresponding A_i, B_i for each one, with the same valuations as player 1 had for A_1, B_1. For each of these pairs of auctions, we can keep the previous rules, meaning

Spe_2 will happen only if player b plays 'correctly' on n auctions, else Spe_1 will be implemented. Since there are no payments to accumulate and c does not have new incentive to disturb b's plan, since valuations are unit demand, we have achieved an equilibrium where n pairs of auctions go as described and Spe_2 concludes the mechanism.

The social welfare of our bad equilibrium is $n\delta + 4$, where as the value of the optimal outcome is $(1-\epsilon)n + 4$, leading to a POA of $\frac{(1-\epsilon)n+4}{n\delta+4}$ which can be made arbitrarily high.

4.6 Weakly Smooth Mechanisms

All the inefficiencies of the second price auction we studied were caused by players bidding, in one way or another, above their valuation. When proving why the second price auction was not a smooth mechanism, recall that one player made a blatantly high bid, which at best could get him 0 utility. In the case of sequential composition, overbidding was rampant but more subtle: player b signaled his preference, by submitting a high bid on an item he didn't specifically need.

Weakly smooth mechanisms attempt to address this issue, by disallowing overbidding, after having precisely defined it. In single price auctions overbidding is easily identified: players bid more than the items value. In more complicated settings where the connection between reported bids, valuations and payments is less direct overbidding needs to be defined. The definition used generalizes no-overbidding assumptions from [BR11, CKKK11, CKS08]. We start of by defining willingness to pay.

Definition 17 (Willingness-to-pay). *Given a mechanism $(\mathcal{A}, X, P)$ a players maximum willingness-to-pay for an allocation x_i is when using strategy a_i is defined as the maximum he could ever pay conditional on allocation x_i:*

$$B_i(a_i, x_i) = \max_{a_{-i}:X_i(a)=x_i} P_i(a) \tag{4.62}$$

This is the most a player could pay for a combination a_i, x_i (as if the other players conspired against him). Based on that, we can move on to no-overbidding.

Definition 18 (No-overbidding). *A randomized action profile* $\mathbf{a}$ *satisfies the no-overbidding assumption if:*

$$\mathbb{E}_{\mathbf{a}}\left[B_i(\mathbf{a}_i, X_i(\mathbf{a}))\right] \le \mathbb{E}_{\mathbf{a}}\left[v_i(X_i(\mathbf{a}_i))\right] \tag{4.63}$$

The meaning of this definition is obvious: no one can expect to pay more than what he expects to get.

Definition 19 (Weakly Smooth Mechanism). *A mechanism is weakly (λ, μ_1, μ_2)-smooth for $\lambda, \mu_1, \mu_2 \ge 0$ if for any type profile $v \in \times_i \mathcal{V}_i$ and for any action profile a there exists a randomized action $\mathbf{a}_i^*(v, a_i)$ for each player i, such that:*

$$\sum_i u_i^{v_i}(\mathbf{a}_i^*(v, a_i), a_{-i}) \ge \lambda OPT(v) - \mu_1 \sum_i P_i(a) - \mu_2 \sum_i B_i(a_i, X_i(a)) \tag{4.64}$$

On it's own this isn't too interesting. For now we relaxed the smoothness condition and made sure there is a more direct connection between risky play and value gained, avoiding pitfalls we encountered in second price auctions. Combined with the no-overbidding assumption we reach interesting theorems.

Theorem 15. *If a mechanism is weakly (λ, μ_1, μ_2)-smooth then any correlated equilibrium in the full information setting and any mixed Bayes-Nash equilibrium in the incomplete information setting that satisfies the no overbidding assumption achieves efficiency at least $\frac{\lambda}{\mu_2+\max\{1,\mu_1\}}$ of the expected optimal*

It's proof is fairly similar to the one we already did for the complete information setting. Moreover, sequential and simultaneous composition theorems also carry over.

Theorem 16. *Consider the mechanism defined by the sequential composition of m mechanisms. Suppose that each mechanism j is weakly (λ, μ_1, μ_2)-smooth when the mechanism restricted valuations of the players come from a class of valuations $(\mathcal{V}_i^j)_{i\in[N]}$.*

- *If the valuation $v_i : \mathcal{X}_i \to \mathbb{R}^+$ of each players across mechanisms is XOS with component valuations $v_{ij}^l \in \mathcal{V}_i^j$ then the global valuation is also (λ, μ_1, μ_2)-smooth*
- *If the valuation $v_i : \mathcal{X}_i \to \mathbb{R}^+$ of each players across mechanisms is unit-demand with valuations from $v_{ij} \in \mathcal{V}_i^j$ then the global valuation is also $(\lambda, \mu_1 + 1, \mu_2)$-smooth*

Proof. The proof is identical to the one for smooth mechanisms. In those proofs, our technique was to express randomized deviations as a function of induced valuations, add over mechanisms and players and try to salvage as much of the optimal as possible while grouping payments together, to reach the correct form of a smoothness condition.

For this proof we can do exactly the same, but we will also have to account for the $B_i(\cdot, \cdot)$ terms. Fortunately, due to the independence of action spaces and allocations across mechanisms, willingness-to-pay of a player is additive.

$$\sum_j B_i^j(a_i^j, x_i^j) = \sum_j \max_{a_{-i}^j : X_i^j(a_i^j) = x_i^j} P_i^j(a^j) \tag{4.65}$$

$$= \max_{a_{-i} : X_i(a_i) = x_i} \sum_j P_i^j(a^j) \tag{4.66}$$

$$= \max_{a_{-i} : X_i(a_i) = x_i} P_i(a) = B_i(a_i, x_i) \tag{4.67}$$

Using this the new term can be handled exactly like the payments. □

We will now show that the second price auction is $(1, 0, 1)$-smooth. Assume a second price auction with n players. The highest player valuation is v_h and the highest bid submitted is b. The randomized deviation will be player h bidding v_h

and everybody else bidding 0. Only player h can have positive utility. The sum of willingness-to-pay is b because players who did not submit the highest bid win nothing, while the highest bidder might pay up to b. Plugging into the smoothness condition we have:

$$u_h(v_h, a_{-h}) \geq v_h - 0 - b \tag{4.68}$$

If $v_h < b$ then player h wins nothing and the right hand side is negative. Else, player h has utility $v_h - b$.

Thus, the second price auction is a smooth game and under the no overbidding assumption has certain decent compositional properties. However, it's POA is $\frac{1}{2}$ which is worse than the first price auction's $\frac{e}{e-1}$

Chapter 5

Lower Bounds through Computational Complexity

5.1 Motivation

In this chapter we will try to derive lower bounds on the POA using techniques from computational complexity, instead of arguing directly about the players actions and strategic behavior. This foundations of this line of research have been laid by Tim Roughgarden in [Rou14] and the main proofs shown in this chapter appear in this paper.

Up to this point, all solution concepts to games were studied with no mention of the computational aspects of the underlying process. We did this, despite several important results of algorithmic game theory hinting that computing Nash equilibria might even be intractable [CDT09, DGP09, EY10, HM10]. This may imply that computing Nash equilibria can solve problems that have no efficient solution.

The POA connects the cost of the worst equilibrium of our solution concept with the best optimal solution. This means that in some problems, easy to compute, 'tractable' equilibria may be significantly suboptimal, especially if the underlying optimization problem is intractable. With any luck, this can lead to lower bounds on the POA, in a way similar to computing polynomial approximate solution to hard problems. We will examine this relationship in two cases, starting off with the relatively tame Congestion Game, following up on our analysis from previous chapters and then setting the basics of multi-party communication protocols to establish similar lower bound proofs for auctions.

5.2 Cost Minimization in Congestion Games

We start by a refresher on Congestion Games. We define a Congestion Game by:

- A ground set of resources E
- A set of n players, each of them with action sets $A_1, \dots, A_n \in 2^E$

- A cost function $c_e : \mathbb{Z}^+ \to \mathbb{R}^+$ for each resource, which is a function of the number of players using it

For any given action profile $\mathbf{a} \in A_1 \times \ldots \times A_n$, we define the load x_e of resource $e \in E$ as the number of players using it. The cost to player i is defined as $C_i(\mathbf{a}) = \sum_{e \in a_i} c_e(x_e)$ and the total cost of an outcome:

$$C(\mathbf{a}) = \sum_{i=1}^{n} \sum_{e \in a_i} c_e(x_e) = \sum_{e \in E} c_e(x_e) x_e$$

As we have already mentioned, the intricacies of calculating the POA for congestion games lies in the chosen set of C of cost functions, which has already been thoroughly investigated by [ADG$^+$06, CK05, Rou09] to name a few.

To find a lower POA bound, we first need to convert this game theoretic problem in an optimization one. Let us restrict our analysis for polynomial cost functions with nonnegative coefficients of degree d contained in the set C_d. There exist tight lower bounds following an intricate construction from [ADG$^+$06]. The optimization problem derived will be called CostMinimization(d) or CM(d) for short: given a description of a congestion games with cost functions in C_d, compute the cost of the optimal outcome. Suppose the players strategies and the polynomial coefficients are part of the input. This problem is known to be NP-complete by reduction from 3-Partition, proven by [MS12]. We will use this to show that there are no 'easy' to compute Nash equilibria with good POA.

Theorem 17. *Fix any $d \geq 1$.*

1. *There is a polynomial time reduction R from an NP-complete problem Π that computes a parameter C^* and maps 'yes' and 'no' of Π to instances of CM(d) with cost at most C^* and at least ρC^* respectively.*

2. *NP $\neq$ CoNP*

Then the worst case pure POA in congestion games with cost function in C_d is at least ρ.

The proof of this theorem relies on some simple propositions. The first, proven by Rosenthal ([Ros73]) in 1973:

Proposition 17.1. *Every Congestion Game has at least one PNE.*

The second proposition is:

Proposition 17.2. *Deciding whether an action profile of a Congestion Game is a PNE can be solved in polynomial time.*

Proof. Just remember that the action available for each player are part of the input. Thus, given an action profile $\mathbf{a}$ we can go over all players one by one and try every different action, stopping if we find one that decreases this player's cost. □

The last proposition is the simplest:

Proposition 17.3. *The problem of computing the cost of an action profile can be solved in polynomial time.*

Proof. Fairly obvious. We can go through each players action and compute the loads for each resource in polynomial time. Then we need to evaluate the costs of each resource (also done in polynomial time) and add everything together. □

Armed with these propositions, we can go on and prove the theorem.

Proof of 17. Let α be the worst case POA of congestion games with functions in C_d and consider the following nondeterministic protocol $\mathcal{P}$ for instances of the CM(d) problem produced by the reduction R:

1. Given the Congestion Game produced by R, nondeterministically guess a PNE **a**
2. Deterministically verify that it is indeed a PNE. This step is necessary, because a nondeterministic Turing Machine tries to find the path that outputs 'yes', thus we need this filter
3. Compute the cost C of said PNE
4. Output 'yes' if and only if $C < \rho C^*$, where C^* is the parameter given by the reduction

By the definition of the POA the cost C is at most α the cost of the optimal outcome. Thus, if $\alpha < \rho$ then our procedure outputs 'yes' if there is an outcome with cost less than C^*, since $C \leq \alpha C^* \leq \rho C^*$. It outputs 'no' whenever every outcome has cost at least C^*. This reduces an problem in NP to finding equilibria for Congestion Games, which is a PLS complete problem. Thus, as proven in [JPY88], this would imply NP = coNP leading to a contradiction. So we have $\alpha \leq \rho$. □

This theorem reduced proving POA lower bounds for Congestion Games to proving hardness of approximation results for the optimization problems CM(d). This class of optimization problems (either for polynomials or for other cost functions) has not been well studied. However, we present this method as it is interesting on it's own and provides insights for the more complicated case with auctions.

5.3 Welfare Maximization in Combinatorial Auctions

In this section we will attempt to provide a similar lower bound POA argument, but for some types of combinatorial auctions. The main differences from the previous analysis will be the super-polynomial number of player actions (which do not need super-polynomial space to define as input for the optimization problem) and that lower bounds will come from communications and not computational complexity, although it can be argued that problems which are difficult to prove in a communication setting have an underlying computationally hard computational problem.

We begin by defining the combinatorial auction and it's accompanying communication setting. A combinatorial auction setting is defined by:

- n bidders and a set U of m distinct items.
- Each bidder i has a valuation function $v_i : 2^U \rightarrow \mathbb{R}^+$. We will consider valuations that satisfy $v(\emptyset) = 0$ and monotonicity $v(S) \leq v(T)$ for $S \subseteq T$.
- We also assume that each valuation is integral and polynomial bounded from above in m, n.

An allocation is a partition of items $S_1, \ldots, S_n$ where each player gets his corresponding set. The welfare of an allocation is

$$w(S_1, \ldots, S_n) = \sum_{i=1}^{n} v_i(S_)$$

The setting we have just described is not a complete auction, but only a welfare maximization setting.

As we did in the previous chapter, the concrete mechanism we will study is the simultaneous first price auction (or S1A for short). We will auction the m items one by one by first price auctions at the same time. The only difference from the standard first price auction is that we will restrict the bids to be integral and bounded from above by V_{max}, the highest valuations. As such, the action set of every player i is $A_i = 0, 1, \ldots, V_{max}{}^m$. The utility of every player is quasilinear with his valuation following the previous restrictions. In the case of a tie, we will give the item to the lexicographically first player. Note that restricting our bids will not be detrimental, since proving a lower bound on the POA of this auctions will obviously extend to the general setting.

The underlying complexity problem is determining the allocation that maximizes (at least approximately) the welfare. This is irrespective of the S1A and is a property of the welfare maximization setting of the combinatorial auction. Instead of measuring computational complexity (as we did with Congestion Games) we will obtain hardness results from communication complexity.

We will use the standard communication complexity model in this setting, defined in [NS06]. The goal is to have the players communicate in a standardized fashion, in order to reach an good allocation they all agree on. This model is also known as the Number in Hand (NIH) model, where everyone knows only his own valuation v_i. Communication between players can be done in many ways, but image having a blackboard where everyone writes, one at a time, for everyone else to see. Communication is defined with a protocol. There are two kinds of protocols in this model. Deterministic protocols specify what players can say and in which order, hoping that players will eventually reach an agreement over the allocation. Nondeterministic protocols, which we will use, start off by having an oracle (which knows all valuations and has infinite computational powers) write some advice on the blackboard for the rest of the players to read. A protocol is deemed tractable if polynomial in both m and n. Remember, that the input for this setting is $n2^m$ since players have a valuation for each subset of items.

Remember that the communication protocol in the end has to serve as proof of the allocations good approximation of the optimal welfare. As a result, just having the oracle write the optimal allocation is not enough because there is no way for the players to verify it. Since players' valuations contain exponential information, it seems logical that an exponential amount of communication bits is needed as proof. Just to give a taste of such theorems we present the following from [DNS10].

Theorem 18. *Let $\delta > 0$ be any arbitrarily small constant. For subadditive player valuations and every $n \leq m^{0.5-\delta}$, every nondeterministic communication protocol that distinguishes between instances with optimal welfare at least $2n$ and instances with optimal welfare at most $n+1$ requires an exponential amount (in m) of bits in the worst case.*

Before stating the theorem of S1A, we need to point out one more difference. First of all, single items auctions may not have PNE's, but by Nash's theorem they have at least one MNE. On top of that, since players have an exponential number of actions (up to $(V_{max}+1)^m$ we need a more compact way to handle equilibria. For this purpose, we define *approximate* mixed Nash equilibria.

Definition 20 (Approximate MNE)**.** *A product distribution over outcomes $\boldsymbol{\sigma} = \times_i \sigma_i$ is an ϵ-MNE if for each player i:*

$$E_{\mathbf{s}_{-i} \sim \boldsymbol{\sigma}_{-i}}[u_i(s'_i, \mathbf{s}_{-i})] \leq E_{\mathbf{s} \sim \boldsymbol{\sigma}}[u_i(\mathbf{s})] + \epsilon$$

for any $s'_i \in S_i$

We also define a t-uniform mixed strategy for player i as a distribution over at most t actions from A_i.

We use the following important theorem, due to Lipton et al. [LMM03].

Theorem 19. *Let G be a game with n players, each with at most N with all payoffs between $-V_{max}$ and V_{max}. For every $\epsilon > 0$, G has a $(12n^2 \ln(n^2 N))/\epsilon^2$-uniform eV_{max}-MNE.*

All the tools are available to state the main theorem.

Theorem 20. *Let $\mathcal{V}$ denote a set of valuation profiles with all valuations bounded above by V_{max}. Assume that any nondeterministic communication protocol that can distinguish between $v \in \mathcal{V}$ having maximum welfare at least W^* or at most W^*/ρ requires communication exponential in m for sufficiently large m, n.*

Then for every polynomial $p(m,n)$ the worst case POA of $p(m,n)^{-1}V_{max}$-MNE in S1A with valuation profiles in $\mathcal{V}$ is at least ρ.

Proof. Fix a polynomial $p(n,m)$ and consider the following nondeterministic protocol $\mathcal{P}$.

1. Using Lipton's theorem, compute a t-uniform, $p(n,m)^{-1}V_{max}$-MNE $\mathbf{x}$ where $t = (12n^2 \ln(n^2 N))p(n,m)^2$ and $N = (V_{max}+1)^m$. Notice t is polynomial in m, n thus the oracle can broadcast this allocation in polynomial bits.

2. Verify that $\mathbf{x}$ is indeed a $p(n,m)^{-1}V_{max}$-MNE. This is done offline by each player, who then broadcasts that the accepts it.

3. Compute the expected welfare W of $\mathbf{x}$. Again, each player broadcasts his welfare using his private valuation.

4. Output 'yes' if and only if $W > W^*/\rho$

This is a well defined, nondeterministic protocol that uses communication polynomial in m, n. As before, assuming the worst case POA is α we have that in the case of $\alpha < \rho$, the protocol outputs 'yes' whenever there is an allocation with welfare at least W^* and 'no' whenever all allocations have welfare less that W^*/ρ. This contradicts the assumption of needing exponential communication for this task for sufficiently large m. Thus, we have that $\alpha \geq \rho$. □

This theorem essentially reduces proving lower POA bounds for S1A to proving exponential lower bounds for nondeterministic communication protocols. Also note that slight variations of the previous theorem may work for many simple auction types, since the reliance on the S1A is just to easily verify the ϵV_{max}-MNE and it's welfare.

Combining **theorem 20** and **theorem 18** we can show that S1A have a POA of at least 2 for subadditive valuations. This had already been proven by Christodoulou et al. [CKST13], while the upper bound was found by Feldman et al. [FFGL13]. In the following table we summarize POA bounds for different valuations of the S1A.

Valuation	Communication Lower Bound	Upper Bound
Subadditive	2([DNS10])	2
XOS	$\frac{e}{e-1}$([DNS10])	$\frac{e}{e-1}$ ([ST13])
Submodular	$\frac{2e}{2e-1}$ ([DV13])	$\frac{e}{e-1}$ ([ST13])

As we can see, for the XOS case the upper and lower bounds match for S1A. The upper bound proof was detailed in the previous chapter, completing the analysis of the simultaneous first price auction. The first price auction, even though it is not truthful nor has PNE's, is a very simple mechanism that achieves a POA of 2 for general subadditive valuations and about 1.58 for XOS which is exceptional performance.

Chapter 6

Conclusion

6.1 Remarks

We have presented 3 modern algorithmic game theory tools and shown their applications for auctions and congestion games. The problem of allocating different items to players using the first price auction has been thoroughly examined and concluded with a matching upper and lower POA bound of $\frac{e}{e-1}$ for XOS valuations. However, the techniques used to reach these bounds are quite general and we hope the reader will be able to apply them to settings of his own.

6.2 Future Work

Although the methods from chapters 3 and 4 are tailored to work with auctions, the idea behind them could be promising in other settings as well, the most obvious of which changing the social welfare to $SW(a) = \max_i c_i(a)$ and trying to apply similar techniques to scheduling, which still contains a variety of open problems and gaps in known POA bounds. But, as with any approach to scheduling, this may be extremely complicated and time consuming and likely will only work for a restricted setting.

A different approach would be to use some other method for determining player preferences other than payments and try to prove similar composability theorems as those in Chapter 2. For example, one could use resource burning or probabilistic verification and punishment afterwards.

Finally, the approach of Chapter 5 gives rise to new problems in complexity theory generated from game theory, which may contain settings that were not considered interesting enough until now. Also, it may motivate finding more accurate complexity results for more restricted classes of problems, such as congestion games with polynomial coefficients.

Bibliography

[AAE05] Baruch Awerbuch, Yossi Azar, and Amir Epstein. The Price of Routing Unsplittable Flow. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*, STOC '05, pages 57–66, New York, NY, USA, 2005. ACM.

[ADG+06] Sebastian Aland, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann. Exact Price of Anarchy for Polynomial Congestion Games. In Bruno Durand and Wolfgang Thomas, editors, *STACS 2006*, number 3884 in Lecture Notes in Computer Science, pages 218–229. Springer Berlin Heidelberg, 2006.

[Aum74] Robert J. Aumann. Subjectivity and Correlation in Randomized Strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

[BHLR08] Avrim Blum, MohammadTaghi Hajiaghayi, Katrina Ligett, and Aaron Roth. Regret minimization and the price of total anarchy. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 373–382. ACM, 2008.

[BR11] Kshipra Bhawalkar and Tim Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 700–709. SIAM, 2011.

[CBL06] NicolG• Cesa-Bianchi and GG•bor Lugosi. *Prediction, learning, and games*. Cambridge University Press, Cambridge; New York, 2006.

[CDT09] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM (JACM)*, 56(3):14, 2009.

[CK05] George Christodoulou and Elias Koutsoupias. The Price of Anarchy of Finite Congestion Games. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*, STOC '05, pages 67–73, New York, NY, USA, 2005. ACM.

[CKKK11] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 81–90. ACM, 2011.

[CKS08] George Christodoulou, AnnamG•ria KovG•cs, and Michael Schapira. Bayesian combinatorial auctions. In *Automata, Languages and Programming*, pages 820–832. Springer, 2008.

[CKST13] George Christodoulou, AnnamG•ria KovG•cs, Alkmini Sgouritsa, and Bo Tang. Tight Bounds for the Price of Anarchy of Simultaneous First Price Auctions. *arXiv:1312.2371 [cs]*, December 2013. arXiv: 1312.2371.

[Cla71] Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, September 1971.

[DGP09] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

[DNS10] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. *Mathematics of Operations Research*, 35(1):1–13, 2010.

[DV13] Shahar Dobzinski and Jan VondrG•k. Communication Complexity of Combinatorial Auctions with Submodular Valuations. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '13, pages 1205–1215, New Orleans, Louisiana, 2013. SIAM.

[EY10] Kousha Etessami and Mihalis Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM Journal on Computing*, 39(6):2531–2597, 2010.

[Fei09] Uriel Feige. On Maximizing Welfare When Utility Functions Are Subadditive. *SIAM Journal on Computing*, 39(1):122–142, January 2009.

[FFGL13] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 201–210. ACM, 2013.

[Gib73] Allan Gibbard. Manipulation of Voting Schemes: A General Result. *Econometrica*, 41(4):587–601, 1973.

[GK14] Yiannis Giannakopoulos and Elias Koutsoupias. Duality and optimality of auctions for uniform distributions. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 259–276. ACM, 2014.

[Gro73] Theodore Groves. Incentives in Teams. *Econometrica*, 41(4):617–631, July 1973.

[HM10] Sergiu Hart and Yishay Mansour. How long to equilibrium? The communication complexity of uncoupled equilibrium procedures. *Games and Economic Behavior*, 69(1):107–126, May 2010.

[JPY88] David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79–100, August 1988.

[KP09] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. *Computer science review*, 3(2):65–69, 2009.

[LLN01] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial Auctions with Decreasing Marginal Utilities. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*, EC '01, pages 18–28, New York, NY, USA, 2001. ACM.

[LMM03] Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 36–41. ACM, 2003.

[LST12] Renato Paes Leme, Vasilis Syrgkanis, and G•va Tardos. Sequential auctions and externalities. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 869–886. SIAM, 2012.

[MS12] Carol A. Meyers and Andreas S. Schulz. The complexity of welfare maximization in congestion games. *Networks*, 59(2):252–260, 2012.

[Mye81] Roger B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[Nas51] John Nash. Non-Cooperative Games. *Annals of Mathematics*, 54(2):286–295, September 1951.

[NS06] Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192–224, July 2006.

[Ros73] Robert W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, December 1973.

[Rou03] Tim Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.

[Rou09] Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 513–522. ACM, 2009.

[Rou14] Tim Roughgarden. Barriers to near-optimal equilibria. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 71–80. IEEE, 2014.

[ST13] Vasilis Syrgkanis and Eva Tardos. Composable and efficient mechanisms. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 211–220. ACM, 2013.

[Syr12] Vasilis Syrgkanis. Bayesian games and the smoothness framework. *arXiv preprint arXiv:1203.5155*, 2012.

[Vet02] A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings*, pages 416–425, 2002.

[Vic61] William Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance*, 16(1):8–37, March 1961.