EFFICIENCY OF MECHANISMS IN COMPLEX MARKETS

A Dissertation
Presented to the Faculty of the Graduate School
of Cornell University
in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

by
Vasileios Syrgkanis

August 2014

EFFICIENCY OF MECHANISMS IN COMPLEX MARKETS

Vasileios Syrgkanis, Ph.D.

Cornell University 2014

We provide a unifying theory for the analysis and design of efficient simple mechanisms for allocating resources to strategic players, with guaranteed good properties even when players participate in many mechanisms simultaneously or sequentially and even when they use learning algorithms to identify how to play and have incomplete information about the parameters of the game. These properties are essential in large scale markets, such as electronic marketplaces, where mechanisms rarely run in isolation and the environment is too complex to assume that the market will always converge to the classic economic equilibrium or that the participants will have full knowledge of the competition.

We propose the notion of a smooth mechanism, and show that smooth mechanisms possess all the aforementioned desiderata in large scale markets. We further give guarantees for smooth mechanisms even when players have budget constraints on their payments. We provide several examples of smooth mechanisms and show that many simple mechanisms used in practice are smooth (such as formats of position auctions, uniform price auctions, proportional bandwidth allocation mechanisms, greedy combinatorial auctions). We give algorithmic characterizations of which resource allocation algorithms lead to smooth mechanisms when accompanied by appropriate payment schemes and show a strong connection with greedy algorithms on matroids. Last we show how inefficiencies of mechanisms can be alleviated when the market grows large in a canonical manner.

BIOGRAPHICAL SKETCH

Vasilis Syrgkanis was born on September 14th, 1986 in Thessaloniki, Greece. He received a diploma in Electrical Engineering and Computer Science from the National Technical University of Athens in July 2009 and expects to receive a Ph.D. in Computer Science with a minor in Applied Mathematics from Cornell University in August 2014.

To my parents, Christos and Evangelia.

ACKNOWLEDGEMENTS

First and foremost, I am truly indebted to my Ph.D. advisor, Éva Tardos. Her insights and advice throughout my Ph.D. years, on all aspects of research, have been invaluable to me. I would like to thank her for the countless research discussions, for her inspiring ideas she generously shared with me, for being patient with all my research endeavors, even those that were bound to fail and for the many research opportunities she strove to make available to me. She has made my years at Cornell an awesome research experience which I will never forget.

I would also like to thank my professors at Cornell and especially Bobby Kleinberg, Jon Kleinberg and David Shmoys for the many research discussions and feedback. I would also like to thank Larry Blume and Emin Gün Sirer for serving in my committee and for their feedback on my research.

I would like to thank my research collaborators for all their help, patience and for the great time we had doing research. First I would like to thank Renato Paes Leme for being an awesome friend and research collaborator from the very first years of my Ph.D.. His help during the first years has been invaluable and I want to thank him for sharing with me many of his research ideas, many of which resulted in core results of this thesis. I would also like to thank my collaborators at MSR New England and at MSR Cambridge, Christian Borgs, Jennifer Chayes, Michal Feldman, Nicole Immorlica and Brendan Lucier, Yoram Bachrach and Milan Vojnović for three amazing summers.

Before coming to Cornell I benefited from the research guidance of my professors at the National Technical University of Athens, Stathis Zachos, Dimitris Fotakis and Aris Pagourtzis. I would like to thank them for introducing me to the field of algorithmic game theory.

I would also like to deeply thank my parents for encouraging me to pursue a research career and for being supportive of all my decisions. Finally, I would like to thank my wife Marieta for putting up with me and for her support throughout my Ph.D. years, as well as all my friends at Cornell for making Ithaca a fun place to be.

This work was supported in part by a Simons Graduate Fellowship in Theoretical Computer Science, ONR grant N00014-98-1-0589 and NSF grant CCF-0729006 and was performed in part while visiting MSR New England, MSR Cambridge, UK, and the Hebrew University of Jerusalem.

TABLE OF CONTENTS

LIST OF FIGURES

Part I

Introduction and Preliminaries

1
INTRODUCTION

1.1 Allocating Resources to Self-Interested Users

How would you allocate resources in a system so as to maximize the total value of the users? If you knew how much each user values each possible allocation of resources then this would be purely an *optimization problem* (most probably a hard one in each generic formulation) and thereby designing the appropriate algorithm for the approximately efficient allocation of the resources, would be the way to go. However, in most situations, the value that users have for each allocation of resources is private information. Hence, the system also has the task to elicit these parameters or some good approximation of them, taking into account that users will behave *strategically* and *selfishly*. The standard approach to solving the *incentive problem* is to introduce payments. Roughly, the combination of an allocation and a payment rule is what we will refer to as a *mechanism*.

Such *mechanisms* are ubiquitous in both economic systems and large-scale computer systems. From the computer science perspective, they can capture most modern electronic markets such as auction marketplaces (e.g. eBay), online advertisement markets (e.g. Google AdWords etc.), crowdousrcing contests (e.g. topCoder), where payments are implicit in the form of costly effort, bandwidth allocation mechanisms (e.g. Kelly's proportional allocation mechanism) and computing resource sharing mechanisms in the "cloud". From the purely economic perspective they can capture settings such as spectrum auctions, government auctions for natural resources (e.g. timber auctions), art auctions (e.g. Sotheby's) and auctions for financial derivatives (e.g. government bonds).

1.1.1 Desiderata in Large-Scale Distributed Markets

The large-scale nature of modern markets, especially those enabled by computer systems, such as electronic marketplaces, introduces new challenges in the theoretical design and analysis of mechanisms. Though mechanism design is a field with a long and distinguished history, starting from the early works of Vickrey, Clarke and Groves [68, 16, 33] in the 60's and 70's, many of the challenges we list below have not been at the forefront of the field.

Composability in the Presence of Multiple Mechanisms. In most markets listed in the previous section, resources are owned by different entities and many mechanisms are running at the same time, with players simultaneously or sequentially participating in many of them (e.g. different sellers on eBay, different online advertisement platforms). Even if the resources are owned by a single entity (e.g. a single online advertisement platform),it is almost infeasible or impractical to run a global centralized mechanism for all the resources and a more decentralized market structure, where small groups of resources are sold via a separate mechanism, is preferred (e.g. the advertisement slots in each search query are sold separately via the means of an auction, called the Generalized Second Price auction).

In these situations, it is crucial that the market as a whole performs reasonably well, i.e. the global allocation of resources is approximately efficient. Thus the local mechanisms used must satisfy some *composability* property: *local properties that imply local efficiency guarantees for each mechanism in isolation, also directly imply global efficiency guarantees.*

Robustness to Learning Behavior and Incomplete Information. In most large-scale markets, the decision problem that each participant is facing is far too complex to assume with certainty that the market will arrive at the classic economic equilibrium, i.e. a state where no participant wants to unilaterally deviate, aka a Nash equilibrium. Rather we need more robust guarantees even if players use learning algorithms to identify how to play in the game. Such learning behavior will not necessarily lead to a Nash equilibrium and could potentially also lead to correlations in the behavior of participants. Any efficiency guarantee of a mechanism should extend to generic enough models of learning behavior.

Moreover, we cannot expect the participants to know all the parameters of the game (e.g. valuations of opponents). Therefore, the mechanism should also be robust with respect to informational assumptions and should be approximately efficient even when players have only partial information (e.g. distributional beliefs) about these parameters.

Simple Rules with Fast Implementation. The Internet environment allows for running millions of mechanisms, which necessitates the use of very simple and intuitive allocation and payment schemes with a fast implementation. As an example, approximately seven thousand search queries happen on Google's search engine every second and each of these search queries triggers an auction among advertisers for the allocation of the special advertisement slots that will appear together with the "organic" search results. Hence, the auction rule that is used, should be able to be computed in a matter of milliseconds.

1.1.2 Thesis Goal: Robust Efficiency Guarantees in Markets Composed of Simple Mechanisms

The goal of this thesis is to provide a theoretical framework for the design and analysis of simple mechanisms for allocating resources to self-interested and strategic users, with guaranteed good properties even when the users participate in multiple different mechanisms simultaneously or sequentially and even when players use learning algorithms and have incomplete information of the market. One of the key questions we will address is:

> *What properties of local mechanisms guarantee global efficiency in a market composed of such mechanisms and even under learning behavior and incomplete information?*

Traditional mechanism design considered mechanisms only in isolation, an assumption not so realistic in many large-scale markets, where players can cover their needs from multiple different mechanisms. As perfectly summarized by two leading economists of the field in the concluding remarks of their seminal paper on competitive bidding:

> *"Most analyses of competitive bidding situations are based on the assumption that each auction can be treated in isolation. This assumption is sometimes unreasonable."*, Milgrom and Weber, 1982

Moreover, mechanism design has mostly focused on truthful mechanisms, where players are incentivized to truthfully reveal all their private parameters to the mechanism. In an environment with several auctions running simultaneously or sequentially, truthfulness of each individual auction loses its appeal, as the global mechanism is no longer truthful, even if each individual part is. The literature's focus on truthful mechanisms is based on the revelation principle, showing that if there are better non-truthful solutions, the mechanism designer can run this alternate solution on the players' behalf. However, the revelation principle is limited to mechanisms running in isolation: with multiple mechanisms run by different parties, there is no global coordinator to implement the solution. Requiring global coordination between mechanisms is not viable and could lead to complicated coordination problems, such as agreeing on ways to divide the global revenue.

From the analysis perspective, a handful of papers in the economic literature have analyzed properties of strategic outcomes of games arising from selling a set of items via auctions simultaneously or sequentially [55, 22, 53, 9, 30, 6, 58] (a special case of our general setup). However, most of the economics literature has made several simplifying assumptions, such as symmetric user properties or small number of users or complete information of the parameters of the market. The main hurdle in extending the analysis to more realistic settings is analytically solving for the equilibrium. In the general setup that we study, analytically solving for the equilibrium is an impossible task and even more importantly it is not true that the equilibrium of the market is always unique. Instead we will follow the *price of anarchy analysis* from the computer science literature [15, 57, 46, 8, 36] that attempts to analyze the efficiency without solving for the equilibrium, as we will describe in subsequent sections.

1.1.3 A Concrete Example

Simultaneous Item Auctions. Consider an example with two sellers A, B, each having one item for sale. For simplicity, the market has two participants α, β and each participant wants only one item (i.e. is unit-demand). The items that are for sale are not completely identical and the participants exhibit some slight preference: player α has value 2 for item A, value 1 for item B and 2 for the bundle of the two items (since he will only use one of them). Player β prefers item B, having value 2 for item B, value 1 for item A and 2 for the bundle. Obviously the optimal allocation in this market is for each player to win his most preferred item, yielding a total value of 4.

What would happen in this market if each seller was using a second-price auction to sell his item (i.e. the highest bidder wins and pays the second highest bid)? The two participants are playing a game where their strategy is to submit a bid on each of the two items. For simplicity, assume that a Nash equilibrium of the game will arise, i.e. a profile of bids such that no player can gain by deviating. Assuming that the utility that the player derives from the interaction is his value for his allocation minus his total payment, then it is easy to see that the following is one equilibrium: player α bids 1 on item B and 0 on item A and player β bids 1 on item A and 0 on item B. Both players derive a utility of 1 and it is easy to see that no unilateral deviation of a player can lead to a better utility. Thus at the equilibrium, the allocation is suboptimal and the total value is only half of the value of the optimal allocation.

One of the main take-away messages of this example is that the nice properties of a single-item second-price auction in isolation, break the moment there are several mechanisms running simultaneously: in a single-item second price auction it is a dominant strategy for the player to bid his true value for the item irrespective of what the opponents are doing and under such truthful behavior the item will go the highest value player. In the simultaneous auction setting, not only players don't have dominant strategies, but even the concept of truthfulness does not make sense, as the players can no longer express their whole valuation function through their bids. Another observation, in the above example is that if a first-price auction (i.e. winner pays his bid) was used instead of a second price, then every Nash equilibrium with deterministic bids would have resulted in the optimal allocation. Thus mechanisms that seem inferior when studied in isolation might perform better in an environment where multiple mechanisms occur at the same time.

1.1.4 Approach and Main Conclusions

We will approach the problem using techniques from the computer science literature and more specifically, the work on the *price of anarchy*, initiated by the seminal papers of [42, 62]. The price of anarchy literature attempts to quantify the efficiency of all possible strategic outcomes without analytically solving for the equilibrium, but rather simply from the fact that if an outcome is an equilibrium then every deviation of a user must lead to lower utility, a.k.a. the *best-response property*. Our work develops a unifying theory of how to analyze mechanisms via such best-response arguments. Special cases of our theory includes some earlier work on the price of anarchy in specific auction settings [15, 57, 46, 8, 36]. Our work will unify and heavily extend the results in these papers in a single theory on the *price of anarchy of mechanisms*. Apart from the applications we present in the thesis, our theory has been used subsequent to our work, in quantifying the efficiency of mechanisms at equilibrium [17, 14, 5].

More formally, we define the notion of a (λ, μ)-*smooth mechanism* and show that *smooth mechanisms* are approximately efficient and possess all the desired properties of composability and robustness under learning behavior and incomplete information. The definition of a smooth mechanism is based on the existence of a "well-behaved" best response action for each player. Intuitively, the mechanism must admit for each player an "optimal" action, such that no matter what the other players are doing, this action guarantees her a good fraction of her optimal allocation and at a price that is comparable to what is currently being paid for that allocation. Our notion of smoothness is focused on mechanisms where players have quasilinear utilities and is closely related to the notion of smooth games introduced by Roughgarden [59].

Our main result is to show that smooth mechanisms compose well and are robust to incomplete information and learning behavior:

> *A market composed of smooth mechanisms running simultaneously is approximately as efficient as each individual mechanism would have been if run in isolation, when players have complement-free valuations across mechanisms. Efficiency is achieved even in learning outcomes, as well as in Bayesian settings with uncertainty about participants.*

We present several other robustness properties of smooth mechanisms, such as composability when mechanisms are run sequentially rather than simultaneously, efficiency properties when players have budget constraints on the payments they can make and how the inefficiencies of some smooth mechanisms can be alleviated if the market becomes large in a canonical manner.

We further show that many well-studied and used mechanisms are smooth, such as several forms of single-item auctions such as first price and all-pay, some formats of ad-auctions, Kelly's [41, 40] proportional bandwidth allocation mechanism, uniform price auctions, as well as a number of other mechanisms. In that respect, we also present algorithmic characterizations of what algorithms for allocating resources, lead to smooth mechanisms when accompanied with appropriate payment schemes and show a strong connection between smoothness and greedy algorithms under well-behaved resource allocation constraints.

1.2 Our Contributions

1.2.1 Robust Efficiency Guarantees for Mechanisms

We define the notion of a (λ, μ)-smooth mechanism and show that any such mechanism achieves at least a $\frac{\lambda}{\max(1,\mu)}$ fraction of the maximum possible social welfare at every Nash equilibrium. Moreover, this guarantee extends directly to any *coarse correlated equilibrium*, which is a superset of Nash equilibria.

No-Regret Learning (Sections 2.4 and 3.1). As is known coarse correlated equilibria have a strong connection to *no-regret learning* in games. Suppose that the mechanism is played repeatedly with the parameters of every player remaining fixed and the players use some update rule to learn how to play the game. All we assume is that the learning rule satisfies the property that in the long run the player doesn't regret having played a fixed strategy in all periods. Then it is known that the empirical distribution of players' actions of any such no-regret sequence of play will converge to a coarse correlated equilibrium of the static game [12]. Thus the efficiency guarantee of a smooth mechanism directly extends to the average welfare of any such no-regret sequence.

Bayesian Incomplete Information (Section 3.2). In addition, we show that this guarantee extends directly to Bayesian settings of incomplete information, where each player's private parameters are drawn independently from some commonly known distribution. In that setting we define a notion of a *Bayesian coarse correlated equilibrium* and we show that the expected welfare of any such equilibrium is at least $\frac{\lambda}{\max\{1,\mu\}}$ of the expected optimal welfare (in expectation over player parameters). Bayesian coarse correlated equilibria are a superset of Bayes-Nash equilibria [34] and similar to coarse correlated equilibria have a strong connection with no-regret dynamics as we explain below.

No-Regret Learning with Stochastic Parameters (Section 3.3). We consider a situation where the game is played repeatedly and at each iteration each player's private parameters are re-sampled independently from some distribution (unlike in the previous repeated game version, where they were fixed). Equivalently, one can view each player as a population, where each "atom" in the population has some fixed parameter and at each time step one player from each population is picked to play in the mechanism. We show that if each player achieves the no-regret property for each possible parameter instantiation (or equivalently each "atom" in the population achieves the no-regret property), then the limit empirical distribution of play converges almost surely to the set of Bayes coarse correlated equilibria that we defined. Therefore, the efficiency guarantee directly extends to the average welfare of any no-regret sequence in the above Bayesian repeated game setting.

Bandit Learning. The no-regret learning guarantees have the extra robustness properties that for a player to achieve the no-regret property he doesn't need to be aware of any parameters of the game, neither the distributions from which the parameters are re-drawn. There are update rules that the player can invoke (e.g. multiplicative weight updates [3]), that only require for the player to be able to calculate his utility from the action he took at each time step. Thus it suffices to know just his value for the allocation he received and his payment.

Efficiency with the No-overbidding Refinement (Chapter 6). For some mechanisms, such as the second price auction, good performance requires that participants do not bid above their value. It is easy to see that even in a single-item second price auction, there exist Nash equilibria where players overbid and whose welfare is arbitrarily worse than the optimal. However, if we consider only the subset of equilibria where players don't bid above their value for the item, then every Nash equilibrium is efficient.

For such "second-price type" mechanisms, we identify the notion of a weakly smooth mechanism. Weakly smooth mechanisms achieve high welfare at any equilibrium that satisfies a generalization of the non-overbidding assumption that we described above for the case of a single-item second price auction. Moreover, this guarantee is equally robust to the guarantees of smooth mechanisms, in the sense that it extends to learning outcomes and Bayesian incomplete information.

1.2.2 Composability of Mechanisms

Simultaneous Composability of Smooth Mechanisms (Section 4.2). We show that smooth mechanisms compose well in parallel: if we run any number of (λ, μ)-smooth mechanisms simultaneously and players valuations over outcomes of different mechanisms satisfy a complement-free condition that we explain in the next paragraph, then the global market can also be viewed as a (λ, μ)-smooth mechanism, and hence achieves a $\lambda / \max(1, \mu)$ fraction of the maximum social welfare in all Bayesian coarse correlated equilibria and for any independent distributions of player parameters.

Complement-Free Valuations Across Mechanisms (Section 4.3). For our simultaneous composability results, we need to assume that user's valuations have no complements across the different mechanisms. At a high-level all we need to assume for the composability property is that the marginal valuation of a player for an allocation from a specific mechanism can only decrease if more mechanisms come into the market and give him some non-empty allocation.

In more detail, we develop a hierarchy of valuations on outcomes that have no complements across mechanisms. Existing valuation hierarchies consider only valuations on sets of items. We identify analogs of complement-free valuations across mechanisms, without making any assumption about the valuations of players' for outcomes within a mechanism.

We define natural generalizations of submodular, fractionally subadditive, XOS and subadditive valuations over outcomes of different mechanisms. In the context of valuations on sets of items, fractionally subadditive is a superset of submodular valuations, and is known to be equivalent to the class of XOS valuations. We show an equivalent connection among the generalized versions of these functions extending the results of Feige [23] and Lehmann et al. [45].

If smoothness of each local mechanism holds only for some restricted class of local valuations, then we will need to make roughly the same assumption for the component-wise marginals of the valuation of a player across mechanisms. For instance, if the allocation space of a mechanism is partially ordered and the smoothness property holds only when the valuations of players are monotone with respect to the partial order, then we will also need to assume that if we fix the allocation from other mechanisms, the valuation of the player across mechanisms is also monotone with respect to the allocation from the specific mechanism. Similarly, if the allocation space forms a lattice and local smoothness holds only for submodular valuations over the lattice, then we will need to assume that the valuation across mechanisms is submodular with respect to the product lattice of allocations of different mechanisms.

Approximate Composability with Restricted Complements (Section 4.5). We also show that in the presence of complementary valuations, the smoothness of the global market degrades smoothly with the "size" of the complements. For the case of set functions, a natural class of complementary relations are those defined via the means of a positively weighted hypergraph, where the value for a set of items is the total weight of hyperedges contained in the set. Then the size of the complements can be defined as the cardinality of the largest hyperedge.

Based on this intuition we define a novel measure of complementarity of a set function and more generally of a valuation function over outcomes of mechanisms and we show an approximate composability property that degrades smoothly with this measure.

Such restricted complement valuations find good application in spectrum auctions where bands in neighboring geographic region exhibit complementary relations (e.g. have extra value when acquired in conjunction). They also find applications in online advertisement auctions where slots in different parts of a webpage can have a complementary effect, as they create an "impression effect" when acquired in conjunction.

Sequential Composability of Smooth Mechanisms (Chapter 5). We also show that smooth mechanisms compose well sequentially, though for a more restrictive assumption on valuations : if we run any number of (λ, μ)-smooth mechanisms sequentially and a player's value is the maximum valued allocation she got among all mechanisms then the global mechanism achieves welfare at least $\lambda/(\mu+1)$ of the optimal social welfare at every Bayes correlated equilibrium (not coarse correlated equilibrium).

To show this theorem we define a more relaxed smoothness condition, denoted as smoothness via swap deviations and show that the global mechanism satisfies this relaxed $(\lambda, \mu+1)$-smoothness condition. We then show that smooth mechanisms via swap deviations guarantee good efficiency at every correlated equilibrium, hence no-swap regret dynamics (i.e. dynamics where in the long-run players don't regret swapping some action with some other) and even under incomplete information. Our efficiency proof for the incomplete information setting uses a bluffing technique to handle the fact that in a sequential mechanism, past actions might reveal information about the private value of a player.

1.2.3 Budget Constraints

The results discussed so far, assume that a participants utility from the mechanism is equal to his value for his allocation minus his payment, i.e. utilities are quasi-linear with respect to money.

The most common non-quasi-linear valuation is when players have budget constraints on their payments. We extend our results to settings where participants have budget constraints. With budget constraints, maximizing welfare is not an achievable goal, as we cannot expect a low budget participant to be effective at maximizing her contribution to welfare. We define a new benchmark in this setting, which we call the optimal "effective welfare"; capping the contribution of each player to the welfare by their budget.

We show that, subject to a minor strengthening of the smoothness definition, dubbed *conservative smoothness* (which all the known and presented smooth mechanisms satisfy), all our results about efficiency for the case of simultaneous mechanisms carry over to bounds for this benchmark when players have budget constraints. For more details see Chapter 7.

1.2.4 Algorithmic Characterizations of Approximately Efficient Mechanisms

The definition of a smooth mechanism is a semantic one, based on an existential property of a best-response action. It does not directly give algorithmic guidelines about what mechanisms are smooth. An analogue is truthfulness, which is also a semantic property; it is useful to have descriptive algorithmic conditions for truthfulness, such as optimal algorithms (as in the VCG mechanism). *Can we give analogous, useful characterizations of algorithmic conditions that guarantee smoothness?*

A common feature in many of the mechanisms that we show are smooth is the *greediness* of the allocation rule. Indeed, an intuition that arises from the line of work on approximately efficient mechanisms is that greedy algorithms lend themselves well to mechanism design, in the sense that they generate auctions with good performance at equilibrium. We formalize this intuition and provide algorithmic characterizations of smoothness.

Specifically, we show that if a greedy allocation rule is used to allocate resources subject to a matroid constraint, and players have submodular[1] preferences over the resources, then the resulting mechanism is smooth and will achieve a constant fraction of the optimal welfare at every Bayes coarse correlated equilibrium. We also provide similar characterizations for greedy and optimal algorithms when the feasibility constraints are intersections of matroids. For more details see Chapter 8.

[1] Our results actually hold for the more general class of fractionally subadditive preferences.

1.2.5 Efficiency in Large Market Limits

We address the question of whether the efficiency guarantees of a mechanism improve as the game grows large in a canonical way. The intuition is that if a player has a negligible effect in the outcome of the market then any strategic manipulation that he might employ, will not ruin social welfare by much. Hence, it is reasonable to expect that as the market grows large the inefficiency of a mechanism will improve.

We propose a smoothness in the limit framework and show a very general full efficiency in the limit result for the case of simultaneous uniform price auctions, with multiple goods and arbitrary monotone combinatorial valuations, assuming that the supply of each good grows as the number of players grows and that each player fails to arrive in the market with some probability δ. For more details see Chapter 9.

1.2.6 Applications

We show that many well-known auctions are *smooth* and can be analyzed in our framework. We list a few representative examples below, and note that our composition result applies when running any set of such auctions simultaneously or sequentially.

Single Item Auctions (Chapter 10). We show that the first price single item auction is $\left(1-\frac{1}{e},1\right)$-smooth, the all-pay auction is $\left(\frac{1}{2},1\right)$-smooth and the second price auction is weakly and $(1,0,1)$-smooth. We also give a smoothness proof for the hybrid auction in which the winner pays a convex combination of her own bid and the second highest bid. Our framework implies that running m simultaneous first price auctions and bidders have fractionally subadditive valuations and budget constraints achieves efficiency at least $1-\frac{1}{e}$ of the optimal effective welfare. All-pay auctions achieve a guarantee of $\frac{1}{2}$. Second price auctions achieve a guarantee of $\frac{1}{2}$ under the no-overbidding assumption. For sequential auctions with unit-demand bidders and no budget constraints the first price, all-pay and second price auctions give guarantees of $\frac{1}{2}(1-\frac{1}{e})$, $\frac{1}{4}$ and $\frac{1}{2}$ respectively.

Position Auctions (Chapter 11). We analyze position auctions for more general valuation spaces than what has been typically considered [21, 13]. We use the model of Abrams et al [2], where each player i has an arbitrary valuation v_{ij} for appearing at position j, that is monotone in the position. Most of the literature in position auctions has considered valuations of the form $v_{ij} = a_j \gamma_i v_i$, i.e. players have only value per click v_i and their click-through-rate is dependent in a separable way on their quality and on the position. The more general class of valuations can capture settings where players have value both for click and for the impression itself, and settings where the click-through-rates are not separable. We show that the following very simple first price analog of the auction of [2] is $(\frac{1}{2}, 1)$-smooth: solicit bids from the players, allocate positions in order of bids and charge each player his bid. The implied guarantee of $\frac{1}{2}$ holds for simultaneous composition when players have monotone fractionally subadditive valuations and budget constraints. Such valuations capture, for instance, settings where bidders have value v_i only for the first k clicks, or settings where the marginal value per-click of a player decreases with the number of clicks he gets. In addition, a bound of $\frac{1}{4}$ is implied for the sequential composition when bidders value is the maximum value among all impressions he got. In contrast, [2] consider the second price analog of this auction, and show that it always has an efficient Nash equilibrium, but do not consider the price of anarchy. We show that the second price version is weakly $(\frac{1}{2}, 0, 1)$-smooth, implying an efficiency guarantee of $\frac{1}{4}$ for simultaneous and sequential composition of such auctions under the no-overbidding assumption. We also consider other variations of the well-studied GFP and GSP mechanisms for the case when players have only values per click.

Greedy Direct Auctions (Chapter 12). Lucier and Borodin [46] considers combinatorial auctions, whose allocation function is based on a restricted class of greedy c-approximation algorithms. When a first price payment is used, they show that such a greedy auction has a $c + O(\log(c))$ efficiency guarantee. We improve this bound, by showing that this mechanism is $(1 - e^{-1/c}, c)$-smooth implying an efficiency guarantee of at least $\frac{1}{c+0.58}$. This bound extends to the simultaneous composition of such mechanisms when bidders have fractionally subadditive valuations across auctions and budget constraints. For example, when each auctions sells only a small number of items, greedy algorithms can do quite well (giving a $\sqrt{k}$-approximation for arbitrary valuations, if each auction sells at most k items). Observe, that fractionally subadditive valuations across auctions allow for complements within the items of a single greedy auction, hence is more general than just assuming that players have fractionally subadditive valuations over the whole universe of items. We also show that the above analysis is a special case of a more general property of *direct auctions*, i.e. auctions where players can report their whole valuation over allocations.

Bandwidth Allocation Mechanisms (Chapter 13). We consider the single-link bandwidth sharing version of the setting studied by Johari and Tsitsiklis [38] where a set of players want to share a resource: an edge with bandwidth C. Each player has a concave valuation $v_i(x_i)$ for getting x_i units of bandwidth. The mechanism studied in [38] is the well-known Kelly Mechanism [41, 40]: solicit bids b_i, allocate to each player bandwidth proportional to his bid $x_i = \frac{b_i}{\sum_j b_j}$, charge each player b_i. We show that this mechanism is $(2-\sqrt{3}, 1)$-smooth, implying an efficiency guarantee of approximately $1/4$ for Bayes coarse correlated eqilibria. We note that [38] considered only Nash equilibria of the complete information setting. Hence, we extend the analysis to incomplete information. Moreover, the same efficiency guarantee extends to the case when we run such mechanisms simultaneously and players have budget constraints and monotone, lattice-submodular valuations on the lattice defined on $\mathbb{R}^m$ by the coordinate-wise ordering. If the valuations are twice differentiable, being monotone and lattice-submodular translates to: every partial derivative is non-negative and every cross-derivative is non-positive.

Multi-Unit Auctions (Chapter 14). For the setting of multi-unit auctions (i.e. all items are identical) where players have concave utilities in the amount of units they get, we give two smooth mechanisms. Recently, Markakis et al. [50] studied a greedy mechanism and showed a $O(\log(m))$ approximation for the case of mixed Bayes-Nash equilibria under a no-overbidding assumption. We show that a first price version of their mechanism where each player is charged his declared marginal bids for the units he acquired is $\left(\frac{1}{2}\left(1-\frac{1}{e}\right), 1\right)$-smooth, while the uniform price version of [50] is weakly $\left(\frac{1}{2}\left(1-\frac{1}{e}\right), 0, 1\right)$-smooth. Therefore our smooth analysis improves the $O(\log(m))$ bound of [50] to a constant

$\frac{1}{4}\left(1-\frac{1}{e}\right)$ and to $\frac{1}{2}\left(1-\frac{1}{e}\right)$ when a first price payment rule is used. It also extends the analysis to the case of simultaneous uniform price auctions, where players have submodular valuations on the product lattice $\mathbb{N}^m$ of vectors of allocated units of each good.

1.3 Comparison to Related Work

In this section we provide an overview of the main work that is related to the general direction of the thesis. Since our thesis touches several subjects, more specific related work is mentioned in each corresponding section, whenever appropriate.

There has been a long line of research on quantifying inefficiency of equilibria starting from Koutsoupias and Papadimitriou [42] who introduced the notion of the price of anarchy. More recently, this analysis technique has also been used to quantify the inefficiency of auction games, including games of incomplete information. A series of papers, Bikhchandani [9], Christodoulou et al [15], Bhawalkar and Roughgarden [8], Hassidim et al [36], Paes Leme et al [56], Syrgkanis and Tardos [66] studied the efficiency of equilibria of non-truthful combinatorial auctions that are based on running separate item auctions (simultaneously or sequentially) for each item. Lucier and Borodin [46] studied Bayes-Nash Equilibria of non-truthful auctions based on greedy allocation algorithms. Caragiannis et al [13] studied the inefficiency of Bayes-Nash equilibria of the generalized second price auction. All this literature can be thought of as special cases of our framework and all the proofs can be understood as smoothness proofs giving the same or even tighter results. A recent exception is the paper by Feldman et al. [25] giving a tighter bound for simultaneous item-auctions with subadditive bidders, than what would follow from our framework.

Roughgarden [59] proposed a framework, which he calls smoothness in games, and showed that a number of classical price of anarchy results (such as routing and valid utility games) can be proved using this framework. Further, he showed that such efficiency proofs carry over to efficiency of coarse correlated equilibria (no-regret learning outcomes). Nadav and Roughgarden [54] give the broadest solution concept for which smoothness proofs apply. Schoppmann and Roughgarden [61] extend the framework to games with continuous strategy spaces, providing tighter results. However, these papers consider only the full information setting and do not capture several of the auctions described previously. Our definition of a smooth mechanism is closely related to the notion of a smooth game. If utilities of the game were always non-negative (which we only assume in expectation) then a (λ, μ)-smooth mechanism can be thought of as a $(\lambda, \mu - 1)$-smooth game. Moreover, our definition of a smooth mechanism has several technical differences and imposes weaker conditions in some respects, so as to allow us to prove our sequential and simultaneous composability results and also give tight efficiency results for many of the applications described so far.

Recent papers offer extensions of the smoothness framework to incomplete information games. Lucier and Paes Leme [47] introduced the concept of semi-smoothness (inspired by their GSP analysis), and showed that efficiency results shown via semi-smoothness extend to the incomplete information version of the game, even if the types of the players are arbitrarily correlated. Semi-smoothness is a much more restrictive property (for instance, not satisfied by the simultaneous item-bidding auctions) than just requiring that every complete information instance of the mechanism is smooth in the complete information setting and applies mostly to mechanisms where players can express their whole valuation profile through their actions. Moreover, semi-smoothness is a property that does not compose: i.e. local semi-smoothness of each mechanism does not imply global semi-smoothness.

Independent to our work in Syrgkanis [65], Roughgarden [60] also offered a similar to ours direct extension theorem of efficiency guarantees from complete to incomplete information. However, the results in [65] and [60] address efficiency in general games and not mechanisms and for that reason they require a stronger smoothness property (called universal smoothness in [65]), which relates utilities of players with different types in a single inequality. Moreover, none of the previous work addresses the issue of learning under incomplete information and provides guarantees only for the static game of incomplete information and only for Bayes-Nash equilibria. Last, the approach used in these papers cannot capture efficiency in sequential games, such as sequential item auctions, which is achieved by our work.

1.4 Bibliographic Notes

Large part of the thesis appears in published or working research papers [56, 66, 65, 67, 27, 48, 49, 24, 26]. The majority of the results in Chapters 3, 4, 5, 6, 7 and Part III appeared in [67]. The main theorem in Section 3.2 is primarily inspired by the main theorem in [65]. The results of section 3.3 are novel results. The results in Section 4.5 appear in [24]. The results in Section 5.3 first appeared in

[56, 66]. The results in subsection 5.3.1 appear in [27]. The results of Chapter 8 appear in [49]. The results of Chapter 9 appear in [26].

2

PRELIMINARIES

2.1 Notational Conventions

We will use bold letters $\mathbf{x}$ to denote a random variable in some probability space. We will denote with $\Delta(\Omega)$ the space of probability distributions over a finite set Ω. Abusing notation we will use $\mathbf{x}$ to denote both the random variable and its distribution, since the distinction will be clear from the context. Moreover, we will write $\textsc{supp}(\mathbf{x})$ for the support of the distribution of $\mathbf{x}$. We will typically use un-indexed letters x to denote vectors $x = (x_1, \ldots, x_n)$ in some product space. We use $\mathbb{R}_+$ for non-negative real numbers.

2.2 Mechanism Design with Quasi-Linear Preferences

Most of this thesis will be considering the following generic setting: a set of resources are to be allocated to a set of n players. The allocation vector $x = (x_1, \ldots, x_n)$ has to lie in a set of feasible allocation vectors $\mathcal{X}$ that is a subset of a product space of allocations $\mathcal{X} \subseteq \mathcal{X}_1 \times \ldots \times \mathcal{X}_n$. We assume that players can be asked to pay for their allocation and thereby a payment vector $p = (p_1, \ldots, p_n) \in \mathbb{R}_+^n$ will also be decided. The pair $o = (x, p)$ of an allocation vector and a payment vector is referred to as an outcome.

Each player i, has a valuation function that maps an allocation to some non-negative real number: $v_i : \mathcal{X}_i \to \mathbb{R}_+$. We will denote with $\mathcal{V}_i$ the set of allowed valuations for player i and with $\mathcal{V} = \mathcal{V}_1 \times \ldots \times \mathcal{V}_n$ the set of allowed valuation profiles. If a player is given allocation x_i and is asked to pay p_i, then the utility that she derives is:

$$u_i(x_i, p_i; v_i) = v_i(x_i) - p_i \tag{2.1}$$

In other words, players have quasi-linear preferences with respect to money.

We will refer to such a mechanism design setting with quasi-linear preferences via the tuple $(n, \mathcal{X}, \mathcal{V})$. Some important examples captured by this generic formulation of a mechanism design setting are:

1. *Combinatorial auctions*: where $\mathcal{X}_i$ is the power set of items and $\mathcal{X}$ is the subset of this product space such that no item is assigned to more than one player,
2. *Combinatorial public projects*: where $\mathcal{X}_i$ is the power set of potential projects to be built and $\mathcal{X}$ is the subset of the product space such that every coordinate is the same and each coordinate satisfies some constraint based on which projects can be simultaneously built (i.e. a set of simultaneously feasible projects is built and shared by the players),
3. *Position auctions*: where $\mathcal{X}_i$ is the set of positions and $\mathcal{X}$ is the subset of the product space where no two coordinates are assigned the same position,
4. *Bandwidth allocation mechanisms*: where $\mathcal{X}_i$ is the portion of the bandwidth assigned to player i and $\mathcal{X}$ is the subset such that the sum of the coordinates is at most the bandwidth capacity.

Definition 2.2.1 (Mechanism)**.** *Given a mechanism design setting $(n, \mathcal{X}, \mathcal{V})$, a mechanism $\mathcal{M}$ is a tuple $(\mathcal{A}, X, P)$, where $\mathcal{A} = \mathcal{A}_1 \times \ldots \times \mathcal{A}_n$ and $\mathcal{A}_i$ is a set of actions available to player i, $X : \mathcal{A} \to \Delta(\mathcal{X})$ is an allocation function that maps each action profile $a = (a_1, \ldots, a_n)$ to a distribution over feasible allocation vectors $\boldsymbol{x} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$ and*

$P : \mathcal{A} \to \Delta(\mathbb{R}_+^n)$ is a payment function that maps each action profile to a distribution over payment vectors $\boldsymbol{p} = (\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n)$. [1]

We will denote with X_i and P_i the i-th coordinate of the allocation and payment functions respectively and with $U_i^{\mathcal{M}} : \Delta(\mathcal{A}) \times \mathcal{V}_i \to \mathbb{R}$ the expected utility of player i from mechanism $\mathcal{M}$:

$$U_i^{\mathcal{M}}(\mathbf{a}; v_i) = \mathbb{E}_{\mathbf{a}, X_i(\mathbf{a}), P_i(\mathbf{a})} \left[v_i \left(X_i(\mathbf{a}) \right) - P_i \left(\mathbf{a} \right) \right] \tag{2.2}$$

We will also denote the expected revenue of the mechanism as:

$$\mathcal{R}^{\mathcal{M}}(\mathbf{a}) = \sum_{i \in [n]} \mathbb{E}_{\mathbf{a}, P_i(\mathbf{a})}[P_i(\mathbf{a})] \tag{2.3}$$

Voluntary Participation Assumption. Throughout the thesis we will assume that players can always drop out of the mechanism and get 0 utility. Thus we will assume that the action space of each player always contains an exit action, under which he pays nothing and he gets an empty allocation, for which he has zero valuation.

2.3 Efficiency Measure

We will measure efficiency of an outcome $o = (x, p)$ in terms of the social welfare, which is the total value that the players have for the given allocation:

$$w(x; v) = \sum_{i \in [n]} v_i(x_i)$$

[1] We could consider mechanisms where the distribution of allocation and payments is correlated but, due to quasi-linearity of utilities and risk-neutrality, assuming independent payment and allocation distributions is wihtout loss of generality.

Equivalently, due to the quasi-linearity of utilities, the welfare can be viewed as the total utility of the players and the payment of the auctioneer(s), i.e. the total utility of all the participants.

Given a mechanism $\mathcal{M}$ we will denote with $SW : \mathcal{A} \times \mathcal{V} \to \mathbb{R}_+$, the social welfare produced by the mechanism under some action profile:

$$SW^{\mathcal{M}}(a; v) = \sum_{i \in [n]} \mathbb{E}_{X_i(a)}[v_i(X_i(a))] \tag{2.4}$$

For any valuation profile $v \in \mathcal{V}$ we will denote with $x^*(v)$ the optimal allocation, i.e. the allocation that maximizes $w(x; v)$ over all feasible allocations $x \in \mathcal{X}$ and we will denote with

$$\textsc{Opt}(v) = w(x^*(v); v) = \sum_{i \in [n]} v_i(x_i^*(v)) \tag{2.5}$$

2.4 Equilibrium Concepts and the Price of Anarchy

Our goal is to provide robust worst-case guarantees on the social welfare achievable by a given mechanism at *rational outcomes* as compared to the optimal welfare. However, the above statement requires a formal definition of what is a *rational outcome* of a mechanism. In other words, it requires as to define a solution concept that predicts which outcomes will arise if players behave rationally, trying to maximize their own utility. However, the problem that a player faces when choosing his action in a mechanism is not a simple optimization problem, since his allocation and payment depend on the actions of others. Therefore, each mechanism defines a game among the n players and we have to analyze non-cooperative equilibrium solution concepts of the resulting game. Our goal would be to provide guarantees for as large sets of rational outcomes as possible and optimally for outcomes that can arise as the empirical distribution of play of some dynamic process if the mechanism is played repeatedly and players follow some decentralized learning process.

If the game is played only once and players have complete knowledge of all the parameters of the game, such as the valuations of all their opponents, then the standard solution concept is that of a (mixed) Nash equilibrium (NE).

Definition 2.4.1 (Nash Equilibrium - NE). *A vector of randomized actions* $\mathbf{a} \in \Delta(\mathcal{A}_1) \times \ldots \times \Delta(\mathcal{A}_n)$*, is a Nash equilibrium, if no player can benefit by deviating unilaterally, i.e. for any* $i \in [n]$ *and* $a_i' \in A_i$*:*

$$U_i^{\mathcal{M}}(\mathbf{a}; v_i) \geq U_i^{\mathcal{M}}(a_i', \mathbf{a}_{-i}; v_i) \tag{2.6}$$

Price of Anarchy. To quantify the worst-case social welfare guarantee achievable by a mechanism at equilibrium, we will use the well-established notion of the *Price of Anarchy* (POA), which was introduced by Koutsoupias and Papadimitriou [42]. In the context of mechanisms, the POA of Nash equilibria is the largest ratio among all possible Nash equilibria of the resulting game, of the optimal social welfare over the expected social welfare at equilibrium:

$$\text{NE-POA} = \sup_{\mathbf{a} \in \text{NE}} \frac{\text{OPT}(v)}{SW^{\mathcal{M}}(\mathbf{a}; v)} \tag{2.7}$$

In other words, if the NE-POA of a mechanism $\mathcal{M}$ is at most ρ, then the expected social welfare at every Nash equilibrium of the mechanism is at least $\frac{1}{\rho} \cdot \text{OPT}(v)$.

Critiques of the Nash Equilibrium. There are two main critiques about the concept of the Nash equilibrium, which can weaken any guarantees that apply only to Nash equilibrium outcomes of a mechanism.

First, it is not clear *how would the players arrive at such an equilibrium*, i.e. it is not always true that there are natural dynamics such that if the game is played repeatedly then the players will converge to a Nash equilibrium. On the contrary other solution concepts such as the Correlated Equilibrium (CE) (see Definition 2.4.4) or the Coarse Correlated Equilibrium (see Definition 2.4.5) always admit such dynamics and arguably natural ones (c.f. Hart and Mas-Colell [35] for an extensive discussion on convergence of simple dynamics to CE and Blum and Mansour [11] for convergence to CCE).

Second, the *complete information* assumption that players know the valuation of all players is almost unreasonable in settings like electronic markets and other complex large-scale marketplaces. The classic approach in such situations is to assume that players have only distributional beliefs about the valuations of their opponents and maximize their utility only in expectation over their beliefs. This is formalized by the notion of a Bayesian Game, introduced by Harsanyi [34] and the corresponding solution concept of a Bayes-Nash equilibrium (BAYES-NE) (see Definition 2.4.2).

In this thesis we will provide efficiency guarantees for solution concepts that address each of these critiques, thereby offering some form of robustness of our results. We will also provide guarantees for the solution concepts of a Bayes-Correlated Equilibrium (BAYES-CE) (see Definition 3.3.2) and Bayes-Coarse Correlated Equilibrium (BAYES-CCE) (see Definition 3.3.3) that address both critiques simultaneously, i.e. they result from decentralized learning dynamics in incomplete information environments. Figure 2.4 gives a quick overview of the robustness of each solution concept. Guarantees that hold for the concept of BAYES-CCE, which is equivalent to the limit empirical distribution of

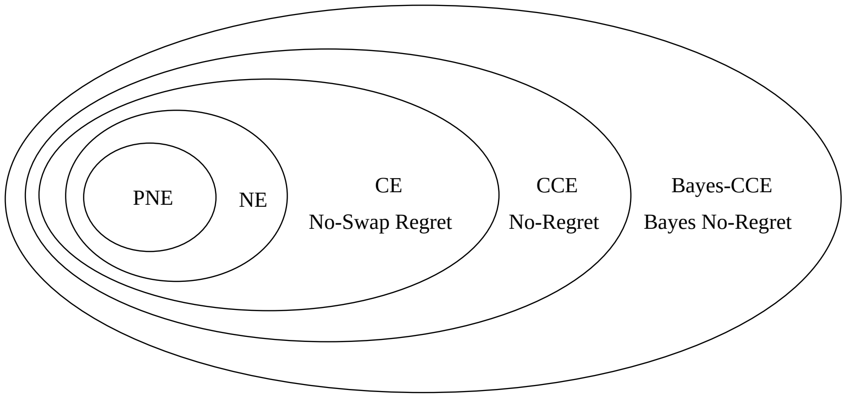

Figure 2.1: Comparison among solution concepts, with respect to robustness of guarantees.

a no-regret sequence under incomplete information, and for any distributional beliefs, is the most robust guarantee. Most of our guarantees would hold for this concept or for the slightly less general one of Bayes-Correlated Equilibrium, which is equivalent to the limit empirical distribution of a no-swap regret sequence under incomplete information. The formal definitions of all these concepts are given in the sections that follow and in Section 3.3.

2.4.1 Incomplete Information and Bayes-Nash Equilibrium

I realized that a major problem in arms control negotiations is the fact that each side is relatively well informed about its own position with respect to various variables ... but may be rather poorly informed about the other side's position in terms of such variables.

– Harsanyi, Nobel Lecture, 1994

In the incomplete information setting, the valuation $\mathbf{v}_i$ of each player is random and drawn independently from some commonly known distribution

$\mathcal{F}_i \in \Delta(\mathcal{V}_i)$ and we will denote with $\mathcal{F} = \mathcal{F}_1 \times \ldots \times \mathcal{F}_n$ the product joint distribution of valuations. In most generality valuations are drawn from some joint correlated distribution, but since almost all of our results would require the independence assumption, we restrict our attention to the case of independent valuation distributions.

Prior to making her decision on which action to play, each player learns his own valuation v_i and nothing else. This defines a Bayesian game, where a player's strategy is a mapping $s_i : \mathcal{A}_i^{\mathcal{V}_i}$ from a valuation $v_i \in \mathcal{V}_i$ to an action $a_i \in \mathcal{A}_i$. We will denote with $\Sigma_i = \mathcal{A}_i^{\mathcal{V}_i}$ the strategy space of each player and with $\Sigma = \Sigma_1 \times \ldots \times \Sigma_n$.

The most commonly used solution concept in Bayesian games, is the following natural generalization of the Nash equilibrium, introduced by Harsanyi [34]:

Definition 2.4.2 (Bayes-Nash Equilibrium - BAYES-NE)**.** *A vector of strategy profiles $s = (s_1, \ldots, s_n) \in \Sigma$ is a Bayes-Nash equilibrium if for any $i \in [n]$, any $v_i \in \mathcal{V}_i$ and $a_i' \in \mathcal{A}_i$:*

$$\mathbb{E}_{\mathbf{v}_{-i}|\mathbf{v}_i=v_i}\left[U_i^{\mathcal{M}}(s(\mathbf{v}); v_i)\right] \geq \mathbb{E}_{\mathbf{v}_{-i}|\mathbf{v}_i=v_i}\left[U_i^{\mathcal{M}}(a_i', s_{-i}(\mathbf{v}_{-i}); v_i)\right] \tag{2.8}$$

Bayes-Nash Price of Anarchy. In the incomplete information setting, we will measure efficiency of a mechanism in expectation over the valuation profile and for any possible distributional beliefs $\mathcal{F}_i$. Our benchmark will also be the expected ex-post optimal welfare over all possible valuation profiles, rather than the ex-post optimal welfare. Specifically, we will analyze the *Bayes-Nash Price of*

Anarchy, defined as:

$$\text{BAYES-NE-POA} = \sup_{\mathcal{F}_1,\ldots,\mathcal{F}_n} \sup_{s \in \text{BAYES-NE}} \frac{\mathbb{E}_{\mathbf{v}}\left[\text{OPT}(\mathbf{v})\right]}{\mathbb{E}_{\mathbf{v}}\left[SW^{\mathcal{M}}(s(\mathbf{v}); \mathbf{v})\right]} \tag{2.9}$$

In other words, if the BAYES-NE-POA of a mechanism $\mathcal{M}$ is at most ρ, then the expected social welfare at every Bayes-Nash equilibrium of the mechanism is at least $\frac{1}{\rho} \cdot \mathbb{E}_{\mathbf{v}}\left[\text{OPT}(\mathbf{v})\right]$.

Similar to the Nash Equilibrium, the Bayes-Nash equilibrium has the drawback that, most of the times, no natural decentralized dynamics is guaranteed to converge to it. In Section 3.3 we will analyze dynamic solution concepts that converge to some notion of equilibrium of the static game that is a superset of the Bayes-Nash equilibrium in the incomplete information setting. Hence, we will address the second critique of Nash equilibria.

2.4.2 Repeated Games, Learning and Correlated Equilibria

> *"The theory of repeated games suggests that collusive behavior in a single auction can be the result of noncooperative behavior in a repeated bidding situation."*
>
> – Milgrom and Weber, 1982, p. 1118

In many scenarios, such as online auctions, bandwidth sharing, or even classic economic settings such as mineral auctions [52], mechanisms are not run only once, but rather are run repeatedly, with players participating in many instances of the same game. In such a repeated setting it is reasonable to analyze the quality of outcomes that arise from learning behavior of the participants.

Moreover, as we shall explain below, in many cases the empirical distribution of actions of such learning behavior will inevitably converge to some static equilibrium concept of the stage game, thereby also offering a dynamic justification to it and giving one explanation of how players could have arrived at such an equilibrium.

One natural form of learning in games is that of no-regret learning, which is a descriptive rather than prescriptive class of learning rules. It simply states that no matter how the players learn how to play in the game, to the least they should, in the long run, have no regret against playing any fixed action all the time.

In this section we will focus on a learning scenario, where the valuations of the players are drawn at the beginning of time and fixed during the learning process. This would correspond, in some sense, to learning under complete information, since as we will discuss, in the long run these dynamics will converge to an *equilibrium* of the static complete information game. In Section 3.3, we will address settings where the valuations of the players that play at each iteration of the learning process are varying and in that setting, learning will correspond to some form of *equilibrium* of the static incomplete information game.

Formally, we analyze the repeated game among a set of n players with valuation profile $v \in \mathcal{V}$, where at each iteration t, each player i picks an action $a_i^t \in \mathcal{A}_i$ and plays mechanism $\mathcal{M}$, incurring utility: $U_i^{\mathcal{M}}(a^t; v_i)$. We will assume that players decide their action a_i^t, using some no-external or no-swap regret algorithm. An algorithm achieves no-external regret, if in the limit as time goes to infinity, a player does not regret switching to playing a fixed action a_i^*. It achieves no-swap regret if the player does have regret against any swap mapping: i.e. swap a_i with a_i' in the history of play.

Definition 2.4.3 (Vanishing External and Swap Regret)**.** *A sequence of action profiles $a^1, \ldots, a^T, \ldots$ has* vanishing external regret *for player i if for any $a_i^* \in \mathcal{A}_i$:*

$$\liminf_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \left(U_i^{\mathcal{M}}(a^t; v_i) - U_i^{\mathcal{M}}(a_i^*, a_{-i}^t; v_i) \right) \geq 0 \tag{2.10}$$

and it has vanishing swap regret *if for any $a_i^* : \mathcal{A}_i \to \mathcal{A}_i$:*

$$\liminf_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \left(U_i^{\mathcal{M}}(a^t; v_i) - U_i^{\mathcal{M}}(a_i^*(a_i^t), a_{-i}^t; v_i) \right) \geq 0 \tag{2.11}$$

Agnostic learning. It is worth pointing that for a player to achieve vanishing external or swap regret, he does not need to know the valuations of his opponents at the beginning of the learning process. In fact the only thing we need to assume is that after each period t, he learns his utility from the action that he played (see Auer et al. [3] for an example of an algorithm that each player can use to achieve vanishing regret under such limited feedback).

Price of Total Anarchy. In the repeated game setting we are interested in comparing the average efficiency of a mechanism at any sequence of play that has vanishing external or swap regret for all players, as compared to the optimal welfare. This is captured by the price of total anarchy introduced by Blum et al.[12]:

$$\sup_{(a^t) \text{ has vanishing external regret}} \limsup_{T\to\infty} \frac{\textsc{Opt}(v)}{\frac{1}{T}\sum_{t=1}^{T} SW^{\mathcal{M}}(a^t; v)} \tag{2.12}$$

and correspondingly for vanishing swap regret sequences.

Optimal Convergence Times and Non-zero Regret. Although we will present our results for the limit empirical distribution as T goes to infinity, where the player actually achieves no-regret, most our results will also have an approximate analogue for any finite time. In other words, suppose that for some time T, the learning algorithm has achieved average external regret of $\epsilon(T)$:

$$\frac{1}{T}\sum_{t=1}^{T}\left(U_i^{\mathcal{M}}(a^t;v_i)-U_i^{\mathcal{M}}(a_i^*,a_{-i}^t;v_i)\right)\geq-\epsilon(T) \tag{2.13}$$

Then if the price of total anarchy is upper bounded by ρ, then in all of our results we will also be able to claim a lower bound on the average welfare of the form:

$$\frac{1}{T}\sum_{t=1}^{T}SW^{\mathcal{M}}(a^t;v)\geq\frac{1}{\rho}\textsc{Opt}(v)-n\cdot\epsilon(T) \tag{2.14}$$

If at each time step the player can query his utility for every possible action $a_i\in\mathcal{A}_i$, then there exist algorithms (e.g. multiplicative weight updates) that can achieve average external regret of $\sqrt{\frac{\log(|\mathcal{A}_i|)}{T}}$ after T time steps. If the player can only query the utility from the action that he actually chose, then there are algorithms that can guarantee regret of $\sqrt{\frac{|\mathcal{A}_i|\log(|\mathcal{A}_i|)}{T}}$ [3]. Thus convergence is sufficiently fast if the action space of the mechanism has reasonable size.

Corresponding static concepts. As is well known (c.f. Blum and Mansour [11], Hart and Mas-Colell [35]), if a sequence of play has vanishing external regret then the empirical distribution of actions converges to the set of coarse correlated equilibria of the static complete information game. While if it satisfies vanishing swap regret then it converges to the set of correlated equilibria.

A correlated equilibrium is a correlated distribution over action profiles, such that no player has incentive to switch to some action a_i', conditional on playing some other action a_i. The correlated equilibrium has the following interpretation: a "correlator" draws a random action profile a according to some correlated probability distribution. Then he proposes to each player i to play action $\mathbf{a}_i$. The distribution is a correlated equilibrium if for each player it is in their own interest to follow the proposed strategy, assuming everyone else also follows that strategy.

In the context of learning, the history of play can be thought as the correlating device. Then there is an obvious connection between a correlated equilibrium and a no-swap regret sequence, since if the player did not regret not having switched to some action a_i' at each time step where previously he was playing action a_i, then this is equivalent to saying that the empirical distribution of the sequence satisfies the correlated equilibrium conditions.

Definition 2.4.4 (Correlated Equilibrium - CE)**.** *A randomized action profile* $\mathbf{a} \in \Delta(\mathcal{A}_1 \times \ldots \times \mathcal{A}_n)$ *is a correlated equilibrium if for any player* $i \in [n]$ *and for any mappings* $a_i'(\mathbf{a}_i)$*:*

$$U_i^{\mathcal{M}}(\mathbf{a}; v_i) \geq U_i^{\mathcal{M}}(a_i'(\mathbf{a}_i), \mathbf{a}_{-i}; v_i) \tag{2.15}$$

A coarse correlated equilibrium is a relaxation of the correlated equilibrium. Specifically, it is a distribution over action profiles such that no player wants to deviate to some other fixed action, unconditionally: i.e. the correlator draws a random action profile a and proposes to each player i, action $\mathbf{a}_i$. Then the player cannot gain by ignoring the correlator's proposal and even before he learns the correlators proposed action, deciding to play some other action a_i'. Observe that because the distribution of a is correlated, then the proposal contains information about the distribution of actions of opponents, and thereby conditioning on the proposal would be a stronger condition. For this reason the set of coarse correlated equilibria is a superset of the set of correlated equilibria.

Similar to correlated equilibria, coarse correlated equilibria have a strong connection to no regret sequences. If a sequence of play satisfies the no-external regret condition, then it is easy to see that the empirical distribution of action profiles satisfies exactly the conditions of a coarse correlated equilibrium.

Definition 2.4.5 (Coarse Correlated Equilibrium - CCE)**.** *A randomized action profile $\mathbf{a} \in \Delta(\mathcal{A}_1 \times \ldots \times \mathcal{A}_n)$ is a coarse correlated equilibrium if for any player $i \in [n]$ and any action $a_i' \in \mathcal{A}_i$:*

$$U_i^{\mathcal{M}}(\mathbf{a}; v_i) \geq U_i^{\mathcal{M}}(a_i', \mathbf{a}_{-i}; v_i) \tag{2.16}$$

Therefore, in order to quantify the price of total anarchy it suffices to provide bounds on the price of anarchy of these static equilibrium concepts, i.e. it suffices to bound the CE-POA and the CCE-POA defined as:

$$\text{CE-PoA} = \sup_{\mathbf{a} \in \text{CE}} \frac{\text{OPT}(v)}{\mathbb{E}_{\mathbf{a}}\left[SW^{\mathcal{M}}(\mathbf{a}; v)\right]} \tag{2.17}$$

and correspondingly for CCE.

Part II

Theory of Smooth Mechanisms

3

SMOOTH MECHANISMS AND EFFICIENCY

"At the core of economics is the concept of efficiency."

– Leibenstein, 1966, p.392

In this chapter we will provide the main definition of a smooth mechanism, which is based on a best-response property. We will give the implications that smoothness has on the efficiency of a mechanism at equilibrium. We will also show that this efficiency extends to learning behavior in a repeated game setting and even when there is incomplete information. At a high level, the results of this chapter can be summarized by the following high-level theorem.

Informal Theorem 1. *Efficiency analysis of mechanisms in the non-robust outcome of complete information pure Nash equilibrium, directly extends to robust learning and incomplete information outcomes, as long as the analysis is based on a simple best-response property.*

3.1 Definition and Efficiency under Complete Information

In this section we introduce the notion of a smooth mechanism. Our notion is similar in spirit to the smoothness of games of Roughgarden [59], but is tailored to the mechanism design settings where players have quasilinear preferences.

Definition 3.1.1 (Smooth Mechanism)**.** *A mechanism $\mathcal{M}$ is (λ,μ)-smooth for some $\lambda,\mu \geq 0$, if for any valuation profile $v \in \mathcal{V}$ and for each player $i \in [n]$ there exists a* *randomized action* $\mathbf{a}_i^*(v)$*, such that for any action profile* $a \in \mathcal{A}$*:*

$$\sum_{i\in[n]} U_i^{\mathcal{M}}(\mathbf{a}_i^*(v), a_{-i}; v_i) \geq \lambda \cdot \textsc{Opt}(v) - \mu \cdot \mathcal{R}^{\mathcal{M}}(a) \tag{3.1}$$

Instantiated to a combinatorial auction setting, the definition of a smooth mechanism has a very natural interpretation as guaranteeing an approximate analog of market cleaning prices for the items. Bikhchandani [9] showed that pure Nash equilibria of the mechanism defined by running independent simultaneous first price auctions for each item, have a one to one correspondence to market equilibria and thereby define market clearing prices, which implies that the outcome is efficient. Market clearing prices are guaranteed when each participant can modify her bid to claim her optimal bundle at the price paid for this bundle in the current solution.

A mechanism is $(1, 1)$-smooth, in essence if the above property is satisfied only in aggregate, but for any outcome of the mechanism, not only at equilibrium. While a (λ, μ)-smooth mechanism satisfies this only approximately, both in terms of the bundle claimed, as well as the price paid for it. In addition, unlike the pure equilibrium analysis, it requires the modified bid to to not depend on the actions of the rest of the players, but can depend on the whole valuation profile, as if the player had complete information on the valuations of his opponents. Thus you can view the smoothness definition as a complete information definition, even though, as we will show in the next section, it directly implies guarantees for incomplete information settings.

We first show that smooth mechanisms have low price of anarchy in the complete information setting and that this result extends to all coarse correlated equilibria (and hence vanishing external regret sequences of play).

Theorem 3.1.2. *If a mechanism is* (λ, μ)*-smooth then every coarse correlated equilibrium achieves social welfare at least* $\frac{\lambda}{\max\{1,\mu\}}$ *of the optimal welfare, i.e.*

$$\text{CCE-PoA} \leq \frac{\max\{1, \mu\}}{\lambda}.$$

Proof. Let $\mathbf{a} \in \Delta(\mathcal{A}_1 \times \ldots \times \mathcal{A}_n)$, be a CCE of the mechanism. Since no players wants to deviate to $\mathbf{a}_i^*(v)$:

$$\sum_{i\in[n]} U_i^{\mathcal{M}}(\mathbf{a}; v_i) \geq \sum_{i\in[n]} U_i^{\mathcal{M}}(\mathbf{a}_i^*(v), \mathbf{a}_{-i}; v_i)$$

By the (λ, μ)-smoothness property for each a in the support of $\mathbf{a}$:

$$\sum_{i\in[n]} U_i^{\mathcal{M}}(\mathbf{a}; v_i) \geq \lambda\textsc{Opt}(v) - \mu\mathcal{R}^{\mathcal{M}}(\mathbf{a})$$

Since players have quasi-linear utilities we have:

$$\mathbb{E}_{\mathbf{a}}\left[SW^{\mathcal{M}}(\mathbf{a}; v)\right] = \sum_{i\in[n]} U_i^{\mathcal{M}}(\mathbf{a}; v_i) + \mathcal{R}^{\mathcal{M}}(\mathbf{a})$$

and thereby:

$$\mathbb{E}_{\mathbf{a}}\left[SW^{\mathcal{M}}(\mathbf{a}; v)\right] \geq \lambda\textsc{Opt}(v) - (\mu - 1)\mathcal{R}^{\mathcal{M}}(\mathbf{a})$$

The result follows if $\mu \leq 1$. When $\mu > 1$, to get the result, we note that $\mathbb{E}_{\mathbf{a},X_i(\mathbf{a})}\left[v_i(X_i(\mathbf{a}))\right] \geq \mathbb{E}_{\mathbf{a},P_i(\mathbf{a})}\left[P_i(\mathbf{a})\right]$, since by the voluntary participation assumption, players always have the possibility to withdraw from the mechanism and get 0 utility. Thus $\mathbb{E}_{\mathbf{a}}\left[SW^{\mathcal{M}}(\mathbf{a}; v)\right] \geq \mathcal{R}^{\mathcal{M}}(\mathbf{a})$. ∎

EXAMPLE 3.1.1. (Single-Item First Price Auction - FPA) In the single item first price auction, one indivisible item is to be allocated among n bidders. Hence, the allocation space is $\mathcal{X}_i \in \{0, 1\}$ and the set of feasible allocations are the ones satisfying $\sum_{i\in[n]} x_i \leq 1$. Each bidder $i \in [n]$ has a value $v_i \in \mathbb{R}_+$ for getting the item. The mechanism asks each player to submit a bid b_i. The highest bidder wins the item and pays his bid (ties are broken arbitrarily).

Lemma 3.1.3. *The first price auction is a* $(1-1/e,1)$*-smooth mechanism.*

Proof. To see why smoothness holds, note that under any valuation profile $v = (v_1,\ldots,v_n)$, the highest value player (wlog player 1) can deviate to submitting a randomized bid $\mathbf{b}_1^*$ drawn from a distribution with density function $f(x) = \frac{1}{v_1-x}$ and support $[0,(1-1/e)v_1]$, while all non-highest value players should just deviate to bidding 0. No matter what the rest of the players are bidding, the utility of the highest bidder from the deviation is:

$$\begin{aligned} U_1^{\text{FPA}}(\mathbf{b}_1^*, b_{-1}; v_1) \geq & \int_{\max_{i\neq 1} b_i}^{\left(1-\frac{1}{e}\right)v_1} (v_1 - x) f(x) dx \geq \left(1-\frac{1}{e}\right) v_1 - \max_i b_i \\ = & \left(1-\frac{1}{e}\right) \textsc{Opt}(v) - \mathcal{R}^{\mathcal{M}}(b) \end{aligned}$$

■

Therefore, we conclude that any coarse correlated equilibrium of a first price auction and hence any vanishing external regret sequence of play in an infinitely repeated first price auction game, will achieve social welfare at least $\left(1-\frac{1}{e}\right) \approx .63$ of the optimal welfare. ■

Comparison to smooth games of [59]. The smoothness property of a mechanism has several similarities with Roughgarden's notion of smoothness of games. To think of a mechanism as a game, we will consider the mechanism also as a player, with utility $\sum_i P_i(a)$ and no strategic decision to make. Our definition of a (λ,μ)-smooth mechanism, is closely related to the game being $(\lambda,\mu-1)$-smooth in the sense of [59], with the difference that we dropped the term $-(\mu-1)\sum_i U_i^{\mathcal{M}}(a;v_i)$ on the right hand side, to make the definition more natural in the context of mechanisms. Note that this change makes the definitions incomparable, as with an arbitrary action profile a, the player utilities

$U_i^{\mathcal{M}}(a; v_i)$ can be negative. Thus a mechanism can be (λ, μ)-smooth under our definition, but the game that it defines might not be $(\lambda, \mu - 1)$-smooth. For instance, the first price auction is not a $\left(1 - \frac{1}{e}, 0\right)$-smooth game (without any technical modification to the smoothness definition, such as viewing the auctioneer as an extra player, assuming players don't overbid etc.), but as we already showed it is a $\left(1 - \frac{1}{e}, 1\right)$-smooth mechanism.

This change also enables our composability of mechanisms results and also enables tighter bounds on the efficiency of mechanisms. Also it allows us to prove a direct extension theorem of the efficiency guarantees to incomplete information games without having to alter the complete information analysis or the complete information smoothness definition. This is not true for smooth games, where to prove extension theorems to incomplete information games, one needs to alter and strengthen the complete information property (see Roughgarden [60] or the universal smoothness of Syrgkanis [65]).

3.2 Extension to Incomplete Information

Next we consider the case where the valuation of each player is drawn from a distribution $\mathcal{F}_i$ over his valuation space $\mathcal{V}_i$. These distributions are independent and are common knowledge. A mechanism $\mathcal{M} = (\mathcal{A}, X, P)$ now defines a game of incomplete information as defined in Section 2.4.1.

The main result of this section is to show that if a mechanism is smooth according to definition 3.1.1 then it achieves a good fraction of the expected optimal social welfare at every Bayes-Nash equilibrium of the incomplete information game, irrespective of the distributions of valuations.

Note that the deviating strategy $\mathbf{a}_i^*(v)$ of player i required by the smoothness property depends on the whole valuation profile v and not only on the valuation of player i. As a result $\mathbf{a}_i^*(v)$ cannot be directly used as deviation for the player in the incomplete information game, as she is not aware of the valuations v_{-i}. We use random sampling to handle the dependence on the values of other players, so as to construct a deviation that depends only on the value of the player and hence is valid for the incomplete information game. Such a random sampling approach was first used in Christodoulou et al. [15] in the context of analyzing the efficiency of simultaneous second price auctions under incomplete information. Here we portray that it is a much more general technique, applying to any smooth mechanism.

Theorem 3.2.1. *If a mechanism $\mathcal{M}$ is (λ, μ)-smooth, then for any vector of independent valuation distributions $\mathcal{F} = (\mathcal{F}_1, \ldots, \mathcal{F}_n)$, every mixed Bayes-Nash Equilibrium has expected social welfare at least $\frac{\lambda}{\max\{1,\mu\}}$ of the expected optimal social welfare, i.e.*

$$\text{BAYES-NE-POA} \leq \frac{\max\{1, \mu\}}{\lambda}.$$

Proof. We will prove it for the case of a pure Bayes-Nash equilibrium $s \in \Sigma$ (the generalization to mixed equilibria is straightforward). Consider the following randomized deviation for each player i that depends only on the information that he has which is his own value v_i and the equilibrium strategies $s(\cdot)$: He random samples a valuation profile $\mathbf{w} \sim \mathcal{F}$. Then he plays according to the randomized action $\mathbf{a}_i^*(v_i, \mathbf{w}_{-i})$, i.e., the player deviates using the randomized action guaranteed by the smoothness property for his true type v_i and the random sample of the types of the others $\mathbf{w}_{-i}$. Since this is not a profitable deviation for player i:

$$\begin{aligned}
\mathbb{E}_{\mathbf{v}}\left[U_i^{\mathcal{M}}(s(\mathbf{v});\mathbf{v}_i)\right] &\geq \mathbb{E}_{\mathbf{v},\mathbf{w}}\left[U_i^{\mathcal{M}}(\mathbf{a}_i^*(\mathbf{v}_i,\mathbf{w}_{-i}),s_{-i}(\mathbf{v}_{-i});\mathbf{v}_i)\right] \\
&= \mathbb{E}_{\mathbf{v},\mathbf{w}}\left[U_i^{\mathcal{M}}(\mathbf{a}_i^*(\mathbf{w}_i,\mathbf{w}_{-i}),s_{-i}(\mathbf{v}_{-i});\mathbf{w}_i)\right] \\
&= \mathbb{E}_{\mathbf{v},\mathbf{w}}\left[U_i^{\mathcal{M}}(\mathbf{a}_i^*(\mathbf{w}),s_{-i}(\mathbf{v}_{-i});\mathbf{w}_i)\right],
\end{aligned}$$

where the first equation is an exchange of variable names and regrouping using independence. Summing over players and using smoothness:

$$\begin{aligned}
\mathbb{E}_{\mathbf{v}}\left[\sum_{i\in[n]} U_i^{\mathcal{M}}(s(\mathbf{v});\mathbf{v}_i)\right] &\geq \mathbb{E}_{\mathbf{v},\mathbf{w}}\left[\sum_{i\in[n]} U_i^{\mathcal{M}}(\mathbf{a}_i^*(\mathbf{w}),s_{-i}(\mathbf{v}_{-i});\mathbf{w}_i)\right] \\
&\geq \mathbb{E}_{\mathbf{v},\mathbf{w}}\left[\lambda\textsc{Opt}(\mathbf{w})-\mu\mathcal{R}^{\mathcal{M}}(s(\mathbf{v}))\right] \\
&= \lambda\mathbb{E}_{\mathbf{w}}\left[\textsc{Opt}(\mathbf{w})\right]-\mu\mathbb{E}_{\mathbf{v}}\left[\mathcal{R}^{\mathcal{M}}(s(\mathbf{v}))\right]
\end{aligned}$$

By quasi-linearity of utility and using the fact that players have the possibility to withdraw from the mechanism, we get the result along the same lines as in the proof of Theorem 3.1.2. ■

EXAMPLE 3.2.1. (Asymmetric First Price Auction) Following on our running example (see Example 3.1.1), consider a first price auction under incomplete information where the valuation v_i of each bidder for the item is drawn independently from some commonly known distribution $\mathcal{F}_i$. Then the strategy space Σ_i is the set of mappings from a value v_i to a bid $b_i(v_i)$, i.e. $\Sigma = \mathbb{R} \to \mathbb{R}$.

Observe that we allow for different bidders to have different valuation distributions. This setting is commonly referred to as an asymmetric first price auction (see [51]) and has a long history in the economic literature (see Section 4.3 of Krishna [43]). It is well established that solving for a Bayes-Nash equilibrium in this setting (unlike the symmetric case) is a daunting task and in most cases a closed form solution to the equilibrium bidding function does not exist. Most papers in the area have considered only special instances of two bidders with specific parametric distributions (i.e. two uniform distributions with different upper and lower bounds [39]).

As we showed in Lemma 3.1.3 the first price auction is a $(1 - 1/e, 1)$-smooth mechanism. Hence, Theorem 3.2.1 implies that any Bayes-Nash equilibrium of the asymmetric first price auction, with arbitrary number of players and arbitrary independent distributions achieves at least $(1 - 1/e) \approx .63$ of the expected optimal welfare. Thus, our smoothness approach yields a bound on the social welfare achievable by any Bayes-Nash equilibrium without having to solve for the equilibrium! In the next chapters we will see how this approach will allow us to go far beyond the single item auction to simultaneous or sequential first price auctions, where solving for an equilibrium seems an even more impossible analytic task.

One natural question is whether the bound of $1-1/e$ is a tight one, i.e. is there an example that achieves such an inefficiency? In the appendix Section A.1 we present a lower bound example where the Bayes-Nash equilibrium achieves $.93$ of the optimal. Closing the gap between upper and lower bound for the Bayes-Nash equilibrium of the asymmetric first price auction is an interesting open question. For the case of correlated valuations we show in the appendix Section A.2 that the bound of $\frac{e}{e-1}$ on the Bayes-Nash price of anarchy is tight. $\frac{e}{e-1}$ is an upper bound even for correlated valuations, since the single-item first price auction actually satisfies the stronger semi-smoothness condition of Lucier et al. [47], which leads to bounds on the price of anarchy that extend to correlated valuations too.

■

3.3 Extension to No-Regret under Incomplete Information

We consider the following setting of a repeated game under incomplete information: there exist n equally sized populations of players. For each $i \in [n]$ population P_i consists of a finite set of r players. Each player $q \in P_i$ has some valuation $V_i(q)$. We denote with $\mathcal{F}_i \in \Delta(\mathcal{V}_i)$ the empirical distribution of values in population P_i and with $\mathbf{v}_i$ a random sample from $\mathcal{F}_i$. In other words, the value of a randomly chosen player from population P_i is distributed according to $\mathcal{F}_i$.

We will describe a repeated game structure and we will argue that any no-regret sequence of the game will, essentially, converge almost surely to the set of BAYES-CCE as time goes to infinity. Then we will show that the efficiency guarantee of a smooth mechanism extends to the set of BAYES-CCE and thereby to the limit average welfare of any no-regret sequence of the repeated game.

Repeated Game 1: Repeated Random Matching Game.

At each iteration t:

1 First, each player $q \in P_i$ in each population P_i picks an action a_{iq}^t. For each population $i \in [n]$ we denote with $\mu_i^t : \mathcal{A}_i^{P_i}$ a function that takes as input a player $q \in P_i$ and outputs his action $\mu_i^t(q) = a_{iq}^t$.

2 Then we pick from each population i one player $q_i^t \in P_i$ uniformly at random. Let $q^t = (q_1^t, \ldots, q_n^t)$ be the chosen profile of players and $\mu^t(q^t) = (\mu_1^t(q_1^t), \ldots, \mu_n^t(q_n^t))$ be the profile of chosen actions.

3 Then each player q_i^t participates in an instance of mechanism $\mathcal{M}$, in the role of player $i \in [n]$, with action $\mu_i^t(q_i^t)$ and experiences a utility of $U_i^{\mathcal{M}}(\mu^t(q^t); V_i(q_i^t))$. The rest of the players experience zero utility.

We assume that each player $q_i \in P_i$, uses some no-regret learning rule to play in this repeated game.

REMARK 3.3.1. We point out that for each player in each population to achieve no-regret he does not need to know the distribution of values in the remainder populations. There exist algorithms that can achieve the no-regret property and simply require an oracle that returns the utility of a player at each iteration. Thus all we need to assume is that each player receives as feedback his utility at each iteration. ■

REMARK 3.3.2. We also note that our results would extend to the case where at each period multiple matchings are sampled independently and players potentially participate in more than one instance of the mechanism $\mathcal{M}$ and potentially with different players from the remaining population. The only thing that the players need to observe in such a setting is their average utility that resulted from their action $\mu_i^t(q) \in \mathcal{A}_i$ from all the instances that they participated at the given period. Such a scenario seems an appealing model in online ad auction marketplaces where players receive only average utility feedback from their bids. ■

Bayesian Price of Total Anarchy. In this repeated game setting we want to compare the average social welfare of any sequence of play where each player uses a vanishing external or swap regret algorithm versus the average optimal welfare. Moreover, we want to quantify the worst-case such average welfare over all possible valuation distributions within each population:

$$\sup_{\mathcal{F}_1,\ldots,\mathcal{F}_n} \limsup_{T\to\infty} \frac{\frac{1}{T}\sum_{t=1}^{T} \textsc{Opt}(V(q^t))}{\frac{1}{T}\sum_{t=1}^{T} SW^{\mathcal{M}}(\mu^t(q^t); V(q^t))} \tag{3.2}$$

Since the valuation of the chosen player is re-drawn independently at every time step, the average optimal welfare will converge almost surely to the expected ex-post optimal welfare $\mathbb{E}_v[\textsc{Opt}(v)]$ of the static incomplete information setting.

In the remainder of this section we will prove the following theorem:

Theorem 3.3.1. *If a mechanism is (λ, μ)-smooth then the Bayesian price of total anarchy of any vanishing external regret sequence of play of the repeated matching game is at most $\frac{\max\{1,\mu\}}{\lambda}$.*

Roadmap of the proof. In the next subsection we define Bayesian generalizations of correlated and coarse correlated equilibria and in subsection 3.3.1, we essentially show that any vanishing external regret sequence of play of the random matching repeated game, will converge almost surely to the Bayesian version of a coarse correlated equilibrium of the static incomplete information game. Therefore the Bayesian price of total anarchy will be upper bounded by the Bayesian price of anarchy of these coarse correlated equilibria. Finally, in the last subsection we show that the price of anarchy bound of smooth mechanisms directly extends to Bayesian coarse correlated equilibria, thereby providing an upper bound on the Bayesian price of total anarchy of the repeated game.

3.3.1 Bayesian Correlated and Coarse Correlated Equilibrium

Defining a correlated equilibrium in an incomplete information game is not such a straightforward task and several notions have been introduced (c.f. Forges [29], Bergemann and Morris [7], Caragiannis et al. [13]). The Bayes-Correlated Equilibrium that we define here is the same as in Section 4.1 of [29] and corresponds to the correlated equilibrium of a corresponding complete information game defined as follows: the utility of a player i in the complete information game is his ex-ante expected utility from the mechanism. The strategy space of player i is $\Sigma_i = \mathcal{A}_i^{\mathcal{V}_i}$. For a strategy profile $s \in \Sigma$, the utility of a player in the complete information game is then:

$$U_i^{\text{ex}}(s) = \mathbb{E}_{\mathbf{v}}\left[U_i^{\mathcal{M}}(s(\mathbf{v}); \mathbf{v}_i\right] \tag{3.3}$$

A BAYES-CE and BAYES-CCE are simply a CE and CCE of this complete information game. For completeness we provide the formal definitions below.

Definition 3.3.2 (Bayes-Correlated Equilibrium - BAYES-CE)**.** *A randomized strategy profile* $\mathbf{s} \in \Delta(\Sigma)$ *is a Bayes-correlated Equilibrium if for every* $s_i' \in \Sigma_i \to \Sigma_i$*:*

$$\mathbb{E}_{\mathbf{s}}\mathbb{E}_{\mathbf{v}}\left[U_i^{\mathcal{M}}(\mathbf{s}(\mathbf{v}); \mathbf{v}_i)\right] \geq \mathbb{E}_{\mathbf{s}}\mathbb{E}_{\mathbf{v}}\left[U_i^{\mathcal{M}}(s_i'(\mathbf{v}_i, \mathbf{s}_i), \mathbf{s}_{-i}(\mathbf{v}_{-i}); \mathbf{v}_i)\right] \tag{3.4}$$

Definition 3.3.3 (Bayes-Coarse Correlated Equilibrium - BAYES-CCE)**.** *A randomized strategy profile* $\mathbf{s} \in \Delta(\Sigma)$ *is a Bayes-coarse correlated Equilibrium if for every* $s_i' \in \Sigma_i$*:*

$$\mathbb{E}_{\mathbf{s}}\mathbb{E}_{\mathbf{v}}\left[U_i^{\mathcal{M}}(\mathbf{s}(\mathbf{v}); \mathbf{v}_i)\right] \geq \mathbb{E}_{\mathbf{s}}\mathbb{E}_{\mathbf{v}}\left[U_i^{\mathcal{M}}(\mathbf{s}_i'(\mathbf{v}_i), \mathbf{s}_{-i}(\mathbf{v}_{-i}); \mathbf{v}_i)\right] \tag{3.5}$$

REMARK 3.3.3. In the above definitions we assumed that the space $\Sigma = \mathcal{A}_1^{\mathcal{V}_1} \times \ldots \times \mathcal{A}_n^{\mathcal{V}_n}$ admits probability distributions. This is not always obvious, since this is a space of functions. When the valuation spaces $\mathcal{V}_i$ are finite and probability distributions on the action spaces are well defined, then this is true, and hence in the above and subsequence definitions that involve probability distributions over functions $\mathcal{A}_i^{\mathcal{V}_i}$ we will assume that the valuation space $\mathcal{V}_i$ is finite. For continuous valuation spaces, the issue is a special case of the problem of defining mixed strategies in extensive form games with infinite information sets and the reader is directed to the classic work of Aumann [4] for one approach. ■

REMARK 3.3.4. We also point out that our definition of BAYES-CCE is inherently different and more restricted than the one defined in Caragiannis et al. [13]. There, a BAYES-CCE is defined as a joint distribution D over $\mathcal{V} \times \mathcal{A}$, such that if $(\mathbf{v}, \mathbf{a}) \sim D$ then for any $v_i \in \mathcal{V}_i$ and $a_i'(v_i) \in \mathcal{A}_i$:

$$\mathbb{E}_{(\mathbf{v},\mathbf{a})}\left[U_i^{\mathcal{M}}(\mathbf{a}; v_i)\right] \geq \mathbb{E}_{(\mathbf{v},\mathbf{a})}\left[U_i^{\mathcal{M}}(a_i'(\mathbf{v}_i), \mathbf{a}_{-i}; v_i)\right] \tag{3.6}$$

The main difference is that the product distribution defined by a distribution in $\Delta(\Sigma)$ and the distribution of values, cannot produce any possible joint distribution over $(\mathcal{V}, \mathcal{A})$, but the type of joint distributions are restricted to satisfy a conditional independence property described by Forges [29]. Namely that a player i's action is conditionally independent of some other player j's value, given player i's type. Such a conditional independence property seems essential for the guarantees that we will present in this thesis to extend to a BAYES-CCE and hence do not seem to extend to the notion given in Caragiannis et al. [13].

However, as we will show in Section 3.3, the no-regret dynamics that we analyze, which are mathematically equivalent to the dynamics in Caragiannis et al. [13], do converge to the smallest set of BAYES-CCE that we define and for which our efficiency guarantees will extend. This extra convergence property is not needed when the mechanism satisfies the stronger *semi-smoothness* property defined in [13] and thereby was not needed to show efficiency bounds in their setting. ■

Last, we also describe the concept of an *Agent-Form Bayes-Correlated Equilibrium* (AGENT-BAYES-CE) (c.f. Forges [29]), which is similar to BAYES-CE, with the restriction that the deviation conditional on a value v_i, can only depend on the previous action $s_i(v_i)$ and not on the whole strategy s_i. In the definition of a BAYES-CE, we assume that the player does not regret switching to some

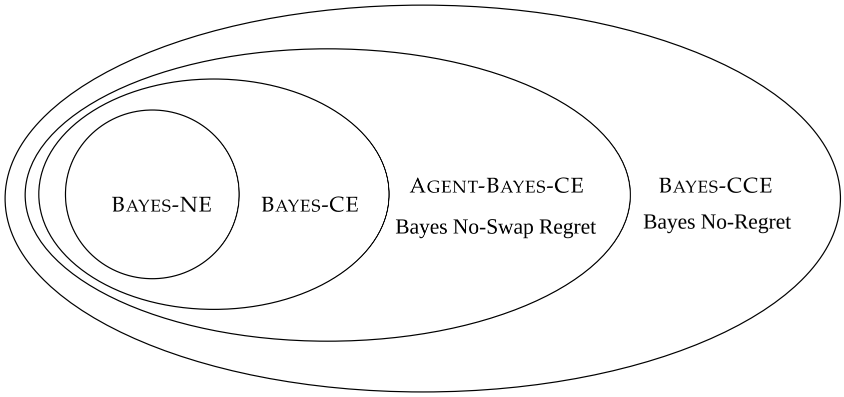

Figure 3.1: Comparison among static solution concepts in the incomplete information setting and connection to no-regret learning under incomplete information.

other strategy s'_i whenever he was using strategy s_i. Thus the switch s'_i can be viewed as a function of the whole strategy s_i, in other words, the deviating action $a'_i = s'_i(v_i)$ of a player with value v_i can depend on the action that a player with value v'_i was previously playing. The AGENT-BAYES-CE does not allow for such dependence. Hence, the set of AGENT-BAYES-CE is a superset of BAYES-CE.

Definition 3.3.4 (Agent-Form Bayes-Correlated Equilibrium - AGENT-BAYES-CE**).** *A randomized strategy profile* $\mathbf{s} \in \Delta(\Sigma)$ *is a Bayes-correlated Equilibrium if for any* $i \in [n]$, $v_i \in \mathcal{V}_i$ *and* $a'_i : \mathcal{A}_i \to \mathcal{A}_i$*:*

$$\mathbb{E}_{\mathbf{s}}\mathbb{E}_{\mathbf{v}|\mathbf{v}_i=v_i}\left[U_i^{\mathcal{M}}(\mathbf{s}(\mathbf{v});\mathbf{v}_i)\right] \geq \mathbb{E}_{\mathbf{s}}\mathbb{E}_{\mathbf{v}|\mathbf{v}_i=v_i}\left[U_i^{\mathcal{M}}(a'_i(\mathbf{s}_i(\mathbf{v}_i)),\mathbf{s}_{-i}(\mathbf{v}_{-i});\mathbf{v}_i)\right] \tag{3.7}$$

Observe that a BAYES-CCE is a superset of AGENT-BAYES-CE, since the type of deviations allowed in the definition of a BAYES-CCE belong to the class of deviations allowed in the definition of a AGENT-BAYES-CE. Thus the stability restrictions are weaker under a BAYES-CCE. Figure 3.3.1 depicts the comparison of the different solution concepts related to games of incomplete information.

3.3.2 Convergence to Bayes-Coarse Correlated Equilibria

For any given sequence of play of the repeated random matching game we define the following sequence of strategy-value pairs (s^t, v^t) where $s : \mathcal{V} \to \mathcal{A}$:

$$s^t(v) = \begin{cases} \mu^t(q^t) & \text{if } V(q^t) = v \\ \text{arbitrary } a \in \mathcal{A} & \text{o.w.} \end{cases} \tag{3.8}$$

and $v^t = V(q^t)$. Then observe that all that matters to compute the average social welfare of the game for any given time step T, is the empirical distribution of (s, v), up till time step T, denoted as D^T, i.e. if $(\mathbf{s}^T, \mathbf{v}^T)$ is a random sample from D^T:

$$\frac{1}{T}\sum_{t=1}^{T} SW^{\mathcal{M}}(\mu^t(q^t); V(q^t)) = \mathbb{E}_{(\mathbf{s}^T,\mathbf{v}^T)}\left[SW^{\mathcal{M}}(\mathbf{s}^T(\mathbf{v}^T); \mathbf{v}^T)\right] \tag{3.9}$$

Lemma 3.3.5 (Almost Sure Convergence to BAYES-CCE). *Let $D \in \Delta(\Sigma \times \mathcal{V})$ be a joint distribution, such that there is a subsequence of $\{D^T\}_T$, converging in distribution to D. Then, almost surely, D is a product distribution, i.e. $D = D_s \times D_v$, with $D_s \in \Delta(\Sigma)$ and $D_v \times \Delta(\mathcal{V})$ such that $D_v = \mathcal{F}$ and $D_s \in$* BAYES-CCE *of the static incomplete information game with distributional beliefs $\mathcal{F}$.*

Proof. For a $q \in P_i$ let $x_i^t(q) = \mathbf{1}_{q_i^t = q}$. Since the sequence has vanishing regret for each player in population P_i, it must be that for any $q_i \in P_i$ and any $s_i^* : \mathcal{V} \to \mathcal{A}$:

$$\sum_{t=1}^{T} x_i^t(q_i)\cdot\left(U_i^{\mathcal{M}}(\mu_i^t(q_i), \mu_{-i}^t(q_{-i}^t); V_i(q_i)) - U_i^{\mathcal{M}}(s_i^*(V(q_i)), \mu_{-i}^t(q_{-i}^t); V_i(q_i))\right) \geq -o(T)$$

Summing over all players in P_i we get:

$$\sum_{t=1}^{T}\left\{U_i^{\mathcal{M}}(\mu^t(q^t); V_i(q_i^t)) - U_i^{\mathcal{M}}(s_i^*(V_i(q_i^t)), \mu_{-i}^t(q_{-i}^t); V_i(q_i^t))\right\} \geq -o(T)$$

Using the definition of s^t from Equation (3.8) and $v_i^t = V_i(q_i^t)$, we can rewrite the above as:

$$\sum_{t=1}^{T}\left\{U_i^{\mathcal{M}}(s^t(v^t); v_i^t) - U_i^{\mathcal{M}}(s_i^*(v_i^t), s_{-i}^t(v_{-i}^t); v_i^t)\right\} \geq -o(T)$$

For any fixed T, let $D_s^T \in \Delta(\Sigma)$ denote the empirical distribution of s^t and let s be a random sample from D_s^T. For each $s \in \Sigma$, let $\mathcal{T}_s \subset [T]$ denote the time steps such that $s^t = s$ for each $t \in \mathcal{T}_s$. Then we get:

$$\mathbb{E}_{\mathbf{s}}\left[\frac{1}{|\mathcal{T}_{\mathbf{s}}|}\sum_{t\in\mathcal{T}_{\mathbf{s}}}\left\{U_i^{\mathcal{M}}(\mathbf{s}(v_i^t);v_i^t) - U_i^{\mathcal{M}}(s_i^*(v_i^t),\mathbf{s}_{-i}(v_{-i}^t);v_i^t))\right\}\right] \geq -\frac{o(T)}{T}$$

For any $s \in \Sigma$, let $\mathcal{T}_{s,v} = \{t \in \mathcal{T}_s : v^t = v\}$. Then we can re-write:

$$\mathbb{E}_{\mathbf{s}}\left[\sum_{v\in\mathcal{V}}\frac{|\mathcal{T}_{\mathbf{s},v}|}{|\mathcal{T}_{\mathbf{s}}|}\left\{U_i^{\mathcal{M}}(\mathbf{s}(v);v_i) - U_i^{\mathcal{M}}(s_i^*(v_i),\mathbf{s}_{-i}(v_{-i});v_i)\right\}\right] \geq -\frac{o(T)}{T} \tag{3.10}$$

Now we observe that $\frac{|\mathcal{T}_{s,v}|}{|\mathcal{T}_s|}$ is the empirical frequency of the valuation vector $v \in \mathcal{V}$, when filtered at time steps where the strategy vector was s. Since at each time step t the valuation vector v^t is picked independently from the distribution of valuation profiles $\mathcal{F}$, this is the empirical frequency of $\mathcal{T}_s$ independent samples from $\mathcal{F}$.

By standard arguments from empirical processes theory, if $\mathcal{T}_s \to \infty$ then this empirical distribution converges almost surely to the distribution $\mathcal{F}$. On the other hand if $\mathcal{T}_s$ doesn't go to ∞, then the empirical frequency of strategy s vanishes to 0 as $T \to \infty$ and therefore has measure zero in the above expectation as $T \to \infty$. Thus for any convergent subsequence of $\{D^T\}$, if D is the limit distribution, then if s is in the support of D, then almost surely the distribution of v conditional on strategy s is $\mathcal{F}$. Thus we can write D as a product distribution $D_s \times \mathcal{F}$.

Moreover, if we denote with **v** the random variable that follows distribution $\mathcal{F}$, then the limit of Inequality (3.10) for any convergent subsequence, will give that:

$$\text{a.s.: } \mathbb{E}_{\mathbf{s}}\mathbb{E}_{\mathbf{v}}\left[U_i^{\mathcal{M}}(\mathbf{s}(\mathbf{v});\mathbf{v}_i) - U_i^{\mathcal{M}}(s_i^*(\mathbf{v}_i),\mathbf{s}_{-i}(\mathbf{v}_{-i});\mathbf{v}_i)\right] \geq 0$$

Thus D_s^T is in the set of BAYES-CCE of the static incomplete incomplete information game among n players, where the valuation profile is drawn from $\mathcal{F}$. ■

Theorem 3.3.6. *The Bayesian price of total anarchy is upper bounded by the Bayesian price of anarchy of Bayesian coarse correlated equilibria.*

Proof. Let $D \in \Delta(\Sigma \times \mathcal{V})$ be a joint distribution, such that there is a subsequence of $\{D^T\}_T$, converging in distribution to D. Then by Lemma 3.3.5, almost surely, D is a product distribution, i.e. $D \in \Delta(\Sigma) \times \Delta(\mathcal{V})$ and that the marginal on $\mathcal{V}$ is equal to $\mathcal{F}$ and the marginal on Σ is a BAYES-CCE of the static incomplete information game with distributional beliefs $\mathcal{F}$.

Therefore, if ρ is the BAYES-CCE $-$ POA of the mechanism, and if $(\mathrm{s}, \mathbf{v})$ is a random sample from D, then almost surely:

$$\mathbb{E}_{\mathbf{s},\mathbf{v}}\left[SW^{\mathcal{M}}(\mathbf{s}(\mathbf{v});\mathbf{v})\right] \geq \frac{1}{\rho}\mathbb{E}_{\mathbf{v}}\left[\textsc{Opt}(\mathbf{v})\right] \tag{3.11}$$

Thus the limit average social welfare of any convergent subsequence will be at least $\frac{1}{\rho}\mathbb{E}_{\mathbf{v}}\left[\textsc{Opt}(\mathbf{v})\right]$, which then implies that almost surely:

$$\liminf_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T} SW^{\mathcal{M}}(\mu^t(q^t); V(q^t)) \geq \frac{1}{\rho}\mathbb{E}_{\mathbf{v}}\left[\textsc{Opt}(\mathbf{v})\right] = \frac{1}{\rho}\lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\textsc{Opt}(V(q^t))$$

Thus for any non-measure zero event, for any ϵ, there exists a $f(\epsilon)$ such that for any $T \geq f(\epsilon)$:

$$\frac{1}{T}\sum_{t=1}^{T} SW^{\mathcal{M}}(\mu^t(q^t); V(q^t)) \geq \frac{1}{\rho}\frac{1}{T}\sum_{t=1}^{T}\textsc{Opt}(V(q^t)) - \epsilon$$

With no loss of generality we can assume that $\mathbb{E}_{\mathbf{v}}\left[\textsc{Opt}(\mathbf{v})\right] > 0$ (o.w. valuations are all zero and theorem holds trivially). The average optimal welfare converges almost surely to $\mathbb{E}_{\mathbf{v}}\left[\textsc{Opt}(\mathbf{v})\right]$, we get that for any non-measure zero event, there exists a $g(\delta)$ such that for $T \geq g(\delta)$, $\frac{1}{T}\sum_{t=1}^{T} \textsc{Opt}(V(q^t))$ is bounded away from zero. Thereby, we can turn the additive error into a multiplicative one, i.e. for any non-measure zero event and for any ϵ' there exists $w(\epsilon')$ such that for any $T \geq w(\epsilon')$:

$$\frac{1}{T}\sum_{t=1}^{T} SW^{\mathcal{M}}(\mu^t(q^t); V(q^t)) \geq \frac{1}{\rho}(1+\epsilon')\frac{1}{T}\sum_{t=1}^{T} \textsc{Opt}(V(q^t))$$

This implies that almost surely:

$$\limsup_{T\to\infty} \frac{\frac{1}{T}\sum_{t=1}^{T} \textsc{Opt}(V(q^t))}{\frac{1}{T}\sum_{t=1}^{T} SW^{\mathcal{M}}(\mu^t(q^t); V(q^t))} \leq \rho = \textsc{Bayes-CCE-PoA}$$

■

3.3.3 Efficiency of Bayes-Coarse Correlated Equilibria

Theorem 3.3.7. *If a mechanism $\mathcal{M}$ is (λ,μ)-smooth, then for any vector of independent distributional beliefs $\mathcal{F}_i$, every Bayesian Coarse Correlated Equilibrium has expected social welfare at least $\frac{\lambda}{\max\{1,\mu\}}$ of the expected optimal social welfare, i.e.*

$$\textsc{Bayes-CCE-PoA} \leq \frac{\max\{1,\mu\}}{\lambda}.$$

Proof. In the proof of Theorem 3.2.1, we essentially showed that for any strategy profile $s \in \Sigma$:

$$\sum_{i\in[n]} \mathbb{E}_{\mathbf{v},\mathbf{w}}\left[U_i^{\mathcal{M}}(\mathbf{a}_i^*(\mathbf{v}_i,\mathbf{w}_{-i}), s_{-i}(\mathbf{v}_{-i}); \mathbf{v}_i)\right] \geq \lambda\mathbb{E}_{\mathbf{w}}\left[\textsc{Opt}(\mathbf{w})\right] - \mu\mathbb{E}_{\mathbf{v}}\left[\mathcal{R}^{\mathcal{M}}(s(\mathbf{v}))\right]$$

where $\mathbf{a}_i^*(v)$ is the smoothness deviation and where w is an independent random sample from $\mathcal{F}$.

Let s $\in \Delta(\Sigma)$, be a $\textsc{Bayes}$-CCE. Since no player i, wants to deviate to any strategy $s_i^*(v_i)$ in the support of the randomized strategy $\mathbf{a}_i^*(\mathbf{v}_i,\mathbf{w}_{-i})$, we get that:

$$\mathbb{E}_{\mathbf{s}}\mathbb{E}_{\mathbf{v}}\left[U_i^{\mathcal{M}}(\mathbf{s}(\mathbf{v}); \mathbf{v}_i)\right] \geq \mathbb{E}_{\mathbf{s}}\mathbb{E}_{\mathbf{v},\mathbf{w}}\left[U_i^{\mathcal{M}}(\mathbf{a}_i^*(\mathbf{v}_i,\mathbf{w}_{-i}), s_{-i}(\mathbf{v}_{-i}); \mathbf{v}_i)\right]$$

Combining the above two inequalities, we get:

$$\mathbb{E}_{\mathbf{s}}\mathbb{E}_{\mathbf{v}}\left[U_i^{\mathcal{M}}(\mathbf{s}(\mathbf{v});\mathbf{v}_i)\right] \geq \lambda\mathbb{E}_{\mathbf{w}}\left[\textsc{Opt}(\mathbf{w})\right] - \mu\mathbb{E}_{\mathbf{s}}\mathbb{E}_{\mathbf{v}}\left[\mathcal{R}^{\mathcal{M}}(\mathbf{s}(\mathbf{v}))\right]$$

By quasi-linearity of utility and using the fact that players have the possibility to withdraw from the mechanism, we get the result, by standard manipulations (see Theorem 3.1.2). ■

4
SIMULTANEOUS COMPOSABILITY

Most analyses of competitive bidding situations are based on the assumption that each auction can be treated in isolation. This assumption is sometimes unreasonable.

– Milgrom and Weber, 1982, p. 1117

Mechanisms rarely run in isolation but rather, several mechanisms take place simultaneously and players typically have valuations that are complex functions of the outcomes of different mechanisms.

In this chapter we analyze a formal model of such a simultaneous occurring mechanism setting and show the following informal theorem:

Informal Theorem 2. *If each individual mechanism is (λ, μ)-smooth and the value of a player over allocations of mechanisms satisfies a complement-free assumption, then the global market consisting of all mechanisms achieves welfare at equilibrium at least $\frac{\lambda}{\max\{1,\mu\}}$ of the optimal.*

The sections in this chapter will elaborate on what we mean by a *complement-free assumption* across mechanisms and will provide the formal proof of this theorem.

4.1 Simultaneous Composition Framework

We consider the following setting: there are n bidders and m mechanisms. Each mechanism $\mathcal{M}_j = (\mathcal{A}^j, X^j, P^j)$ is defined on its own mechanism design setting $(n, \mathcal{X}^j, \mathcal{V}^j)$.

We assume that a player has a valuation over vectors of outcomes from the different mechanisms: $v_i : \mathcal{X}_i \rightarrow \mathbb{R}^+$ where $\mathcal{X}_i = \mathcal{X}_i^1 \times \ldots \times \mathcal{X}_i^m$. A player's utility is still quasi-linear in this extended setting in the sense that his utility from an allocation vector $x_i = (x_i^1, \ldots, x_i^m)$ and payment vector $p_i = (p_i^1, \ldots, p_i^m)$ is given by:

$$u_i(x_i, p_i; v_i) = v_i(x_i^1, \ldots, x_i^m) - \sum_{j=1}^{m} p_i^j \tag{4.1}$$

In this chapter we will consider the case where all mechanisms take place simultaneously (see Chapter 5 for the case of sequential mechanisms). Hence, a player's action space in the game is to report an action a_i^j at each mechanism $j \in [m]$.

The simultaneous composition of m mechanisms can be viewed as a *global mechanism* $\mathcal{M} = (\mathcal{A}, X, P)$, where $\mathcal{A}_i = \mathcal{A}_i^1 \times \ldots \times \mathcal{A}_i^m$, $\mathcal{X} = \mathcal{X}^1 \times \ldots \times \mathcal{X}^m$, $X : \mathcal{A} \rightarrow \mathcal{X}$ is defined as $X(a) = (X^1(a^1), \ldots, X^m(a^m))$ and $P(a) = \sum_{j \in [m]} P^j(a^j)$. Our goal is to give properties of the individual mechanisms that guarantee efficiency of the global mechanism.

4.2 Composability of Smooth Mechanisms

In order to infer good properties of the global mechanism from properties of individual mechanisms, we will need to assume that player valuations satisfies some no complement assumption across outcomes of different mechanisms. Roughly speaking this means that winning an allocation in some mechanism does not increase a player's marginal valuation in a different mechanism. In section 4.5 we will relax this assumption to the case of restricted complements.

When each mechanism is a single item auction, then the value of a player is a set function and the no complement condition is captured by well-understood classes of valuations, such as submodular, fractionally subadditive or XOS, subadditive etc. (see Lehmann et al. [45] for an overview). However, in our setting the valuation of a player is a function on an abstract product space of allocations. Hence, we need to define generalizations of these classes of valuations.

At a high level, for our simultaneous composability theorem we will make the following type of assumption: assume for the moment that a mechanism can be absent from the market (though we will not need it for our main theorem), then we require that the marginal valuation of player for any allocation from some mechanism j, does not increase if a mechanism j' enters the market and awards player i any allocation $x_i^{j'} \in \mathcal{X}_i^j$. This can be viewed as a natural generalization of submodular valuations when the allocation spaces are binary $\mathcal{X}_i^j \in \{0,1\}$. Observe that our generalization of submodular valuations across mechanisms makes no assumption on how the valuation behaves across different allocations from an individual mechanism.

We will show in the next sections that if the valuation space $\mathcal{V}_i^j$ of each mechanism admits single-minded valuations (i.e. positive value only for one allocation) then any valuation that satisfies the above complement-free condition, can be written as a maximum over additively separable valuations, i.e. $v_i(x_i) = \max_{\ell\in\mathcal{L}} \sum_{j\in[m]} v^{j,\ell}(x_i^j)$. We will refer any such valuation as XOS across mechanisms.

To allow for composability of mechanisms where the valuation space does not include single-minded valuations, we will use a generalization of the class of XOS valuation defined as follows: for any product space $\mathcal{C}_i = \mathcal{C}_i^1 \times \ldots \times \mathcal{C}_i^m$, where $\mathcal{C}_i^j$ is a class of valuations over allocations from mechanism j, we will say that a mechanism is XOS-$\mathcal{C}_i$ if it can be expressed as a maximum over additively separable valuations, where the induced valuations $v^{j,\ell}$ fall into the class $\mathcal{C}_i^j$. In Section 4.3, we will give several theorems on the expressiveness of such functions when the classes $\mathcal{C}_i^j$ are interesting special cases, such as all monotone valuations with respect to some partial order on the allocations or the class of all monotone submodular valuations with respect to some lattice defined on the allocation space. Moreover, we will provide the natural generalization of other classes of set functions to products of allocation spaces and show that they are a subclass of XOS valuations, or that they can be well-approximated with XOS valuations.

Definition 4.2.1 (XOS-$\mathcal{C}$)**.** *For any $\mathcal{C} = \mathcal{C}^1 \times \ldots \times \mathcal{C}^m$, with $\mathcal{C}^j \subseteq (\mathcal{X}_i^j \to \mathbb{R}_+)$, a valuation $v : \mathcal{X}_i \to \mathbb{R}_+$ is* XOS-$\mathcal{C}$ *across mechanisms if there exist an index set $\mathcal{L}$ (potentially infinite) of additively separable valuations, such that:*

$$v(x_i) = \max_{\ell\in\mathcal{L}} \sum_{j\in[m]} v^{j,\ell}(x_i^j). \tag{4.2}$$

and for all $\ell \in \mathcal{L}$: $v^{j,\ell} \in \mathcal{C}^j$.

The latter is the generalization of the class of XOS set functions, which are functions $f : 2^{[m]} \rightarrow \mathbb{R}_+$, that can be written as a maximum over additive functions: $f(S) = \max_{\ell \in \mathcal{L}} \sum_{j \in S} w^{j,\ell}$. These are functions that can be written as the XOR of ORs of singleton set functions (see Lehmann et al [45])[1] and they are a superset of submodular set functions. In our generalization we have replaced the signleton valuations, with value functions that depend only on the allocation of a single mechanism $j \in [m]$ and such that these value functions belong to some class of functions $\mathcal{C}^j$. The latter restriction to a class of the local valuations did not even make sense in the case of set functions.

Theorem 4.2.2 (Simultaneous Composition)**.** *Consider a simultaneous composition of m (λ, μ)-smooth mechanisms and let $\mathcal{C}_i = \mathcal{V}_i^1 \times \ldots \times \mathcal{V}_i^m$. If the valuation $v_i : \mathcal{X}_i \rightarrow \mathbb{R}^+$ of each player across mechanisms is* XOS-$\mathcal{C}_i$*, then the global mechanism is also (λ, μ)-smooth.*

Proof. First we show a Lemma, which captures the essence of the proof. Based on the Lemma, it suffices to show composability when the players valuations is additively separable, which is relatively straightforward.

Lemma 4.2.3. *Suppose that a mechanism is (λ, μ)-smooth when valuation profiles come from some class $\mathcal{V}$. Then it is also (λ, μ)-smooth when the valuation of each player i comes from the class* $\max -\mathcal{V}_i$*, consisting of any valuation which can be written as:*

$$v_i(x) = \max_{\ell \in \mathcal{L}_i} v_i^\ell(x) \tag{4.3}$$

with $v_i^\ell \in \mathcal{V}_i$ for each $\ell \in \mathcal{L}_i$.

Proof. Fix a valuation profile $v = (v_1, \ldots, v_n) \in \max -\mathcal{V}$. We remind that $\textsc{Opt}(v)$ is the optimal welfare for v and $x_i^*(v)$ the allocation of player i in the welfare

[1]The XOR of two set functions is defined as: $(v_1 \oplus v_2)(S) = \max\{v_1(S), v_2(S)\}$ and the OR is defined as $(v_1 \vee v_2)(S) = \max_{T \subseteq S} v_1(T) + v_2(S - T)$. A singleton set function is any function of the form: $v(S) = w^j \cdot \mathbf{1}\{j \in S\}$.

optimal allocation. For each player i, denote with $v_i^* \in \mathcal{V}_i$, the valuation v_i^ℓ such that:

$$v_i^*(x_i^*(v)) = \max_{\ell \in \mathcal{L}_i} v_i^\ell(x_i^*(v))$$

Thus $v_i(x_i^*(v)) = v_i^*(x_i^*(v))$ and for any $x_i \in \mathcal{X}_i$: $v_i(x_i) \geq v_i^*(x_i)$. Moreover, let $v^* = (v_1^*, \ldots, v_n^*) \in \mathcal{V}$.

Since the mechanism is (λ, μ)-smooth when valuations come from class $\mathcal{V}$, for each player $i \in [n]$, there exists a randomized action $\mathbf{a}_i^*(v^*) \in \Delta(\mathcal{A}_i)$, such that for any action profile $a \in \mathcal{A}$:

$$\sum_{i \in [n]} U_i^{\mathcal{M}}(\mathbf{a}_i^*(v^*), a_{-i}; v_i^*) \geq \lambda \cdot \textsc{Opt}(v^*) - \mu \cdot \mathcal{R}^{\mathcal{M}}(a)$$

Since for any $x_i \in X_i$: $v_i(x_i) \geq v_i^*(x_i)$, by quasi-linearity of utilities, it is easy to see that for any $a \in \mathcal{A}$: $U_i^{\mathcal{M}}(a; v_i) \geq U_i^{\mathcal{M}}(a; v_i^*)$. Thus:

$$\sum_{i \in [n]} U_i^{\mathcal{M}}(\mathbf{a}_i^*(v^*), a_{-i}; v_i) \geq \sum_{i \in [n]} U_i^{\mathcal{M}}(\mathbf{a}_i^*(v^*), a_{-i}; v_i^*) \geq \lambda \cdot \textsc{Opt}(v^*) - \mu \cdot \mathcal{R}^{\mathcal{M}}(a)$$

Moreover, $\textsc{Opt}(v^*) \geq w(x^*(v); v^*)$. Since $v_i(x_i^*(v)) = v_i^*(x_i^*(v))$: $w(x^*(v); v^*) = w(x^*(v); v) = \textsc{Opt}(v)$. Thus: $\textsc{Opt}(v^*) \geq \textsc{Opt}(v)$. Combining with the previous inequality gives us:

$$\sum_{i \in [n]} U_i^{\mathcal{M}}(\mathbf{a}_i^*(v^*), a_{-i}; v_i) \geq \lambda \cdot \textsc{Opt}(v) - \mu \cdot \mathcal{R}^{\mathcal{M}}(a)$$

Thus the strategies $\mathbf{a}_i^*(v^*)$ (i.e. deviate as if your valuation was the one matching your value for the optimal allocation) constitute the actions required by the smoothness definition. Since the above holds for any initial valuation profile $v \in \max -\mathcal{V}$, the Lemma follows. ■

By Lemma 4.2.3 it suffices to prove the theorem for the case where the valuation of each player i is of the form:

$$v_i(x) = \sum_{j\in[m]} v_i^j(x_i^j) \tag{4.4}$$

But in this case the utility of a player essentially decomposes into the sum of his utilities from each mechanism $j \in [m]$. Therefore, if we denote with $v^j = (v_1^j, \ldots, v_n^j)$ the local valuation profile for each mechanism $j \in [m]$, by considering $\mathbf{a}_i^*(v) = (\mathbf{a}_i^{1,*}(v^1), \ldots, \mathbf{a}_i^{m,*}(v^m))$, then by (λ, μ)-smoothness of each mechanism, for any action profile $a \in \mathcal{A}$:

$$\begin{aligned}
\sum_{i\in[n]} U_i^{\mathcal{M}}(\mathbf{a}_i^*(v), a_{-i}; v_i) = & \sum_{j\in[m]}\sum_{i\in[n]} U_i^{\mathcal{M}_j}(\mathbf{a}_i^{j,*}(v^j), a_{-i}^j; v_i^j) \\
\geq & \sum_{j\in[m]} \lambda \textsc{Opt}^j(v^j) - \mu \mathcal{R}^{\mathcal{M}_j}(a^j) \\
= & \ \lambda \textsc{Opt}(v) - \mu \mathcal{R}^{\mathcal{M}}(a)
\end{aligned}$$

Thus the deviations $\mathbf{a}_i^*(v)$, which correspond to the independent local randomized deviations designated by smoothness of each individual mechanism, imply smoothness of the global mechanism.

■

In the last part of the thesis we will give a number of applications of this result, while at the end of this chapter we will give a concrete example of how the theorem applies to the case of simultaneous single-item first price auctions.

Approximate XOS valuations and approximate composability. To enlarge the applicability of our theorem, it is useful to also consider a relaxation of the theorem, when the valuations are not XOS $-\mathcal{V}_i$, but rather are approximated by them.

Definition 4.2.4 (β-XOS-$\mathcal{C}_i$)**.** *A valuation* $v_i : \mathcal{X}_i \to \mathbb{R}_+$ *is* $\beta - \text{XOS} - \mathcal{C}_i$ *across mechanisms for some* $\beta \geq 1$*, if there exist an* XOS-$\mathcal{C}_i$ *valuation* $\tilde{v}_i$*, such that for all* $x_i \in \mathcal{X}_i$*:*

$$\beta \cdot \tilde{v}_i(x_i) \geq v_i(x_i) \geq \tilde{v}_i(x_i) \tag{4.5}$$

An easy adaptation of the proof of Theorem 4.2.2 yields the following approximate version of it:

Theorem 4.2.5. *Consider a simultaneous composition of* m (λ, μ)*-smooth mechanisms and let* $\mathcal{C}_i = \mathcal{V}_i^1 \times \ldots \times \mathcal{V}_i^m$*. If the valuation* $v_i : \mathcal{X}_i \to \mathbb{R}^+$ *of each player across mechanisms is* β-XOS-$\mathcal{C}_i$*, then the global mechanism is* $\left(\frac{\lambda}{\beta}, \mu\right)$*-smooth.*

The essence of the proof of this theorem is a relaxed version of Lemma 4.2.3, which will be useful in other parts of the thesis.

Lemma 4.2.6. *Suppose that a mechanism is* (λ, μ)*-smooth when valuation profiles come from some class* $\mathcal{V}$*. Then it is* $\left(\frac{\lambda}{\beta}, \mu\right)$*-smooth when the valuation of each player* i *comes from the class* β-$\max$-$\mathcal{V}_i$*, consisting of any valuation which satisfies:*

$$\beta \cdot \max_{\ell \in \mathcal{L}_i} v_i^\ell(x) \geq v_i(x) \geq \max_{\ell \in \mathcal{L}_i} v_i^\ell(x) \tag{4.6}$$

for some index set $\mathcal{L}_i$ *and with* $v_i^\ell \in \mathcal{V}_i$ *for each* $\ell \in \mathcal{L}_i$*.*

4.3 Complement Free Valuations across Mechanisms

In this section we delve into the structure of complement-free valuations across mechanisms and examine the expressiveness of XOS valuations across mechanisms.

Since we focus on the valuation of a specific player i, for notational simplicity we will drop the index i in the current section. Hence, we will analyze classes of complement free valuations $v : \mathcal{X} \to \mathbb{R}_+$ defined on a product space of allocations $\mathcal{X} = \mathcal{X}_1 \times \ldots \times \mathcal{X}_m$. In a mechanism composition setting, $\mathcal{X}_j$ is the set of possible allocations to some player i from mechanism $\mathcal{M}_j$.

4.3.1 Fractionally Subadditive $\equiv$ XOS across Mechanisms

We will first analyze the expressiveness of the class of unconstrained XOS valuations, i.e. XOS-$\mathcal{C}$ where $\mathcal{C}_j$ includes all possible valuation functions and their approximate version of β-XOS. We will define the class of fractionally subadditive valuations across mechanisms, which is a generalization of fractionally subadditive set functions. Then we show that it is equivalent to the class of unconstrained XOS functions, generalizing the result of Feige [23], who showed the result for the case of set functions. In fact we show that it is equivalent to the class of XOS-$\mathcal{C}$ functions where $\mathcal{C}_j$ is any set that contains all *single-minded valuations*, i.e. valuations that are non-zero at only one allocation $x_j \in \mathcal{X}_j$ and 0 for any other allocation.

Definition 4.3.1 (Fractionally Subadditive). *A valuation is fractionally subadditive across mechanisms if*

$$v(x) \leq \sum_{\hat{x} \in \mathcal{X}} \alpha_{\hat{x}} v(\hat{x}),$$

whenever each coordinate x_j is covered by coefficients $\alpha = (\alpha_{\hat{x}})_{\hat{x} \in \mathcal{X}}$, that is $\sum_{\hat{x}: x_j = \hat{x}_j} \alpha_{\hat{x}} \geq 1$.

For comparison, in the case of set functions, a function $f : 2^{[m]} \to \mathbb{R}_+$ is fractionally subadditive if for any set $S \subseteq [m]$ and for any fractional cover $(\alpha, \mathcal{T})$

of S (i.e. a weighted collection of sets $\mathcal{T}$ such that $\sum_{T\in\mathcal{T}:j\in T}\alpha_T \geq 1$ for all $j \in S$), $v(S) \leq \sum_{T\in\mathcal{T}}\alpha_T v(T)$.

The main theorem of this section is that Fractionally Subadditive $\equiv$ XOS. We will actually show the stronger theorem that a relaxed version of fractionally subadditive valuations, denoted as *β-fractionally subadditive* is equivalent to the class of β-XOS.

Definition 4.3.2 (β-Fractionally Subadditive)**.** *A valuation is β-fractionally subadditive across mechanisms if*

$$v(x) \leq \beta \cdot \sum_{\hat{x}\in\mathcal{X}} \alpha_{\hat{x}} v(\hat{x}),$$

whenever each coordinate x_j is covered by coefficients $\alpha = (\alpha_{\hat{x}})_{\hat{x}\in\mathcal{X}}$, that is $\sum_{\hat{x}:x_j=\hat{x}_j}\alpha_{\hat{x}} \geq 1$.

Theorem 4.3.3 (β-XOS $\equiv$ β-Fractionally Subadditive)**.** *A valuation is β-fractionally subadditive across mechanisms if and only if it is β-*XOS*.*

Proof. We will give the proof of the general version of the theorem for β-fractionally subadditive valuation. We first show that if the valuation is β-XOS *across mechanisms* then it is also *β-fractionally subadditive across mechanisms*.

Suppose that there exists a set of additive valuations $\mathcal{L}$ such that: $\max_{\ell\in\mathcal{L}}\sum_{j\in[m]} v_j^\ell(x_j) \leq v(x) \leq \beta \max_{\ell\in\mathcal{L}}\sum_{j\in[m]} v_j^\ell(x_j)$. Now consider an allocation x^* and a fractional cover $(a_x)_{x\in X}$ of x^*, i.e. for all $j \in [m]$, $\sum_{x:x_j=x_j^*} a_x \geq 1$.

Now we have:

$$\begin{aligned}
\sum_{x\in\mathcal{X}} a_x v(x) \geq & \sum_{x\in\mathcal{X}} a_x \max_{\ell\in\mathcal{L}} \sum_{j\in[m]} v_j^\ell(x_j) \geq \max_{\ell\in\mathcal{L}} \sum_{x\in\mathcal{X}} a_x \sum_{j\in[m]} v_j^\ell(x_j) \\
= & \max_{\ell\in\mathcal{L}} \sum_{j\in[m]} \sum_{x\in\mathcal{X}} a_x v_j^\ell(x_j) \geq \max_{\ell\in\mathcal{L}} \sum_{j\in[m]} v_j^\ell(x_j^*) \sum_{x\in\mathcal{X}:x_j=x_j^*} a_x \\
\geq & \max_{\ell\in\mathcal{L}} \sum_{j\in[m]} v_j^\ell(x_j^*) \geq \frac{1}{\beta} v(x^*)
\end{aligned}$$

Thus the valuation is also β-fractionally subadditive.

Now we prove the opposite direction: if a valuation is *β-fractionally subadditive across mechanisms* then it is also *β-XOS across mechanisms*. Consider the following linear program associated with an outcome $x^* \in \mathcal{X}$:

$$\begin{aligned}
V(x^*) = & \min_{a_x} \sum_{x\in X} a_x v(x) \\
\text{s.t.} & \sum_{x:x_j=x_j^*} a_x \geq 1 \quad \text{for all } j\in[m] \\
& a_x \geq 0 \quad \text{for all } x\in\mathcal{X}
\end{aligned} \tag{4.7}$$

By the property of β-fractionally subaddtive valuations, since the set of feasible solutions to the above linear program, constitutes a fractional cover of x^* we know that $\beta V(x^*) \geq v(x^*)$. In addition we know that we can achieve $v(x^*)$ by just setting $a_{x^*} = 1$ and $a_x = 0$ for any other $x \in \mathcal{X}$. Hence, $V(x^*) \leq v(x^*) \leq \beta V(x^*)$.

Now consider the dual of the above linear program:

$$\begin{aligned}
C(x^*) = & \max_{t_j} \sum_{j\in[m]} t_j \\
\text{s.t.} & \sum_{j:x_j=x_j^*} t_j \leq v(x) \quad \text{for all } x\in\mathcal{X} \\
& t_j \geq 0 \quad \text{for all } j\in[m]
\end{aligned} \tag{4.8}$$

By LP duality we know that $V(x^*) = C(x^*)$. Let $t_j^{x^*}$ be an optimal solution to the dual associated with allocation x^*. Now consider the following additive valuation: $v_j^{x^*}(x_j) = t_j$ if $x_j = x_j^*$ and 0 otherwise. By the constraints of the dual we know that $v(x) \geq \sum_{j:x_j=x_j^*} t_j = \sum_{j\in[m]} v_j^{x^*}(x_j)$. Therefore:

$$v(x) \geq \max_{x^*\in X} \sum_{j\in[m]} v_j^{x^*}(x_j)$$

In addition by LP duality we know:

$$\max_{x^*\in\mathcal{X}} \sum_{j\in[m]} v_j^{x^*}(x_j) \geq \sum_{j\in[m]} v_j^{x}(x_j) = \sum_{j\in[m]} t_j^{x} = C(x) = V(x) \geq \frac{1}{\beta} v(x)$$

Therefore we get that:

$$\max_{x^*\in X} \sum_{j\in[m]} v_j^{x^*}(x_j) \leq v(x) \leq \beta \max_{x^*\in\mathcal{X}} \sum_{j\in[m]} v_j^{x^*}(x_j)$$

Hence, the valuation is also β-XOS. ■

REMARK 4.3.1. Observe that in the theorem above we used single-minded component valuations $v_j : \mathcal{X}_j \to \mathbb{R}_+$ of the form: $v_j^{\hat{x}}(x_j) = c$ if $x_j = \hat{x}_j$ and 0 otherwise. Thus this gives the corollary:

Corollary 4.3.4. *A valuation is β-fractionally subadditive across mechanisms if and only if it is β-XOS-$\mathcal{C}$, with $\mathcal{C}_j$ being the class of all single-minded valuations on $\mathcal{X}_j$.*

■

REMARK 4.3.2. Combining the main theorem of this section with the simultaneous composability Theorem 4.2.2, we get the following composability corollary for smooth mechanisms:

Corollary 4.3.5. *Consider the simultaneous composition of m (λ, μ)-smooth mechanisms, such that for each $j \in [m]$, $\mathcal{V}_i^j$ contains all single-minded valuations. Then the* *composition is a $\left(\frac{\lambda}{\beta}, \mu\right)$-smooth mechanism, when players have β-fractionally subadditive valuations across mechanisms.*

■

4.3.2 Hierarchy of Valuations across Mechanisms

In this section we define the classes of *set-submodular and set-subadditive valuations across mechanisms*, which are appropriate generalizations of submodular and subadditive set functions and we show the following generalized version of the corresponding hierarchy of set functions:

$$\text{set-submodular} \subset \text{XOS} \subset \text{set-subadditive} \subset H_m - \text{XOS} \tag{4.9}$$

where H_m is the m-th harmonic number.

To define generalizations of submodular and subadditive valuations, we will assume that each mechanism has a player-specific empty outcome $\perp_j \in \mathcal{X}_j$, which intuitively corresponds to: "the mechanism is not existent for player i". These outcomes don't affect the way the mechanism works (e.g. we don't impose that these outcomes be picked by the mechanism for some strategy profile) but it just serves as a reference point for the valuations of the bidders: we assume that $v(\perp_1, \ldots, \perp_m) = 0$. We will also use the notation x_S to denote the outcome vector that is x_j for all $j \in S$ and $\perp_j$ otherwise. We start with the generalization of subadditivity of set valuations:

Definition 4.3.6 (Set-Subadditive). *A valuation $v : \mathcal{X} \to \mathbb{R}_+$ is set-subadditive if and only if for any two sets $S_1, S_2 \subseteq [m]$ and any $x \in \mathcal{X}$:*

$$v(x_{S_1}) + v(x_{S_2}) \geq v(x_{S_1 \cup S_2})$$

In addition we define the notion of set-submodularity which extends submodularity of set valuations as follows: the marginal benefit from receiving an allocation x_j at some mechanism $\mathcal{M}_j$ decreases as the set of mechanisms from which the agent has received a non-empty allocation becomes larger.

Definition 4.3.7 (Set-Submodular)**.** *A valuation $v : \mathcal{X} \to \mathbb{R}_+$ is set-submodular if and only if, for any $x \in \mathcal{X}$, for any two sets $S \subseteq T \subseteq [m]$ and for any $j \in [m]$:*

$$v\left(x_{S+\{j\}}\right) - v\left(x_S\right) \geq v\left(x_{T+\{j\}}\right) - v\left(x_T\right)$$

Last, we will make the intuitive assumption that if a player wins a non-empty allocation in more mechanisms then his valuation cannot decrease: a valuation is *set-monotone* if for any two sets $S \subseteq T$: $v\left(x_S\right) \leq v\left(x_T\right)$. We show that the relation between these classes of valuations mirrors the relations of the analogous classes for set functions.

Theorem 4.3.8 (Set-Submodular $\subseteq$ XOS)**.** *If a valuation is set - monotone and set-submodular then it is* XOS*.*

Proof. We will prove that there exist a set of additive valuations $\mathcal{L}$ such that

$$\forall x : v(x) = \max_{\ell \in \mathcal{L}} \sum_{j \in [m]} v_j^\ell(x_j)$$

Each additive valuation in $\mathcal{L}$ will be associated with an outcome x. Denote with $M_j = \{j' \in [m] : j' \leq j\}$. The additive valuation associated with an outcome x will then be:

$$v_j^x(\tilde{x}_j) = \begin{cases} v(x_{M_j}) - v(x_{M_{j-1}}) & \text{if } \tilde{x} = x_j \\ 0 & \text{o.w.} \end{cases} \tag{4.10}$$

First observe that:

$$\sum_j v_j^x(x_j) = \sum_j v(x_{M_j}) - v(x_{M_{j-1}}) = v(x)$$

Next we will show that for all $\tilde{x}, x \in \mathcal{X} : v(\tilde{x}) \geq \sum_j v_j^x(\tilde{x}_j)$. The latter two facts together will establish that for all $\tilde{x} \in \mathcal{X} : v(\tilde{x}) = \max_x \sum_j v_j^x(\tilde{x}_j)$ which will complete theorem.

Fix two outcome vectors $\tilde{x}, x$. Let $S = \{j \in [m] : \tilde{x}_j = x_j\}$ and $S_j = \{j' \in S : j' \leq j\} = M_j \cap S$. By the set-monotonicity of the valuation we have:

$$v(\tilde{x}) = v(\tilde{x}_S, \tilde{x}_{-S}) = v(x_S, \tilde{x}_{-S}) \geq v(x_S)$$

Now observe that:

$$v(\tilde{x}) \geq v(x_S) = \sum_{j \in S} v(x_{S_j}) - v(x_{S_{j-1}})$$

Since for all $j \in S : S_j \subseteq M_j$, by set-submodularity we get that:

$$\begin{aligned} v(\tilde{x}) &\geq v(x_S) = \sum_{j \in S} v(x_{S_j}) - v(x_{S_{j-1}}) \\ &\geq \sum_{j \in S} v(x_{M_j}) - v(x_{M_{j-1}}) \\ &= \sum_{j \in S} v_j^x(\tilde{x}_j) = \sum_{j \in [m]} v_j^x(\tilde{x}_j) \end{aligned}$$

The latter equality follows from the fact that for all $j \notin S : v_j^x(\tilde{x}_j) = 0$, by definition. ■

Theorem 4.3.9 (Set-Subadditive $\subseteq$ H_m-XOS). *If a valuation is set-monotone and set-subadditive then it is H_m-XOS, where H_m is the m-th harmonic number.*

Proof. This proof is the generalization of the analogous proof for the case of valuations defined on sets, presented in [8].

We will show that there exists a set of additive valuations $\mathcal{L}$ such that:

$$\max_{\ell \in \mathcal{L}} \sum_j v_j^\ell(x_j) \leq v(x) \leq H_m \max_{\ell \in \mathcal{L}} \sum_j v_j^\ell(x_j)$$

Each additive valuation in $\mathcal{L}$ will be associated with an outcome $x \in \mathcal{X}$. We define the additive valuation associated with outcome x using the iterative process presented in Algorithm 2.

ALGORITHM 2: Procedure for computing the additive valuation associated with each outcome $x \in \mathcal{X}$.

Input : An outcome $x \in \mathcal{X}$ and a valuation $v : \mathcal{X} \to \mathbb{R}$
Output: Monotone valuations $v_j^x : \mathcal{X}_j \to \mathbb{R}$ for each $j \in [m]$

Set $C = \emptyset$;
while $C \neq [m]$ **do**
  Pick $A = \arg\min_{A' \subseteq [m]} \frac{v(x_{A'})}{|A'-C|}$;
  for *each* $j \in A - C$ **do**

$$v_j^x(\tilde{x}_j) = \begin{cases} \frac{v(x_A)}{|A-C|H_m} & \text{if } \tilde{x} = x_j \\ 0 & \text{o.w.} \end{cases} \tag{4.11}$$

  $C = C \cup A$;

First we argue that for each $x, \tilde{x} \in \mathcal{X}$: $v(\tilde{x}) \geq \sum_j v_j^x(\tilde{x}_j)$. Let $S = \{j \in [m] : \tilde{x}_j = x_j\}$. By the set-monotonicity of the valuation we have: $v(\tilde{x}) \geq v(x_S)$. Hence, it suffices to show that:

$$v(x_S) \geq \sum_j v_j^x(\tilde{x}_j) = \sum_{j \in S} v_j^x(\tilde{x}_j)$$

Consider the iteration t of Algorithm 2 at which the k-th element of S is added in C. Since at that iteration the algorithm chose A_t we have:

$$v_k^x(\tilde{x}_j) = \frac{v(x_{A_t})}{|A_t - C|H_m} \leq \frac{v(x_S)}{(|S| - k + 1)H_m}$$

Therefore:

$$\sum_{j \in S} v_j^x(\tilde{x}_j) = \sum_{k=1}^{|S|} v_k^x(\tilde{x}_j) \leq \frac{v(x_S)}{H_m} \sum_{k=1}^{|S|} \frac{1}{|S| - k + 1} = v(x_S)$$

Hence, for all $\tilde{x} \in \mathcal{X} : v(\tilde{x}) \geq \max_{x \in \mathcal{X}} \sum_j v_j^x(\tilde{x}_j)$.

Now we show that for all $\tilde{x} \in X : v(\tilde{x}) \leq H_m \max_{x \in \mathcal{X}} \sum_j v_j^x(\tilde{x}_j)$. Suppose that the algorithm takes T iterations to complete the computation and that

$A_1, \ldots, A_T$ are the sets picked at each iteration. Then:

$$\begin{aligned}\max_{x\in\mathcal{X}} \sum_j v_j^x(\tilde{x}_j) \geq & \sum_j v_j^{\tilde{x}}(\tilde{x}_j) = \sum_{t=1}^{T} \sum_{j\in A_t} v_j^{\tilde{x}}(\tilde{x}_j) \\ = & \sum_{t=1}^{T} \frac{v(\tilde{x}_{A_t})}{H_m} \geq \frac{v(\tilde{x}_{\cup_{t=1}^T A_t})}{H_m} = \frac{v(\tilde{x})}{H_m}\end{aligned}$$

■

4.3.3 Partially Ordered Allocation Spaces

In many settings, such as position auctions or combinatorial auctions or bandwidth allocation mechanisms, smoothness of a mechanism holds only when the valuations come from some restricted class of functions that do not contain single-minded valuations, i.e. only for monotone valuations in the position or only for concave valuations on allocated bandwidth.

For such applications, we want to understand the expressiveness of XOS-$\mathcal{C}$ valuations, when $\mathcal{C}$ doesn't contain single-minded valuations, but doesn't contain some general class of functions. In this section we will focus on the case where the valuation restriction $\mathcal{C}^j$ corresponding to each mechanism, contains only functions that are monotone with respect to some partial order on the allocation space $\mathcal{X}^j$.

EXAMPLE 4.3.1. (Position Auctions with Monotone Valuations) Suppose that each mechanism $\mathcal{M}_j$ is a position auction, i.e. it allocates a set of positions $\{1, \ldots, k\}$ to the players, i.e. $\mathcal{X}_i^j = \{1, \ldots, k\}$, such that no two players get the same position. Such mechanism design settings find application in online advertisement auction settings, where the positions correspond to advertisement slots presented together with the organic results of a search query. In such settings it makes sense to assume that a slot that appears higher in the web-page can only imply a higher value for the advertiser. In this case the allocation space is completely ordered and the valuation of a player is monotone non-decreasing in the position. As we will see in Chapter 11 some natural position auctions are smooth only under such a monotonicity assumption on the allowable valuations. ■

EXAMPLE 4.3.2. (Combinatorial Auctions with Monotone Valuations) Suppose that each mechanism $\mathcal{M}_j$ is a combinatorial auction, i.e. it partitions a set of items $\{1, \ldots, k\}$ to the players, i.e. $\mathcal{X}_i^j = 2^{[k]}$. In that case, the allocation space has a natural partial order defined by the subset relation, i.e. an allocation $S \in \mathcal{X}_i^j$ is smaller than $T \in \mathcal{X}_i^j$ if and only if $S \subseteq T$. For such settings, it is reasonable to assume that the local valuations of the players will satisfy *free-disposal*, i.e. more items cannot decrease my valuation. In fact, as we will see in Chapter 12 many mechanisms for combinatorial auctions are smooth only under such an assumption. ■

Here we will assume that the allocation space $\mathcal{X}_i^j$ of each mechanism admits some generic partial order $\succeq_j$ and we will show a stronger equivalence between fractionally subadditive valuations and XOS valuation, subject to monotonicity constraints with respect to this partial order.

Definition 4.3.10 (Monotone). *Consider a partial order $\succeq_j$ of each allocation space $\mathcal{X}_j$ which defines a poset $(\mathcal{X}_j, \succeq_j)$. The coordinate-wise partial order $\succeq$ of the product space $\mathcal{X} = \mathcal{X}_1 \times \ldots \times \mathcal{X}_m$ is the coordinate-wise ordering. A valuation $v : \mathcal{X} \to \mathbb{R}^+$ is monotone if and only if: $x \succeq \tilde{x} \implies v(x) \geq v(\tilde{x})$.*

Theorem 4.3.11. *A valuation $v : \mathcal{X} \to \mathbb{R}_+$ is monotone with respect to a coordinate-*

wise partial order $\succeq$ and β-fractionally subadditive if and only if it is β-XOS-$\mathcal{C}$, where $\mathcal{C}_j$ contains all valuations $v_j : \mathcal{X}_j \to \mathbb{R}_+$ that are monotone with respect to $\succeq_j$.

Proof. The *if* direction is easy to see, since as we showed in Theorem 4.3.3 if a function is β-XOS-$\mathcal{C}$ for any $\mathcal{C}$, then it is β-fractionally subadditive. Then monotonicity of the valuation follows from the monotonicity of the local valuations v_j^ℓ.

For the other direction, consider a valuation v that is *β-fractionally subadditive across mechanisms* and monotone with respect to the coordinate-wise ordering $(\succeq_j)_{j\in[m]}$. We will prove that it is also *β-XOS across mechanisms* and such that each induced valuation $v_j^\ell : \mathcal{X}_j \to \mathbb{R}^+$, used in the XOS representation, is monotone with respect to $\succeq_j$.

Consider the following variation of the linear program (4.7) used in the proof of Theorem 4.3.3 associated with an outcome $x^* \in \mathcal{X}$:

$$\begin{aligned} V(x^*) = \; & \min_{a_x} \sum_{x\in\mathcal{X}} a_x v(x) \\ & \text{s.t.} \sum_{x: x_j \succeq_j x_j^*} a_x \geq 1 \quad \text{for all } j \in [m] \\ & \qquad a_x \geq 0 \quad \text{for all } x \in \mathcal{X} \end{aligned}$$

Observe that in this variation the first set of constraints is altered to include a summation over outcomes greater than or equal to x_j^* and not only on outcomes equal to x_j^* as in LP (4.3.3).

Consider a feasible solution a_x to the above linear program. For $x \in \mathcal{X}$, let $S(x) = \{j \in [m] : x_j \succeq x_j^*\}$. By the monotonicity of the valuation we know that:

$$\sum_{x\in\mathcal{X}} a_x v(x) \geq \sum_{x\in\mathcal{X}} a_x v(x_{S(x)}^*, x_{-S(x)}) = \sum_{x\in\mathcal{X}} \tilde{a}_x v(x)$$

where, it is easy to see that:

$$\sum_{x:x_j=x_j^*} \tilde{a}_x = \sum_{x:x_j \succeq_{ij} x_j^*} a_x \geq 1$$

where the last inequality follows from the constraints of the linear program.

Therefore $\tilde{a}_x$ is a fractional cover of x^*. Hence, by the property of β-fractionally subaddtive valuations we know that

$$\beta \sum_{x \in \mathcal{X}} a_x v(x^*_{S(x)}, x_{-S(x)}) \geq v(x^*)$$

Since, this holds for any feasible a_x we get that $\beta V(x^*) \geq v(x^*)$. In addition we know that we can achieve $v(x^*)$ by just setting $a_{x^*} = 1$ and $a_x = 0$ for any other $x \in \mathcal{X}$. Hence, $V(x^*) \leq v(x^*) \leq \beta V(x^*)$.

Now consider the dual of the above linear program:

$$\begin{aligned} C(x^*) = \ & \max_{t_j} \sum_{j \in [m]} t_j \\ & \text{s.t.} \sum_{j:x_j \succeq_j x_j^*} t_j \leq v(x) \quad \text{for all } x \in \mathcal{X} \\ & \qquad t_j \geq 0 \quad \text{for all } j \in [m] \end{aligned}$$

By LP duality we know that $V(x^*) = C(x^*)$. Let $t_j^{x^*}$ be an optimal solution to the dual associated with allocation x^*.

Now consider the following induced valuations: $v_j^{x^*}(x_j) = t_j$ if $x_j \succeq_j x_j^*$ and 0 otherwise. By the constraints of the dual we know that $v(x) \geq \sum_{j:x_j \succeq_j x_j^*} t_j = \sum_{j \in [m]} v_j^{x^*}(x_j)$. Therefore:

$$v(x) \geq \max_{x^* \in X} \sum_{j \in [m]} v_j^{x^*}(x_j)$$

In addition by LP duality we know:

$$\max_{x^*\in\mathcal{X}}\sum_{j\in[m]} v_j^{x^*}(x_j) \geq \sum_{j\in[m]} v_j^{x}(x_j) = \sum_{j\in[m]} t_j^x = C(x) = V(x) \geq \frac{1}{\beta}v(x)$$

Therefore we get that:

$$\max_{x^*\in X}\sum_{j\in[m]} v_j^{x^*}(x_j) \leq v(x) \leq \beta \max_{x^*\in X}\sum_{j\in[m]} v_j^{x^*}(x_j)$$

Hence, the valuation is also β-XOS.

To complete the theorem we just need to prove that the valuations $v_j^x(\cdot)$ are monotone under the partial order $\succeq_j$, for each $x \in \mathcal{X}$ and for each $j \in [m]$. Observe that $v_j^x(\cdot)$ takes the same value (namely t_j^x) for all $\tilde{x}_j \succeq_j x_j$ and is 0 for any other $\tilde{x}_j$.

Consider two outcomes $\tilde{x}_j, \hat{x}_j \in \mathcal{X}_j$, such that $\tilde{x}_j \succeq_j \hat{x}_j$. If $\hat{x}_j \succeq_j x_j$ then by transitivity of $\succeq_j$ we also have that $\tilde{x}_j \succeq_j x_j$ and therefore $v_j^x(\tilde{x}_j) = v_j^x(\hat{x}_j)$. Otherwise, by definition we have $v_j^x(\hat{x}_j) = 0$ and therefore trivially $v_j^x(\tilde{x}_j) \geq v_j^x(\hat{x}_j)$. ■

REMARK 4.3.3. In particular in the proof of Theorem 4.3.11 we show that every monotone β-fractionally subadditive valuation can be expressed using induced valuations that are *step valuations*: $v^{j,\ell}(x_j) = c$ if $x_j \succeq_j \hat{x}_j$ and 0 otherwise.

Corollary 4.3.12. *A valuation $v : \mathcal{X} \to \mathbb{R}_+$ is monotone with respect to a coordinate-wise partial order $\succeq$ and β-fractionally subadditive if and only if it is β-XOS-$\mathcal{C}$, where $\mathcal{C}_j$ contains all step valuations with respect to $\succeq_j$.*

In the case where the allocation space $\mathcal{X}_j$ is the power set of a set of items, i.e. a combinatorial auction setting, then a step valuation function corresponds to a *single-minded bidder*, where the value function $f : 2^{[m]} \to \mathbb{R}_+$ takes the form:

$f(S) = c$ if $S \supseteq T$ and 0 otherwise and T is referred to as the players interest set.

■

REMARK 4.3.4. Combining the main theorem of this section with the simultaneous composability Theorem 4.2.2, we get the following composability corollary for smooth mechanisms:

Corollary 4.3.13. *Consider the simultaneous composition of m (λ, μ)-smooth mechanisms, such that for each $j \in [m]$, $\mathcal{V}_i^j$ contains all step valuations with respect to some partial order $\succeq_i^j$. Then the composition is a $\left(\frac{\lambda}{\beta}, \mu\right)$-smooth mechanism, when players have β-fractionally subadditive valuations across mechanisms that are monotone with respect to the coordinate-wise order $\succeq_i$.*

■

4.3.4 Lattice Allocation Spaces

If each partially ordered set $(\mathcal{X}_j, \succeq_j)$ forms a lattice[2] then it is natural to consider valuations that have *diminishing marginal returns* over this lattice: i.e. for any $z \succeq y$ and $t \in \mathcal{X}$

$$v(t \vee y) - v(y) \geq v(t \vee z) - v(z)$$

If the lattice is distributive and the valuation is monotone then the above class of valuations is equivalent to the class of *submodular valuations over the lattice* as we show below.

[2]A partially ordered space $\mathcal{X}$ forms a lattice if any two elements $x, x' \in \mathcal{X}$ have is a least upper bound $y = x \vee x'$, referred to as the join or supremum (i.e. $y \succeq x$ and $y \succeq x'$ and for any y' satisfying the latter conditions, $y' \succeq y$ and any two elements have a greatest lower bound $y = x \wedge x'$, referred to as the meet or the infimum (i.e. $x \succeq y$ and $x' \succeq y$ and for any y' satisfying this inequalities, $y \succeq y'$).

Definition 4.3.14 (Lattice-Submodular)**.** *If each poset $(\mathcal{X}_j, \succeq_j)$ forms a lattice then a valuation is lattice-submodular if and only if it is submodular on the product lattice of outcomes:*

$$\forall x, \tilde{x} \in \mathcal{X} : v(x \vee \tilde{x}) + v(x \wedge \tilde{x}) \leq v(x) + v(\tilde{x}) \tag{4.12}$$

Lemma 4.3.15. *If a valuation satisfies the diminishing marginal property with respect to a lattice structure then it is also lattice-submodular. If the lattice is distributive and the valuation is monotone then the inverse also holds.*

Proof. Consider two outcomes $x, \tilde{x} \in \mathcal{X}$. Since $\tilde{x} \succeq x \wedge \tilde{x}$, by the diminishing marginal returns property we have:

$$\begin{aligned} v(x \vee \tilde{x}) - v(\tilde{x}) \leq \; & v(x \vee (x \wedge \tilde{x})) - v(x \wedge \tilde{x}) \\ = \; & v(x) - v(x \wedge \tilde{x}) \end{aligned}$$

By rearranging we get that:

$$v(x \vee \tilde{x}) + v(x \wedge \tilde{x}) \leq v(x) + v(\tilde{x})$$

Since, $x, \tilde{x}$ where arbitrary the latter holds for any pair and therefore the valuation is submodular across mechanisms.

Now consider a valuation that is monotone and submodular over outcomes and in addition the lattice $(\mathcal{X}, \succeq)$ is distributive. We will show that for any $z \succeq y$ and for any $t \in X$ the diminishing marginal property holds. Invoke the submodular property for $x = t \vee y$ and $\tilde{x} = z$:

$$v((t \vee y) \vee z) + v((t \vee y) \wedge z) \leq v(t \vee y) + v(z)$$

Since $z \succeq y$ we have $t \vee y \vee z = t \vee z$. In addition, by distributivity of the lattice: $(t \vee y) \wedge z = (t \wedge z) \vee (y \wedge z) = (t \wedge z) \vee y$. Thus:

$$v(t \vee z) + v((t \wedge z) \vee y) \leq v(t \vee y) + v(z)$$

Now by monotonicity of the valuation we know that

$$v((t \wedge z) \vee y) \geq v(y)$$

Thus we get:

$$v(t \vee z) + v(y) \leq v(t \vee z) + v((t \wedge z) \vee y) \leq v(t \vee y) + v(z)$$

By rearranging we get the diminishing marginal property:

$$v(t \vee z) - v(z) \leq v(t \vee y) + v(y)$$

■

Below we give two examples of lattice submodular valuations that find natural applications in mechanism design settings.

EXAMPLE 4.3.3. (Bandwidth Allocation with Concave Valuations) Suppose that each mechanism M_j is a bandwidth allocation mechanism, i.e. it's goal is to split a bandwidth capacity C across the players. Thus the allocation space of each player is any portion between 0 and C, i.e. $\mathcal{X}_i^j = [0, C]$. This space is completely ordered (and thereby distributive) lattice and the class of diminishing marginal valuations corresponds to any concave function $f : [0, C] \rightarrow \mathbb{R}_+$. In bandwidth allocation settings it is normally assumed that the valuation of a player is concave and mechanisms for the bandwidth allocation setting, such as Kelly's proportional bandwidth allocation mechanism are smooth only under the concavity assumption (see Chapter 13). ■

EXAMPLE 4.3.4. (Combinatorial Auctions with Submodular Valuations) Suppose that each mechanism is a combinatorial auction setting, described in Example 4.3.2. The combinatorial allocation space $\mathcal{X}_i^j = 2^{[k]}$ is a lattice and in fact it is usually referred to as the binary lattice, where the join of two allocations is their union and the meet is their intersection. This lattice is also distributive. In many situations it makes sense to assume that even locally within each mechanism the players have no-complements across items and that their valuations are locally submodular over the items of the mechanism. In fact several mechanisms, that could be used as local mechanisms, are smooth, for reasonable λ and μ, only for complement-free valuations (e.g. simultaneous first-price auctions). ■

Under such structural assumptions on the allocation space, we can show a stronger connection between lattice submodular valuations and XOS valuations.

Theorem 4.3.16. *If a valuation is monotone and satisfies the diminishing marginal returns property with respect to a distributive product lattice $(\mathcal{X}, \succeq)$ then it can be expressed as an XOS valuation using valuations $v_j^\ell : \mathcal{X}_j \to \mathbb{R}^+$ that are capped marginal valuations:*

$$v_j^\ell(x_j) = v(x_j \wedge \hat{x}_j, \hat{x}_{-j}) - v(\perp_j, \hat{x}_{-j})$$

(for some $\hat{x} \in \mathcal{X}$ associated with each ℓ) and satisfy the diminishing marginal returns property with respect to $(\mathcal{X}_j, \succeq_j)$.

Proof. Suppose that each $(\mathcal{X}_j, \succeq_j)$ forms a lattice and the valuation $v : \mathcal{X} \to \mathbb{R}_+$ is monotone and satisfies the diminishing marginal returns property with respect to the product lattice. We will modify the definition of $v_j^x(\tilde{x}_j)$ used in Theorem 4.3.8 as follows:

$$v_j^x(\tilde{x}_j) = v(\tilde{x}_j \wedge x_j, x_{M_{j-1}}, \perp_{-M_j}) - v(x_{M_{j-1}}) \quad (4.13)$$

We first show that the set of additive valuations satisfy the XOS definition. Fix two outcomes $x, \tilde{x} \in \mathcal{X}$ and let $\hat{x} = \tilde{x} \wedge x$. By the monotonicity of the valuations we have that:

$$v(\tilde{x}) \geq v(\hat{x}) \tag{4.14}$$

Now by the diminishing marginal returns property of the function over the product lattice and the fact that for all j, $\hat{x}_j \succeq_j x_j$ we have:

$$\begin{aligned} v(\tilde{x}) \geq\ & v(\hat{x}) = \sum_j v(\hat{x}_j, \hat{x}_{M_{j-1}}, \perp_{-M_j}) - v(\hat{x}_{M_{j-1}}) \\ \geq\ & \sum_j v(\hat{x}_j, x_{M_{j-1}}, \perp_{-M_j}) - v(x_{M_{j-1}}) = \sum_j v_j^x(\tilde{x}_j) \end{aligned}$$

Now we show that each $v_j^x(\cdot)$ satisfies the diminishing marginal returns with respect to the lattice $(\mathcal{X}_j \succeq_j)$. Observe that the negative part in the definition of $v_j^x(\tilde{x}_j)$ is independent of $\tilde{x}_j$. Thus it suffices to show that the first part satisfies the diminishing marginal returns. For that it suffices to show that the following function:

$$v_j(\tilde{x}_j) = v(\tilde{x}_j \wedge x_j, x_{-j}) \tag{4.15}$$

satisfies the diminishing marginal returns as a function of $\tilde{x}_j$ for any $x \in \mathcal{X}$, whenever $v(\cdot)$ satisfies the diminishing marginal returns with respect to the product lattice. Since, we assumed that the valuation is monotone and the lattices are distributive we will equivalently show that $v_j(\cdot)$ is lattice-submodular whenever $v(\cdot)$ is:

$$v_j(y_j \wedge z_j) + v_j(y_j \vee z_j) =$$

$$v(y_j \wedge z_j \wedge x_j, x_{-j}) + v((y_j \vee z_j) \wedge x_j, x_{-j}) =$$

$$v((y_j \wedge x_j) \wedge (z_j \wedge x_j), x_{-j}) + v((y_j \wedge x_j) \vee (z_j \wedge x_j), x_{-j}) \geq$$

$$v(y_j \wedge x_j, x_{-j}) + v(z_j \wedge x_j, x_{-j}) =$$

$$v_j(y_j) + v_j(z_j)$$

Where the second equality follows from the distributivity of the lattice and the inequality follows from the submodularity of $v(\cdot)$. ■

REMARK 4.3.5. Combining the main theorem of this section with the simultaneous composability Theorem 4.2.2, we get the following composability corollary for smooth mechanisms:

Corollary 4.3.17. *Consider the simultaneous composition of m (λ, μ)-smooth mechanisms, such that each for each $j \in [m]$, $\mathcal{X}_i^j$ forms a lattice. If each player i has a monotone submodular valuation $v_i : \mathcal{X}_i \to \mathbb{R}_+$ on the product lattice and for each $j \in [m]$ capped marginals of v_i belong to $\mathcal{V}_i^j$, then the composition is a $\left(\frac{\lambda}{\beta}, \mu\right)$-smooth mechanism.*

■

4.4 Example: Simultaneous First Price Auctions

One important special case of the results in this chapter is when each individual mechanism is a single-item auction. Several papers, both in the economics and computer science literature, have analyzed the efficiency of the game defined by simultaneous single-item auctions.

Brief overview of literature. In economics, Engelbrecht-Wiggans and Weber [22], were the first to analyze a game of simultaneous second-price auctions with unit-demand bidders in the complete information setting, where every bidder has a value of one for getting any of the items. They showed that if each bidder is restricted to bid in only one of the auctions, then there exists a symmetric mixed Nash equilibrium where the social welfare is only a $1 - \frac{1}{e}$ fraction of the optimal social welfare, i.e. the price of anarchy of mixed Nash equilibria is at least $\frac{e}{e-1}$. More recently, Bikhchandani [9], analyzed simultaneous first price auctions with more general valuations and where players can bid on more than one auction and showed that in the complete information setting, every pure Nash equilibrium (if it exists) must be fully efficient and correspond to a Walrasian equilibrium.

In computer science, Christodoulou et al. [15] first analyzed the efficiency of Bayes-Nash equilibria of simultaneous second-price auctions. They showed that when players have XOS valuations drawn independently from aribtrary distributions, then every Bayes-Nash equilibrium (assuming that player's don't bid above their valuations) has expected welfare at least half of the expected optimal welfare, i.e. the Bayes-Nash price of anarchy is at most 2. The result was later extended to $2 \cdot \beta$ for β-XOS valuations by [8], implying a $O(\log(n))$ bound for subadditive valuations. The result was improved to 4 by Feldman et al. [25]. For first-price auctions Hassidim et al. [36] gave a bound of 4 on the Bayes-Nash price of anarchy for the case of XOS valuations and Feldman et al. [25] a bound of 2 for subadditive valuations.

Smoothness analysis. Following our running Example 3.1.1, in this section we will analyze the case when each auction is a first price auction (c.f. Chapter 10 for other auction formats).

From Lemma 3.1.3 we know that the first price auction is a $\left(1 - \frac{1}{e}, 1\right)$-smooth mechanism. Thus theorem 4.2.2 directly implies that the mechanism defined by running m first price auctions simultaneously is also a $\left(1 - \frac{1}{e}, 1\right)$-smooth mechanism, when players have XOS (or equivalently fractionally subadditive) valuations over the items. Thus the price of anarchy of the game is at most $\frac{e}{e-1}$ and this holds even at no-regret learning outcomes and even under incomplete information, i.e. at every Bayes coarse correlated equilibrium.

Composability proof unraveled. To demystify the above result, we break apart the composability proof for the special case of simultaneous first price auctions with unit-demand bidders. We describe the smoothness deviations that the composability proof of Theorem 4.2.2 constructs and which are the *certificates* that the game is approximately efficient.

More formally, in this expository example, we will consider the case where each player $i \in [n]$ has a value w_{ij} for item $j \in [m]$. Each player is unit-demand, i.e. wants only one item, and if he wins more than one item then his valuation is the highest value item he won:

$$v_i(S) = \max_{j \in S} w_{ij} \tag{4.16}$$

This is a special case of an XOS valuation: for each item $j^* \in [m]$, there is one additively separable valuation $v_i^{j^*}$ in the index set $\mathcal{L}$, and $v_i^{j,j^*}(x_j) = w_{ij^*} \cdot \mathbf{1}\{j = j^*\} \cdot x_j$, with $x_j \in \{0, 1\}$. Then observe that $v_i(x) = \max_{j^* \in \mathcal{L}} \sum_{j \in [m]} v_i^{j,j^*}(x_j)$.

The composability proof constructs the smoothness deviation for the global mechanism as follows: for each valuation profile v, consider the optimal allocation $x^*(v)$. In the case of unit-demand bidders, the optimal allocation will be a matching. Thus each player is allocated an item $j^*(i)$. Then consider the additively separable valuation that matches a player's value for his optimal allocation. In this case it will trivially be $v_i^{j^*(i)}$. Now consider the valuation profile where each player's true valuation is replaced with the latter additively separable valuation. Denote this profile $v^* = (v_1^*, \ldots, v_n^*)$.

Observe that under this profile, at each item $j \in [m]$, only one player i has non-zero value, and this is exactly the player that is matched to that item in the optimal matching allocation. Moreover, his value is w_{ij}. Therefore, he is locally the optimal player for that item under this new valuation profile.

The global deviation asks from each player i, to submit his local smoothness deviation at each item $j \in [m]$, under valuation profile $v^{j,*} = (v_1^{j,j^*(1)}, \ldots, v_n^{j,j^*(n)})$. This means submitting a zero bid at every item $j \neq j^*(i)$ and submitting a random bid with density function $f(x) = \frac{1}{w_{i,j^*(i)} - x}$ and support $[0, (1 - 1/e)w_{i,j^*(i)}]$ at item $j^*(i)$.

Observe that by local smoothness of each first price auction we have that the utility of each player i from his optimal item $j^*(i)$ under this deviation is at least: $\left(1 - \frac{1}{e}\right) w_{i,j^*(i)} - \mathcal{R}^{\mathcal{M}_{j^*(i)}}(b^{j^*(i)})$. The global smoothness then follows by summing over all players and observing that $\sum_{i \in [n]} w_{i,j^*(i)} = \textsc{Opt}(v)$ and $\sum_i \mathcal{R}^{\mathcal{M}_{j^*(i)}}(b^{j^*(i)}) = \mathcal{R}^{\mathcal{M}}(b)$, since $j^*(\cdot)$ is a matching.

4.5 Composability under Restricted Complements

In many scenarios, the value of a player might exhibit some limited complementarities across allocations from different mechanisms. For instance, if each mechanism is a combinatorial auction then maybe two items (left and right shoe) that have value for the player only when acquired in conjunction, might be sold by two different mechanisms. Our complement-free assumption that is implicit in the definition of XOS valuations does not allow for such complementary relations.

However, in many practical scenarios, such as spectrum auctions, bidders valuations do exhibit complements, albeit restricted ones. Hence, it is natural to ask how does the efficiency of a market composed of smooth mechanisms degrades in the presence of complements.

We introduce a measure of the size of a complement and show an approximate composability theorem that will yield efficiency results even in the presence of complements. Intuitively, if the measure of complementarity is k, then it means that the type of "conjunctive" complementary relations can occur only across k-tuples of mechanisms. Thus in the case of a left and a right shoe, the measure is two. For easier understanding of our complementarity measure, we first describe it in the context of simultaneous single-item auctions, where the value of a player is a set function. We then give the generalized definition of the measure for general mechanisms.

4.5.1 Maximum over Positive Hypergraph Set Functions

Given a set M of m items, consider a set function $v : 2^M \to \mathbb{R}_+$ that is *normalized*, i.e. $v(\emptyset) = 0$ and *monotone*, i.e. $v(T) \geq v(S)$ whenever $S \subseteq T \subseteq M$. A hypergraph representation of a set function $v : 2^M \to \mathbb{R}_+$ is a (normalized but not necessarily monotone) set function $h : 2^M \to \mathbb{R}$ that satisfies $v(S) = \sum_{T \subseteq S} h(T)$. It is easy to verify that any set function v admits a unique hypergraph representation and vice versa. A set S such that $h(S) \neq 0$ is referred to as a *hyperedge* of h. Pictorially, the hypergraph representation can be thought of as a weighted hypergraph, where every vertex is associated with an item in M, and the weight of each hyperedge $e \subseteq M$ is $h(e)$. Then the value of the function for any set $S \subseteq M$, is the total value of all hyperedges that are contained in S.

The *rank* of a hypergraph representation h is the largest cardinality of any hyperedge. Similarly, the *positive rank* (respectively, *negative rank*) of h is the largest cardinality of any hyperedge with strictly positive (respectively, negative) value. The rank of a set function v is the rank of its corresponding hypergraph representation, and we refer to a function v with rank r as a *hypergraph-r* function. Last, if the hypegraph representation is non-negative, i.e. for any $S \subseteq M$, $h(S) \geq 0$, then we refer to such a function as a *positive hypergraph-r* (PH-r) function .

We define a parameterized hierarchy of set functions, with a parameter that corresponds to the degree of complementarity.

Definition 4.5.1 (Maximum Over Positive Hypergraph-k (MPH-k) class)**.** *A monotone set function $v : 2^M \to \mathbb{R}_+$ is* Maximum over Positive Hypergraph-k *(*MPH-k*) if it can be expressed as a maximum over a set of* PH-k *functions. That is, there exist* PH-k *functions $\{v_\ell\}_{\ell \in \mathcal{L}}$ such that for every set $S \subseteq M$,*

$$v(S) = \max_{\ell \in \mathcal{L}} v_\ell(S), \tag{4.17}$$

where $\mathcal{L}$ is an arbitrary index set.

The MPH hierarchy is a complete one and thereby the lowest level of the hierarchy that a function belongs to is a valid measure of complementarity for any set function. The two extreme cases of MPH-k functions coincide with two important classes of valuations. Specifically, MPH-1 is the class of functions that can be expressed as the maximum over a set of additive functions. This is

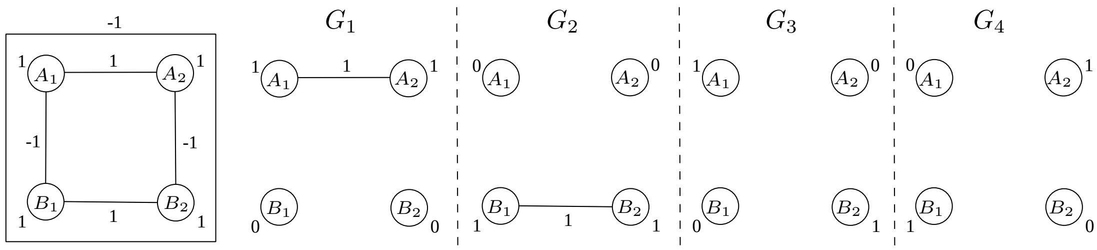

Figure 4.1: The left figure depicts a spectrum auction inspired hypergraph valuation with positive edges and negative hyperedges, which can be expressed as the maximum over the positive graphical valuations on the right.

exactly the class of XOS valuations [45] analyzed in the complement-free valuation section. On the other side, MPH-m coincides with the class of all monotone functions,[3] and so the hierarchy is complete. For intermediate values of k, MPH-k is monotone; namely, for every $k < k'$ it holds that MPH-$k \subset$ MPH-k'. We get the following hierarchy:

$$\text{Submodular} \subset XOS = \text{MPH-1} \subset \cdots \subset \text{MPH-}m = \text{Monotone} \qquad (4.18)$$

A simple example. Consider the example depicted in Figure 4.5.1, which has an intuitive interpretation in the context of FCC spectrum auctions. Suppose that A, B are two spectrum bands and that A_i, B_i are auctions representing band A or B at location i. Locations 1 and 2 are neighboring geographic regions and therefore, a bidder gets a much larger value for getting the same band in both regions. Therefore, A_1 and A_2 have a complementary relationship and similarly B_1 and B_2. However, each A_i has a substitute relationship with B_i and additionally the pair (A_1, A_2) has a substitute relationship with the pair (B_1, B_2), since a bidder will only utilize one pair of bands. This valuation can be represented

[3]Simply create a separate PH-$|S|$ function for each set S with a single hyperedge equal to the set S and with weight $f(S)$. Then, by monotonicity, the maximum of these functions is equal to the initial valuation.

as a hypergraph, as in the left-most diagram in Figure 4.5.1. Also, as illustrated in Figure 4.5.1, this valuation can be represented as a maximum over positive hypergraph valuations of rank 2.

Fractionally "Subadditive" Characterization of MPH-k. We show that the definition of MPH-k functions has a natural analogue as an extension of *fractionally subadditive* functions. More formally, consider a set S of items and let $S|_k$ be all the subsets of S of size at most k. We say that a collection of sets $\mathcal{T} \subseteq 2^S$ together with a weight a_T for each $T \in \mathcal{T}$ is a fractional cover of all the subsets of size at most k (k-fractional cover) of S if $\forall s \in S|_k : \sum_{T \in \mathcal{T}: T \supseteq s} a_T \geq 1$. A valuation $v : 2^M \to \mathbb{R}_+$ is k-fractionally subadditive if for every $S \subseteq M$ and every k-fractional cover $(a_{\mathcal{T}}, \mathcal{T})$ of S, we have $v(S) \leq \sum_{T \in \mathcal{T}} a_T \cdot v(T)$.

Theorem 4.5.2. *The class of monotone k-fractionally subadditive valuations is equivalent to the class of* MPH-k *valuations.*

Proof. First it is easy to observe that any MPH-k valuation is k-fractionally subadditive:

$$\begin{aligned}\sum_{T \in \mathcal{T}} a_T \cdot v(T) = & \sum_{T \in \mathcal{T}} a_T \cdot \max_{\ell \in \mathcal{L}} \sum_{s \in T|_k} w_s^\ell \geq \max_{\ell \in \mathcal{L}} \sum_{T \in \mathcal{T}} a_T \sum_{s \in T|_k} w_s^\ell \\ = & \max_{\ell \in \mathcal{L}} \sum_{s \in S|_k} w_s^\ell \sum_{T \in \mathcal{T}: T \supseteq s} a_T \geq \max_{\ell \in \mathcal{L}} \sum_{s \in S|_k} w_s^\ell = v(S)\end{aligned}$$

To show that any monotone k-fractionally subadditive valuation is an MPH-k valuation, we follow a similar analysis to that carried by Feige [23], as follows. For every set S, we construct a hypergraph-k valuation associated with the set S, and denote it by $\ell(S)$. The set of valuations is then $\mathcal{L} = \cup_{S \subseteq [m]} \ell(S)$. The hypergraph valuation $\ell(S)$ is constructed such that: (i) $v(S) = \sum_{s \in S|_k} w_s^{\ell(S)}$, and

(ii) for any subset $T \subseteq S : v(T) \geq \sum_{s \in T|_k} w_s^{\ell(S)}$. Monotonicity then implies that for any set S, $v(S) = \max_{\ell \in \mathcal{L}} \sum_{s \in S|_k} w_s^\ell$, as desired.

It remains to construct the valuation $\ell(S)$. To this end, we consider the following linear program and its dual:

$$
\begin{aligned}
V(S) = & \min_{(a_T)_{T \subseteq S}} \sum_{T \subseteq S} a_T \cdot v(T) & \qquad C(S) = & \max_{(w_s)_{s \in S|_k}} \sum_{s \in S|_k} w_s \\
\forall s \in S|_k : & \sum_{T \supseteq s} a_T \geq 1 & \qquad \forall T \subseteq S : & \sum_{s \in T|_k} w_s \leq v(T) \\
\forall T \subseteq S : & \ a_T \geq 0 & \qquad \forall s \in S|_k : & \ w_s \geq 0
\end{aligned}
$$

By definition, every feasible solution to the primal program constitutes a fractional cover of every subset of size at most k of S. Therefore, it follows by k-fractional subadditivity that $V(S) \geq v(S)$. Since $v(S)$ can be obtained by setting $a_S = 1$ and $a_T = 0$ for any $T \subset S$, we get that $V(S) = v(S)$. Duality then implies that $C(S) = v(S)$. Thus if we set $(w_s^{\ell(S)})_{s \in S|_k}$ to be the solution to the dual, then the conditions that need to be hold for $\ell(S)$ are satisfied by the constraints of the dual and the duality. ■

4.5.2 Restricted Complements across Mechanisms

To present our composability theorem in the general mechanism composition setting, we first generalize the class of MPH-k set functions to valuations across mechanisms.

Definition 4.5.3 (Positive Hypergraph-k across mechanisms). *A valuation $v : \mathcal{X}_1 \times \ldots \times \mathcal{X}_m \to \mathbb{R}_+$ is positive hypergraph-k across mechanisms if for any $x \in$*

$\mathcal{X}_1 \times \ldots \times \mathcal{X}_m$

$$v(x) = \sum_{e \in E} v^e(x^e) \tag{4.19}$$

where $E \subseteq \{S \subseteq M : |S| \leq k\}$, $x^e = (x^j)_{j \in e}$ *is the vector of allocations on the mechanisms in the set* e *and for all* $e \in E$, $v^e(x^e) \geq 0$.

The latter class is the generalization of positive hypergraph-k set functions (see Abraham et al [1]). The above class of valuations allows exactly k-wise complementary relations across mechanisms. However, it does not allow for arbitrary substitute relations, i.e. it does not even include XOS valuations across mechanisms. To achieve this we define the more general class of maximum over positive hypergraph-k valuations across mechanisms.

Definition 4.5.4 (Maximum over Positive Hypergraph-k across mechanisms)**.** *A valuation* $v : \mathcal{X}_1 \times \ldots \times \mathcal{X}_m \to \mathbb{R}_+$ *is* MPH-k *across mechanisms if there exists a set* $\mathcal{L}$ *of positive hypregraph-*k *valuations, such that for any* $x \in \mathcal{X}_1 \times \ldots \times \mathcal{X}_m$

$$v(x) = \sup_{\ell \in \mathcal{L}} v^{\ell}(x) \tag{4.20}$$

We conclude the section with an example of a valuation with restricted complements across mechanisms, which is not a set function example. The example is motivated by "impression effects" that can arise in online advertisement auctions.

EXAMPLE 4.5.1. (Position Auctions with Impression Effects) For many search queries in Google or Bing, there are more than one group of advertisement slots that are auctioned to advertisers. Typically, there will be a small set of top slots that are presented above the organic search results and another group of side slots that are presented to the right of the organic search results. Though the exact mechanism that takes place to allocate these slots is rather complicated, one can approximate it by the abstraction that each group of slots is auctioned via a separate position auction, such as the generalized second price auction (i.e. advertisers submit a separate bid for the top slots and a separate bid for the side slots). Thus we can view it as a composition of two position mechanisms.

In such a setting an advertiser might have an extra "impression" value for winning an the top slot in both groups, as this will create an impression effect to the web user. Thus the value of the advertiser if he is allocated slot j among the top slots and slot j' among the side slots could look like:

$$v_i(j, j') = a_j \cdot w_i^{\mathrm{c}} + \tilde{a}_{j'} \cdot w_i^{\mathrm{c}} + w_i^{\mathrm{im}} \cdot \mathbf{1}\{j = j' = 1\},$$

where a_j is the click probability of top slot j, $\tilde{a}_{j'}$ is the click probability of side slot j', w_i^{c} is the per-click value of the advertiser and w_i^{im} is the value for the impression effect.

Observe that this valuation is an MPH-2 valuation across mechanisms. Moreover, we can take an even more global view and consider all the position auctions that the player participates in the platform (e.g. other ad campaigns or other keywords). Assuming that across different impressions the value of the advertiser is monotone and fractionally subadditive (i.e. view the two position auctions of each impression as a single mechanism and then consider the composition of the impression mechanisms, then the player's value is monotone and fractionally subadditive across mechanisms), then the whole valuation of the advertiser still remains MPH-2 as the complementarities only appear across the two auctions for the same impression. ■

4.5.3 Composability Theorem with Complements

To prove composability under restricted complements we will need to assume that the allocation space of each mechanism is partially ordered and that the valuation space for which smoothness holds includes all step valuations with respect to the partial order.

Then we show that if the valuation v_i of each player is MPH-k and the value functions $v_i^{e,\ell}(x_i^e)$ used to express his valuation are monotone coordinate-wise with respect to this partial order, then local (λ, μ)-smoothness of each mechanism, implies global $(1 - k + \min\{\lambda, 1\} \cdot k, \mu)$-smoothness.

Theorem 4.5.5. *Consider the simultaneous composition of m mechanisms each being (λ, μ)-smooth for any step valuation with respect to some partial order of the allocation space. If players have* MPH-k *valuations across mechanisms such that $v_i^{e,\ell}(\cdot)$ are monotone coordinate-wise with respect to each partial order, then the composition is $(1 - k + \min\{\lambda, 1\} \cdot k, \mu)$-smooth*

Proof. By Lemma 4.2.3 it suffices to show smoothness of the global mechanism only for positive hypergraph-k valuation, i.e. for each player i we have:

$$v_i(x) = \sum_{e \in E_i} v_i^e(x_i^e), \tag{4.21}$$

where $E_i \subseteq \{S \subseteq M : |S| \leq k\}$ and $x_i^e = (x_i^j)_{j \in e}$. However, as compared to the proof of Theorem 4.2.2, smoothness of the global mechanism for positive hypergraph-k valuations is not as straightforward as the case of additively separable valuations, since the utility of a player doesn't decompose. Hence, in the remainder of the proof we prove the desired smoothness. To achieve this we need to construct the special randomized actions $\mathbf{a}_i^*(v)$ required by the smoothness property.

Let $\tilde{x}_i = (\tilde{x}_i^j)_{j\in[m]}$ be the optimal allocation of each player i. Consider an action profile $a = (a^j)_{j\in M}$ on each auction j and each player deviating to some strategy $\tilde{a}_i = \left(\tilde{a}_i^j\right)_{j\in[m]}$. Then we can obtain the following lower bound a player's utility from the deviation:

$$
\begin{aligned}
&u_i(\tilde{a}_i, a_{-i}) = \\
&\sum_{e\in E_i}\sum_{x_i^e} v_i^e(x_i^e)\cdot \Pr\left(X_i^e(\tilde{a}_i, a_{-i}) = x_i^e\right) - \sum_{j\in M} P_i^j(\tilde{a}_i^j, a_{-i}^j) \geq \\
&\sum_{e\in E_i} v_i^e(\tilde{x}_i^e)\cdot \Pr\left(X_i^e(\tilde{a}_i, a_{-i}) \succeq \tilde{x}_i^e\right) - \sum_{j\in M} P_i^j(\tilde{a}_i^j, a_{-i}^j) \geq \\
&\sum_{e\in E_i} v_i^e(\tilde{x}_i^e)\cdot \left(1 - \sum_{j\in e}\left(1 - \Pr\left(X_i^j(\tilde{a}_i^j, a_{-i}^j) \succeq \tilde{x}_i^j\right)\right)\right) - \sum_{j\in M} P_i^j(\tilde{a}_i^j, a_{-i}^j) = \\
&\sum_{e\in E_i} v_i^e(\tilde{x}_i^e)\cdot(1-|e|) + \sum_{e\in E_i} v_i^e(\tilde{x}_i^e)\sum_{j\in e}\Pr\left(X_i^j(\tilde{a}_i^j, a_{-i}^j) \succeq \tilde{x}_i^j\right) - \sum_{j\in M} P_i^j(\tilde{a}_i^j, a_{-i}^j) = \\
&\sum_{e\in E_i} v_i^e(\tilde{x}_i^e)\cdot(1-|e|) + \sum_{j\in M}\left\{\left(\sum_{e\ni j} v_i^e(\tilde{x}_i^e)\right)\cdot \Pr\left(X_i^j(\tilde{a}_i^j, a_{-i}^j) \succeq \tilde{x}_i^j\right) - P_i^j(\tilde{a}_i^j, a_{-i}^j)\right\}.
\end{aligned}
$$

Summing up over all players we observe that the second summand in the above expression will correspond to deviating utilities of individual single-item auctions, where each player unilaterally deviates to $\tilde{a}_i^j$ and in which every player has a valuation of $\sum_{e\ni j} v_i^e(\tilde{x}_i^e)$ for getting any allocation $x_i^j \succeq \tilde{x}_i^j$ and 0 otherwise. The latter is a step valuation and hence we can set the local deviating actions at each mechanism to the smoothness deviations for the latter step valuation profiles and get:

$$
\begin{aligned}
\sum_i u_i(\tilde{a}_i, a_{-i}) \geq\ & \sum_i\sum_{e\in E_i} v_i^e(\tilde{x}_i^e)(1-|e|) + \lambda\sum_i\sum_{e\in E_i} v_i^e(\tilde{x}_i^e)\cdot|e| - \mu\mathcal{R}^{\mathcal{M}}(a) \\
=\ & \sum_i\sum_{e\in E_i} v_i^e(\tilde{x}_i^e)\cdot(1-(1-\lambda)|e|) - \mu\mathcal{R}^{\mathcal{M}}(a)
\end{aligned}
$$

If $\lambda < 1$ then we use the fact that $|e| \leq k$ to get the $(1-k+\lambda k, \mu)$-smoothness property, otherwise we can simply ignore the term $(1-\lambda)|e|$ and get the $(1,\mu)$-smoothness property ■

REMARK 4.5.1. We need to point out that the above theorem would lead to meaningful price of anarchy bound only when $\lambda > 1 - \frac{1}{k}$. Hence, each individual mechanism must be (λ, μ)-smooth for values of λ arbitrarily close to 1. For instance, the $(1 - 1/e, 1)$-smoothness property that we gave in Example 3.1.1 for the first price auction is not sufficient to yield any reasonable efficiency guarantee under complements.

However, for many mechanisms it is possible to show that the mechanism is (λ, μ)-smooth for values of λ arbitrarily close to 1, at the expense of increasing μ. This is, for instance, the case for the first price auction where as we will see in the next section it is $(\beta(1 - e^{-1/\beta}), \beta)$-smooth for any $\beta \geq 1$. For appropriate value of β, this would then lead to an $O(k)$ price of anarchy bound for simultaneous auctions under k-wise complements. Similar approach will apply for other auctions such as position auctions or greedy combinatorial auctions as we will see in Part III. ■

4.6 Example: Simultaneous First Price Auctions with Complements

As an example, we first analyze the efficiency of simultaneous first price auctions when players have MPH-k valuations over the items. In Section 4.4 we showed that when players have XOS $\equiv$ MPH-1 valuations then the price of anarchy is at most $\frac{e}{e-1}$. Here, we extend this result to show that for MPH-k valuations the price of anarchy is $O(k)$.

We first show a stronger smoothness property of the first price auction and then we will apply Theorem 4.5.5, to get an $O(k)$ price of anarchy for maximum over positive hypergraph-k valuations.

Lemma 4.6.1. *The first price auction is a* $\left(\beta \cdot \left(1 - e^{-1/\beta}\right), \beta\right)$*-smooth mechanism for any* $\beta > 0$.

Proof. Consider valuation profile $v = (v_1, \ldots, v_n)$. The highest value player (wlog player 1) can deviate to submitting a randomized bid $\mathbf{b}_1^*$ drawn from a distribution with density function $f(x) = \frac{\beta}{v_1 - x}$ and support $[0, (1 - e^{-1/\beta})v_1]$, while all non-highest value players should just deviate to bidding 0. No matter what the rest of the players are bidding, the utility of the highest bidder from the deviation is:

$$\begin{aligned} U_1^{\text{FPA}}(\mathbf{b}_1^*, b_{-1}; v_1) \geq & \int_{\max_{i \neq 1} b_i}^{\left(1-e^{-1/\beta}\right)v_1} (v_1 - x) f(x) dx \geq \beta \left(1 - e^{-1/\beta}\right) v_1 - \beta \max_i b_i \\ = & \ \beta \left(1 - e^{-1/\beta}\right) \textsc{Opt}(v) - \beta \sum_{i \in [n]} P_i(b) \end{aligned}$$

■

By applying Theorem 4.5.5, we get that the simultaneous first price auction with MPH-k valuations is $(1 - (1 - \beta \cdot (1 - e^{-1/\beta})) \cdot k, \beta)$-smooth for any $\beta > 0$. For $\beta = \frac{1}{\log(\frac{k}{k-1})}$ we get the optimal price of anarchy bound of $\frac{1}{1-(k-1)\log(\frac{k}{k-1})} \leq k(2 - e^{-k})$.

4.7 Example: Simultaneous Position Auctions

We now consider the setting of simultaneous position auctions with restricted complement valuations. One such instance, was given in Example 4.5.1, where we allowed for the players to exhibit complementarities across the two auctions that happen for the same impression and that corresponded to an impression effect. Here we will assume that each position auction happens via the means of a first-price pay-per-impression auction: each player submits a bid b_i at each position auction and the players are allocated positions in decreasing order of their bids. If a player is allocated a slot then he pays his bid. We will examine the efficiency of m simultaneous such auctions, when players have MPH-k valuations across mechanisms. In Chapter 11 we examine the efficiency of other position mechanisms and give a more extensive analysis.

We first analyze the smoothness of the first-price pay-per-impression auction.

Lemma 4.7.1. *The first-price pay-per-impression position auction is a* $(1-\frac{1}{2\beta},\beta)$*-smooth mechanism for any* $\beta \geq 1$.

Proof. Consider a bid profile b and let j_i^* be the optimal position of player i and let $\pi(j)$ be the player that gets slot j under bid profile b. Suppose that each player deviates to bidding a random bid b_i', uniformly in $[0, \frac{v_{ij_i^*}}{\beta}]$ and let $f(t)$ denote the density function of the random bid. If the random bid t of a player is $b_{\pi(j_i^*)} < t$ then player i wins his optimal slot or a higher slot and hence his value is at least $v_{ij_i^*}$ by monotonicity of the valuation.

Thus a player's utility from this deviation is at least:

$$\begin{aligned} u_i(b_i', b_{-i}) \geq & \int_{b_{\pi(j_i^*)}}^{\frac{v_{ij_i^*}}{\beta}} v_{ij_i^*} f(t) dt - \frac{v_{ij_i^*}}{2\beta} \\ = & \int_{b_{\pi(j_i^*)}}^{\frac{v_{ij_i^*}}{\beta}} \beta \cdot dt - \frac{v_{ij_i^*}}{2\beta} = \left(1 - \frac{1}{2\beta}\right) v_{ij_i^*} - \beta \cdot b_{\pi(j_i^*)} \end{aligned}$$

Summing over all players we get the $(1-\frac{1}{2\beta},\beta)$-smoothness property. ■

Applying Theorem 4.5.5, we get that for any MPH-k valuation across position auctions the simultaneous position auction mechanism is $(1 - \frac{k}{2\beta}, \beta)$-smooth, yielding a price of anarchy bound of $2k$ for $\beta = k$. Thus for instance, we get that for the valuations presented in Example 4.5.1 every equilibrium achieves at least $\frac{1}{4}$ of the optimal welfare.

5

SEQUENTIAL COMPOSABILITY

...With the next issuance of the 10 year DSL, on March 9th, the Dutch Sequential Auction will therefore be applied for the first time. The DSTA plans three auctionettes, according to the following time table:

1. *from 10.00 to 10.15am bids can be submitted for the first auctionette; results around 10.20am;*
2. *10.30 to 10.45am bids can be submitted for the second auctionette; results around 10.50am;*
3. *11.00 to 11.15am bids can be submitted for third and last auctionette; results around 11.20am.*

– Dutch State Treasure Agency, February 1999

In many real world scenarios, ranging from electronic markets like eBay to auctions for art, mechanisms take place sequentially rather than simultaneously. In this section we examine the efficiency of such sequential markets. Specifically, we analyze a model where many mechanisms, in the generic sense, take place one after the other.

The crucial difference with the simultaneous counterpart, analyzed in the previous chapter, is that in sequential settings, players have the ability to respond to deviations of their opponents. For this reason, the analysis of the previous section would break apart. Hence, we need new techniques to show that global efficiency results from local smoothness of mechanisms. In this chapter we will show that smooth mechanisms, still compose well sequentially, but for a more restricted class of valuations. As we will show this restriction is almost necessary.

Informal Theorem 3. *A market consisting of running m (λ, μ)-smooth mechanisms sequentially achieves welfare at equilibrium at least $\frac{\lambda}{1+\mu}$ of the optimal, if the value of a player is* unit-demand across mechanisms.

To prove our theorem we will first relax the smoothness condition to allow the deviating strategies to depend on a player's current action from which he is deviating. This relaxed version of smoothness will allow us to capture efficiency in sequential games where a good deviation might require to simulate a player's previous action until the "right moment" arrives.

Then we show that if a mechanism satisfies this relaxed notion of smoothness then it has similar robust efficiency guarantees as a smooth mechanism, with the exception that the guarantee extends to Bayes correlated rather than coarse correlated equilibria.

Last we show that the global mechanism defined by running m (λ, μ)-smooth mechanisms sequentially, is $(\lambda, \mu + 1)$-smooth under this relaxed notion of smoothness.

5.1 Smoothness via Swap Deviations

Definition 5.1.1 (Smooth Mechanism via Swap Deviations)**.** *A mechanism $\mathcal{M}$ is (λ, μ)-smooth via swap deviations if for any valuation profile $v \in \mathcal{V}$, there exists a* *mapping* $\mathbf{a}_i^*(v, \cdot) : \mathcal{A}_i \to \Delta(\mathcal{A}_i)$*, such that for any action profile* $a \in \mathcal{A}$

$$\sum_{i \in [n]} U_i^{\mathcal{M}}(\mathbf{a}_i^*(v, a_i), a_{-i}; v_i) \geq \lambda \textsc{Opt}(v) - \mu \sum_{i \in [n]} P_i(a) \tag{5.1}$$

Obviously if a mechanism is smooth then it is also smooth via swap deviations, hence properties that we prove about smooth mechanisms via swap deviations automatically extend to any smooth mechanism.

If the mechanism is smooth via swap deviations then we show that the price of anarchy of correlated (rather than coarse correlated) equilibria and hence vanishing swap regret sequences of play is small. The reason why the stronger notion of a correlated equilibrium is needed is due to the fact that the swap deviation used in the smoothness definition depends on the previous action of the player. Hence, when considering whether the smoothness deviation is profitable the player must condition on his previous action. The coarse correlated equilibrium condition does not condition on the action of the player.

Theorem 5.1.2. *If a mechanism is* (λ, μ)*-smooth via swap deviations then*

$$\text{CE-}\textsc{PoA} \leq \frac{\max\{1, \mu\}}{\lambda}.$$

Proof sketch. The proof is identical to the proof of Theorem 3.1.2, with the sole extra observation that if $\mathbf{a}$ is a correlated equilibrium then a player's utility is at least as high as his utility from any swap deviation. ■

Swap Deviations and Sequential Games. For smooth mechanisms via swap deviations, we allow the deviating action $\mathbf{a}_i^*(v, a_i)$ to depend both on the valuation vector v and the current action of the deviating player i. This difference causes our Theorem 5.1.2 to only hold for correlated equilibria and not coarse correlated equilibria. Allowing the deviating strategy to depend on a_i makes it possible to prove a composability theorem for sequential mechanisms. Such a relaxation of the smoothness condition allows us to use deviating strategies in the efficiency analysis, where the player follows his old strategy until the "right moment" to deviate arrives. This way the intention of the player to deviate is revealed to the other players only after the deviation and thereby the player is facing a competition at the deviating moment that is identical to the one under his old strategy.

More specifically, by simulating the equilibrium strategy, guarantees that when some special item arrives the distribution of prices that the player faces is equal to the equilibrium prices. Had he deviated from the equilibrium path, then the other players might respond to this deviation and cause the prices on a special item to rise much higher than equilibrium. Thereby, we wouldn't be able to charge these high prices to the payment of some equilibrium winner. In order to simulate the equilibrium behavior (or his previous behavior) the player needs to tailor his deviation to his current action.

The response of players to early deviations is exactly the reason why the unit-demand assumption is needed, as we show in the example presented in Section 5.3.1. For the case when some players are additive, an additive player needs to deviate at many items to grab his optimal allocation. But a deviation at an early item raises the future prices by a significant amount, rendering the deviation unprofitable. However, these raised future prices do not correspond to equilibrium prices and thereby cannot be charged to the payment of some player at equilibrium. This leads to high inefficiency.

5.1.1 Extension to Incomplete Information and Bluffing

We show that the swap deviation smoothness property leads to efficiency guarantees that are almost as robust as those of smooth mechanisms. Specifically, the efficiency directly extends to Bayes correlated equilibria, but not Bayes coarse correlated equilibria.

Theorem 5.1.3. *If a mechanism $\mathcal{M}$ is (λ, μ)-smooth, then for any vector of independent valuation distributions $\mathcal{F} = (\mathcal{F}_1, \ldots, \mathcal{F}_n)$, every Bayes correlated equilibrium has expected social welfare at least $\frac{\lambda}{\max\{1,\mu\}}$ of the expected optimal social welfare, i.e.*

$$\text{BAYES-CE-POA} \leq \frac{\max\{1, \mu\}}{\lambda}.$$

Proof. We will first show that for any strategy profile $s \in \Sigma$, there exists for each player i, a randomized mapping $\mathbf{s}_i^*(v_i, s_i) \in \Delta(\mathcal{A}_i)$, such that:

$$\sum_{i\in[n]} \mathbb{E}_{\mathbf{v}}\left[U_i^{\mathcal{M}}(\mathbf{s}_i^*(\mathbf{v}_i, s_i), s_{-i}(\mathbf{v}_{-i}); \mathbf{v}_i)\right] \geq \lambda \mathbb{E}_{\mathbf{w}}\left[\text{OPT}(\mathbf{w})\right] - \mu \mathbb{E}_{\mathbf{v}}\left[\mathcal{R}^{\mathcal{M}}(s(\mathbf{v}))\right] \quad (5.2)$$

Consider the following randomized deviation for each player i that depends only on the information that he has which is his own value v_i and the equilibrium strategies $s(\cdot)$: He random samples a valuation profile $w \sim \times_i F_i$. Then he plays $\mathbf{s}_i^*(v_i, s_i) = \mathbf{a}_i^*((v_i, \mathbf{w}_{-i}), s_i(\mathbf{w}_i))$, i.e., the player considers the current action profile $s(w)$, using the randomly sampled type (including the random sample of his own type), and deviates from this action profile using the action given by the smoothness property for his true type v_i, the random sample of the types of the others w_{-i}, and the equilibrium action $s_i(w_i)$ of his randomly sampled type w_i. Using the action $s_i(w_i)$ as the base, is particularly meaningful in sequential mechanisms, where it corresponds to a bluffing technique, where player i "pretends" that his valuation was w_i until he deviates.

The utility of the player under this deviation, in expectation over valuations, can be lower bounded as follows:

$$\begin{aligned}\mathbb{E}_{\mathbf{v}}\left[U_i^{\mathcal{M}}(\mathbf{s}_i^*(\mathbf{v}_i, s_i), s_{-i}(\mathbf{v}_{-i}); \mathbf{v}_i)\right] &= \mathbb{E}_{\mathbf{v},\mathbf{w}}\left[U_i^{\mathcal{M}}(\mathbf{a}_i^*((\mathbf{v}_i, \mathbf{w}_{-i}), s_i(\mathbf{w}_i)), s_{-i}(\mathbf{v}_{-i}); \mathbf{v}_i)\right]\\ &= \mathbb{E}_{\mathbf{v},\mathbf{w}}\left[U_i^{\mathcal{M}}(\mathbf{a}_i^*((\mathbf{w}_i, \mathbf{w}_{-i}), s_i(\mathbf{v}_i)), s_{-i}(\mathbf{v}_{-i}); \mathbf{w}_i)\right]\\ &= \mathbb{E}_{\mathbf{v},\mathbf{w}}\left[U_i^{\mathcal{M}}(\mathbf{a}_i^*(\mathbf{w}, s_i(\mathbf{v}_i)), s_{-i}(\mathbf{v}_{-i}); \mathbf{w}_i)\right],\end{aligned}$$

where the second equation is an exchange of variable names and regrouping using independence. By summing over all players and using the smoothness property:

$$\begin{aligned}\sum_{i\in[n]}\mathbb{E}_{\mathbf{v}}\left[U_i^{\mathcal{M}}(\mathbf{s}_i^*(\mathbf{v}_i, s_i), s_{-i}(\mathbf{v}_{-i}); \mathbf{v}_i)\right] &= \mathbb{E}_{\mathbf{v},\mathbf{w}}\left[\sum_{i\in[n]} U_i^{\mathcal{M}}(\mathbf{a}_i^*(\mathbf{w}, s_i(\mathbf{v}_i)), s_{-i}(\mathbf{v}_{-i}); \mathbf{w}_i)\right]\\ &\geq \mathbb{E}_{\mathbf{v},\mathbf{w}}\left[\lambda \textsc{Opt}(\mathbf{w}) - \mathcal{R}^{\mathcal{M}}(s(\mathbf{v}))\right]\end{aligned}$$

which is the initially claimed property.

We now proceed to the final part of the theorem. Let s $\in \Delta(\Sigma)$, be a BAYES-CE. Since no player i, wants to deviate to any mapping $s_i^*(\cdot,\cdot) : \Sigma_i \to \Sigma_i$ in the support of the randomized mapping $\mathbf{s}_i^*(\cdot,\cdot)$, we get that:

$$\begin{aligned}\mathbb{E}_{\mathbf{s}}\mathbb{E}_{\mathbf{v}}\left[U_i^{\mathcal{M}}(\mathbf{s}(\mathbf{v}); \mathbf{v}_i)\right] &\geq \mathbb{E}_{\mathbf{s}}\mathbb{E}_{\mathbf{v}}\left[U_i^{\mathcal{M}}(\mathbf{s}_i^*(\mathbf{v}_i, \mathbf{s}_i), s_{-i}(\mathbf{v}_{-i}); \mathbf{v}_i)\right]\\ &\geq \lambda\mathbb{E}_{\mathbf{w}}\left[\textsc{Opt}(\mathbf{w})\right] - \mu\mathbb{E}_{\mathbf{s}}\mathbb{E}_{\mathbf{v}}\left[\mathcal{R}^{\mathcal{M}}(s(\mathbf{v}))\right]\end{aligned}$$

By quasi-linearity of utility and using the fact that players have the possibility to withdraw from the mechanism, we get the result, by standard manipulations (see Theorem 3.1.2). ■

5.2 Sequential Composability of Smooth Mechanisms

We consider a setting where m mechanisms take place sequentially in the predefined alphabetical order. We will view the normal form representation of the defined sequential game as another global mechanism. An interesting aspect of the sequential composition is that the strategy of a player in the global mechanism is no longer just an action $a_i^j \in \mathcal{A}_i^j$ for each mechanism but rather a whole contingency plan of what action she will submit to mechanism $\mathcal{M}_j$ conditional on any observed history of play. Our result doesn't depend on what part of the history is observed by the players, whether players just observe their own allocation, or all allocations, or also all prices, or bids. We don't even need that all players observe the same things.

We prove that if each mechanism $\mathcal{M}_j$ is (λ, μ)-smooth via swap deviations (obviously also if it is simply smooth), then the global sequential mechanism is $(\lambda, \mu + 1)$-smooth via swap deviations, if an agent's valuation is the best of her allocations over the different mechanisms:

$$v_i(x_i) = \max_{j \in [m]} v_i^j(x_i^j) \tag{5.3}$$

Theorem 5.2.1 (Sequential Composition)**.** *Consider a 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 $(\lambda, \mu + 1)$-smooth, independent of the information released to players during the sequential rounds.*

Proof. Consider a valuation profile v and an action profile a of the sequential composition. Remember that in the sequential composition a_i is not a strategy a_i^j for each j but rather a whole contingency plan of what action $a_i^j(h_i^j)$ to use at mechanism $\mathcal{M}_j$, conditional on the observed history of play by player i up till mechanism $\mathcal{M}_j$.

Let x^* be the optimal allocation for valuation profile v. As stated, we assume that players have unit-demand valuations of the form:

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

where $v_i^j \in \mathcal{V}_i^j$. We will denote with $j_i^* = \arg\max_{j\in[m]} v_i^j(x_{ij}^*)$, i.e. $v_i(x_i^*) = v_i^{j_i^*}(x_{ij_i^*}^*)$.

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

$$\sum_i U_i^{\mathcal{M}}(\mathbf{a}_i^*(v, a_i), a_{-i}; v_i) \geq \lambda \textsc{Opt}(v) - (1+\mu)\mathcal{R}^{\mathcal{M}}(a)$$

Remember that this will be a randomization over contingency plans.

Consider the following type of randomized deviation $\mathbf{a}_i^* = \mathbf{a}_i^*(v, a_i)$ for player i (the deviation will also be a contingency plan for each history of play): he plays exactly as in a_i until mechanism $j = j_i^*$ and then he plays some randomized action $\mathbf{a}_{ij}^*$ that will be determined later on and will be related to the smoothness of mechanism $\mathcal{M}_j$. The utility of player i from this deviation is at least:

$$\begin{aligned} U_i^{\mathcal{M}}(\mathbf{a}_i^*(v, a_i), a_{-i}; v_i) \geq\ & E_{\mathbf{a}_{ij}^*}[v_i^j(X_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))) - P_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))] - \sum_{j'<j} \mathbb{E}[P_i^j(a)] \\ \geq\ & E_{\mathbf{a}_{ij}^*}[v_i^j(X_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))) - P_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))] - \mathbb{E}[P_i(a)] \end{aligned}$$

Also $a_{-i}^j(h_{-i}^j)$ is the action profile submitted by the rest of the players at mechanism j when players use contingency plan a in the global game and hence each observes a history h_i^j produced by this plan.

Summing over all players we get:

$$\sum_i U_i^{\mathcal{M}}(\mathbf{a}_i^*(v, a_i), a_{-i}; v_i) \geq$$
$$\sum_j \sum_{i: j_i^* = j} E_{\mathbf{a}_{ij}^*}[v_i^j(X_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))) - P_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))] - \mathcal{R}^{\mathcal{M}}(a)$$

Now observe that

$$I \triangleq \sum_{i: j_i^* = j} E_{\mathbf{a}_{ij}^*}[v_i^j(X_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))) - P_i^j(\mathbf{a}_{ij}^*, a_{-i}^j(h_{-i}^j))]$$

represents the sum of utilities when each player unilaterally deviates to $\mathbf{a}_{ij}^*$ in mechanism $\mathcal{M}_j$ while previously everyone was playing $a_i^j(h_i^j)$ (the action that they are submitting to mechanism $\mathcal{M}_j$ when each seeing a history of play h_i^j in previous mechanisms) and when all players with $j = j_i^*$ have valuations $v_i^j : \mathcal{X}_i^j \rightarrow \mathbb{R}_+$, while the rest of the players have value 0 for any outcome. Observe that this history of play h_i^j is the same as the history of play produced by contingency plan a of the global mechanism, since the deviation of player i didn't change the history up till mechanism $\mathcal{M}_j$. Therefore $a^j(h^j) = (a_i^j(h_i^j))_{i \in [n]}$ is the action profile that would have been played at mechanism m under the contingency plan a.

In fact, the above sum is at least the sum of these utilities, since we also need to subtract the payments of the players with 0 valuation to get exactly the sum of the utilities. Let $v^{j,*}$ be the valuation profile consisting of the above induced valuations v_i^j on mechanism $\mathcal{M}_j$.

The value of the optimal outcome in such a setting for mechanism $\mathcal{M}_j$ is at least the value of outcome $x^{j,*}$. Hence, the smoothness of mechanism $\mathcal{M}_j$ says that there must exist a strategy $\mathbf{a}_{ij}^* = \mathbf{a}_{ij}^*(v^j, a_i^j(h_i^j))$ such that:

$$I \geq \lambda \sum_{i: j_i^* = j} v_i^j(x_{ij}^*) - (1 + \mu) \sum_i P_i^j(a^j(h^j))$$

Since the utilities of the agents with 0 valuation never help the left hand side of the above sum, smoothness actually implies that there exist strategies $\mathbf{a}_{ij}^* = \mathbf{a}_{ij}^*(v^j, a_i^j(h_i^j))$ only for the players with non-zero valuation, such that the sum over the utilities of only those agents is at least the right hand side in the above equation.

Note again that $P_i^j(a^j(h^j))$ is the payment made at mechanism j under strategy profile a of the global game, since the deviation of the player didn't change the history of play.

Thus if we set the randomized strategies of the players to follow the above smoothness deviation and sum over all players we will get the global smoothness property. Observe that the deviation $\mathbf{a}_{ij}^*(v_j, a_i^j(h_i^j))$ is a whole contingency plan: play until mechanism j and then observe h_i^j; conditional on h_i^j figure out which action you would have played under your initial strategy a_i; then use the smoothness deviation corresponding to this action. ■

5.3 Example: Sequential First Price Auctions

We instantiate the theory presented in this chapter for the case of sequential first price single item auctions. Such an auction game has a long history in the economics literature, starting from the seminal works of Milgrom and Weber [53] and Weber [69]. However, almost all of the literature makes several simplifying assumptions (e.g. symmetry among items and among bidders, two items or two bidders etc.) so as to make an equilibrium characterization feasible. Moreover, most these assumptions either lead to a fully efficient allocation (e.g. symmetry) or are limited in their applicability (e.g. two items or two bidders).

The results of this chapter allows us to quantify the inefficiency of sequential first price auctions with incomplete information and arbitrary asymmetric value distributions. Specifically, since by Lemma 3.1.3 the first price auction is a $\left(1-\frac{1}{e},1\right)$-smooth mechanism, Theorem 5.2.1 implies that the sequential first price auction mechanism is $\left(1-\frac{1}{e},2\right)$-smooth via swap deviations. Subsequently Theorem 5.1.3 implies that the expected welfare at every Bayes-Nash (and Bayes-correlated) equilibrium, under unit-demand (non-identical items) valuations drawn from independent asymmetric distributions is at least $\frac{1}{2}\left(1-\frac{1}{e}\right)\approx 0.316$ of the expected optimal matching allocation. Moreover, this result holds irrespective of what information is revealed at the end of each auction.

We point out that inefficiency can arise even at subgame perfect equilibria of the complete information game, when items are not identical and players have asymmetric valuations. Consider the example given in Figure 5.1 of a sequential first price auction of three items among four players. Player b prefers to loose the first item, anticipating that he might get a similar item for a cheaper price later. This gives an example where the Price of Anarchy is $3/2$. Notice that this is the only equilibrium using non weakly-dominated strategies.

Sequential composability proof unraveled. For better understanding of the main result of this chapter, we describe here the smoothness deviation that is constructed in the proof of Theorem 5.2.1 for the special case of sequential first price auctions with unit-demand bidders.

Denote with $j^*(i)$ the item allocated to player i in the optimal matching allocation. Then the smoothness deviation is as follows: behave exactly as your

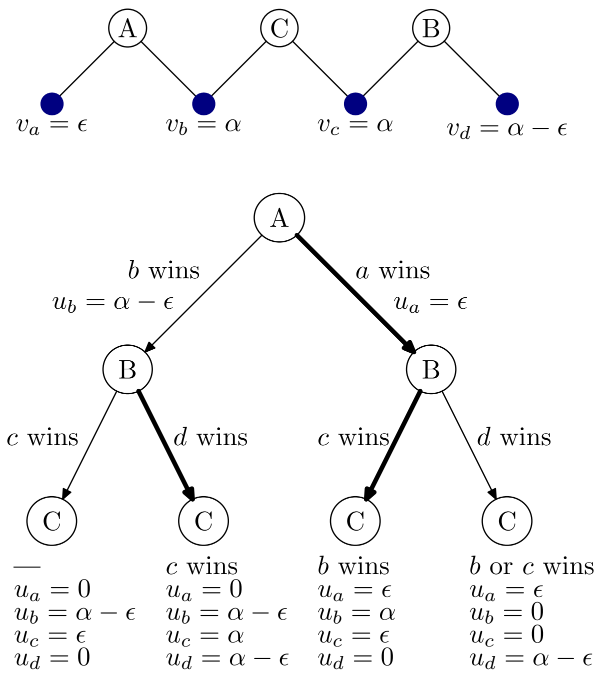

Figure 5.1: Sequential Multi-unit Auction generating POA $3/2$: there are 4 players $\{a, b, c, d\}$ and three items that are auctioned first A, then B and then C. The optimal allocation is $b \to A$, $c \to C$, $d \to B$ with value $3\alpha - \epsilon$. There is a subgame perfect equilibrium that has value $2\alpha + \epsilon$. In the limit when ϵ goes to 0 we get POA $= 3/2$.

previous action a_i, until the auction for item j_i^* arrives. Then at that auction submit your local smoothness deviation for value $w_{i,j^*(i)}$ and assuming you are the only player with non-zero value. Thus submit a random bid with density function $f(x) = \frac{1}{w_{i,j^*(i)} - x}$ and support $[0, (1 - 1/e)w_{i,j^*(i)}]$ at item $j^*(i)$, then drop out from the remaining auctions.

From local smoothness, the utility that the player derives from auction $j^*(i)$ is at least: $\left(1 - \frac{1}{e}\right) w_{i,j^*(i)} - \mathcal{R}^{\mathcal{M}_{j^*(i)}}(b^{j^*(i)})$. Moreover, while simulating his previous action he paid at most his payment $\mathbb{E}[P_i(a)]$ under the previous action profile. Therefore, the overall utility from the deviation is at least: $\left(1 - \frac{1}{e}\right) w_{i,j^*(i)} - \mathcal{R}^{\mathcal{M}_{j^*(i)}}(b^{j^*(i)}) - \mathbb{E}[P_i(a)]$. The global $(\lambda, \mu + 1)$-smoothness via swap deviations follows by summing over all players and observing that $\sum_{i\in[n]} w_{i,j^*(i)} = \textsc{Opt}(v)$ and $\sum_i \mathcal{R}^{\mathcal{M}_{j^*(i)}}(b^{j^*(i)}) = \mathcal{R}^{\mathcal{M}}(b)$, since $j^*(\cdot)$ is a matching.

Bluffing and extension to incomplete information unraveled. When extending the efficiency guarantee to incomplete information, the deviation that is constructed by the extension Theorem 5.1.3 is as follows: randomly sample a valuation profile w_{-i} for your opponents and target the item that you are allocated in the optimal matching for valuation profile (v_i, w_{-i}). Denote this with $j^*(i)$. Then draw a sample w_i of your own value and perform as that randomly sampled player would have performed at the previous strategy profile, until item $j^*(i)$ arrives. Then deviate to the smoothness deviation described in the previous paragraph.

Randomly sampling your own type and playing as a random sample corresponds to a bluffing deviation, since a player pretends to be some other value, until his item of interest arrives. The reason for this bluffing trick is so that the rest of the players will not be able to infer anything about his true valuation through the observed history of play. Thus when item $j^*(i)$ arrives the player is facing a random price that is drawn from the same distribution as the ex-ante equilibrium distribution and not the equilibrium distribution conditional on his true valuation, which can potentially be arbitrarily higher.

5.3.1 Necessity of Unit-Demand Assumption

We will show that if we depart from the unit-demand assumption then we can no longer hope the local smoothness property of a mechanism to imply global guarantees that don't degrade with the size of the market. Specifically, we give a complete information example of a sequential first price auction with some players being additive and some being unit-demand, where the inefficiency grows linearly with the number of players or the number of items. Thus even the presence of additive players destroys the sequential composability property of smooth mechanisms.

Theorem 5.3.1. *The price of anarchy of the sequential first-price item auctions with additive and unit-demand bidders is* $\Omega(\min\{n, m\})$. *Moreover, this result persists even if we consider only subgame perfect equilibria that survive iterated elimination of weakly dominated strategies.*

Informal Description. Before we delve into the details of the proof of Theorem 5.3.1, we give a high-level idea of the type of strategic manipulations that lead to inefficiency and compare to the corresponding simultaneous auction.

Consider an auction instance where two additive bidders have identical values for most of the items for sale, but their valuations differ only on the last few items that are sold. Specifically, assume that there are two items Z_1 and Z_2, auctioned last, such that only player 1 has value for Z_1 and only player 2 has value for Z_2. We will refer to these items as the *non-competitive items* and to all other items as the *competitive items*. The additive bidders know that it is hopeless to try to achieve any positive utility from the competitive items on which they have identical interests. The only utility they can ever derive is from the last, non-competitive items on which they don't compete with each other. If these were the only two players in the auction, then we would obtain the optimal outcome: the two bidders would simply compete on each of the competitive items, with one of them acquiring each competitive item at zero utility.[1]

We now imagine adding unit-demand bidders to the auction in order to perturb the optimality. Specifically, suppose there is a unit-demand bidder that has value for the two *non-competitive items*, with the value for item Z_i being slightly less than player i's value for Z_i, $i \in \{1, 2\}$. This endangers the additive bidders' hopes of getting non-negligible utility, since competition from the unit-demand player may drive up the prices of Z_1 and Z_2. The only hope that the additive bidders have is that the unit-demand bidder will have his demand satisfied prior to these final two auctions, in which case the unit-demand bidder would not bother to bid on them. Hence, the two additive bidders would do anything in their power to guide the auction to such an outcome, even if that means sacrificing all the competitive items! This is exactly the effect that we achieve in our construction. Specifically, we create an instance where this competing unit-demand bidder has his demand satisfied prior to the auctions for Z_1 and Z_2 if and only if a very specific outcome occurs: the additive bidders don't bid at all on all the competitive items, but rather other small-valued bidders acquire the competitive items instead. These small-valued bidders contribute almost nothing to the welfare, and therefore all of the welfare from the competitive items is lost.

It is useful to compare this example with what would happen if the auctions were run simultaneously, rather than sequentially. This uncovers the crucial property of sequential auctions that leads to inefficiency: the *ability to respond to deviations*. If all auctions happened simultaneously, then the behavior of the additive bidders that we described above could not possibly be an equilibrium: one additive bidder, knowing that his additive competitor bids 0 on all the com-

[1]In fact, optimality is always achieved when all bidders are additive, in general.

petitive items, would simply deviate to outbid him on the competitive items and get a huge utility. However, because the items are sold sequentially, this deviation cannot be undertaken without consequence: the moment one of the additive bidders deviates to bidding on the competitive items, in all subsequent auctions the competitor will respond by bidding on subsequent competitive items, leading to zero utility for the remainder of the auctions. Moreover, this response need not be punitive, but is rather the only rational response once the auction has left the equilibrium path (since the additive bidders know that there is no way to obtain positive utility in subsequent auctions). Thus, in a sequential auction, an additive player can only extract utility from at most one competitive item, which is not sufficient to counterbalance the resulting utility-loss due to the increased competition on the last non-competitive item.

The Lower Bound. We now proceed with a formal proof of Theorem 5.3.1. Consider an instance with 2 additive players, $k+1$ unit-demand players and $k+3$ items. Denote with $\{a, b\}$ the two additive players and with $\{p_0, p_1, \ldots, p_k\}$ the $k+1$ unit-demand players. Also denote the items with $\{I_1, \ldots, I_k, Y, Z_1, Z_2\}$. The valuations of the additive players are represented by the following table of v_{ij}, where $\epsilon > 0$ is an arbitrarily small constant:

	I_k	$\ldots$	I_1	Y	Z_1	Z_2
a	$1+\epsilon$	$\ldots$	$1+\epsilon$	0	10	0
b	1	$\ldots$	1	0	0	10

In addition, the unit-demand valuations for the remaining $k+1$ players are given by the table of v_{ij} that follows (an empty entry corresponds to a 0 valuation), though now a valuation of a player when getting a set S is $\max_{j \in S} v_{ij}$:

	I_k	I_{k-1}	I_{k-2}	$\ldots$	I_2	I_1	Y	Z_1	Z_2
p_0				$\ldots$			$10-\epsilon$	$10-\epsilon$	$10-\epsilon$
p_1				$\ldots$		δ_1	10		
p_2				$\ldots$	δ_2	δ_2			
$\ldots$									
p_{k-1}		δ_{k-1}	δ_{k-1}	$\ldots$					
p_k	δ_k	δ_k		$\ldots$					

The constants $\delta_1, \ldots, \delta_k$ are chosen to satisfy the following condition:

$$\delta_k > \delta_{k-1} > \ldots > \delta_2 > \delta_1 > \epsilon \tag{5.4}$$

Note that, by taking ϵ to be arbitrarily small, we can take each δ_i to be arbitrarily small as well.

In the optimal allocation, player a gets all the items $I_1, \ldots, I_k$ and Z_1, player b gets Z_2 and player p_1 gets Y. The resulting social welfare is $k(1+\epsilon)+30$. We assume that the auctions take place in the order depicted in the valuation tables: $\{I_k, \ldots, I_1, Y, Z_1, Z_2\}$. We will show that there is a subgame perfect equilibrium for this auction instance such that the unit-demand players win all the items $I_1, \ldots, I_k$. Specifically, player p_i wins item I_i, player a wins Z_1, player b wins Z_2, and player p_0 wins Y, resulting in a social welfare of $30 - \epsilon + \sum_{i=1}^{k} \delta_i$. Taking δ sufficiently small, this welfare is at most 31. This will establish that the price of anarchy for this instance is at least $\frac{k(1+\epsilon)+30}{31} = O(k)$, establishing Theorem 5.3.1. Furthermore, we will show that this subgame perfect equilibrium is *natural*, in the sense that it survives iterated deletion of weakly dominated strategies.

The intuition is the following: after the first k auctions have been sold, player p_0 has to decide if he will target (and win) item Y, or if he will instead target items Z_1 and/or Z_2. If he targets item Y, he competes with player p_1 and afterwards lets players a and b win items Z_1, Z_2 for free. This decision of player p_0 depends on whether player p_1 has won item I_1, which in turn depends on the outcomes of the first $k-1$ auctions. In particular, player p_1 can win item I_1 only if player p_2 has won item I_2. In turn, p_2 can win I_2 only if p_3 has won item I_3 and so on. Hence, it will turn out that in order for p_0 to want to target item Y, it must be that each item I_i is sold to bidder p_i. Thus, if either player a or b acquires any of the items $I_1, \ldots, I_k$, they will be guaranteed to obtain low utility on items Z_1 and Z_2. This will lead them to bidding truthfully on all subsequent I_i auctions, leading to a severe drop in utility gained from future auctions.

In the remainder of this section, we provide a more formal analysis of the equilibrium in this auction instance. We begin by examining what happens in the last three auctions of Y, Z_1 and Z_2, conditional on the outcomes of the first k auctions. We first examine the outcome of auctions Y, Z_1, Z_2 conditional on the outcome of auction I_1:

Case 1: p_1 has won I_1. Player p_1 has marginal value of $10-\delta_1$ for item Y. Hence, he is willing to bid at most $10-\delta_1$ on item Y. Player p_0 knows that if he loses Y then in the subgame perfect equilibrium in that subgame he will bid $10-\epsilon$ on Z_1 and Z_2 and lose. Thus he expects no utility from the future if he loses Y. Thus he is willing to pay at most $10-\epsilon$ for item Y. Since by assumption (5.4) $\delta_1 > \epsilon$, player p_0 will win Y at a price of $10-\delta_1$. Then players a, b will win Z_1 and Z_2 for free. Thus the utilities in this case from this subgame are: $u(a) = 10$, $u(b) = 10$, $u(p_0) = \delta_1 - \epsilon$, $u(p_1) = 0$.

Case 2: p_1 has lost I_1. Player p_1 has marginal value of 10 for item Y. Hence, he is willing to bid at most 10 on item Y. Player p_0 performs the exact same thinking as in the previous case and thereby is willing to bid at most $10-\epsilon$ for item Y. Thus in this case p_1 will win item Y at a price of $10-\epsilon$. Then, as predicted, p_0 will bid $10-\epsilon$ on Z_1 and Z_2 and lose. Thus the utilities of the players in this case are: $u(a)=\epsilon$, $u(b)=\epsilon$, $u(p_0)=0$, $u(p_1)=\epsilon$.

Now we focus on the auction of item I_1. As was explained in Paes Leme et al. [56] this auction will be an auction with externalities where each player has a different utility for each different winner outcome. These utilities can be concisely expressed in a table of v_{ij}'s where v_{ij} is the value of player i when player j wins. The only players that potentially have any incentive to bid on item I_1 are a, b, p_0, p_1, p_2. The following table summarizes their values for each possible winner outcome of auction I_1 as was calculated in the previous case-analysis (we point that in the diagonal we also add the actual value that a player acquires from item I_1 to his future utility conditional on winning I_1) .

$$[v_{ij}] = \begin{array}{c|ccccc} & a & b & p_0 & p_1 & p_2 \\ \hline a & 1+2\epsilon & \epsilon & \epsilon & 10 & \epsilon \\ b & \epsilon & 1+\epsilon & \epsilon & 10 & \epsilon \\ p_0 & 0 & 0 & 0 & \delta_1-\epsilon & 0 \\ p_1 & \epsilon & \epsilon & \epsilon & \delta_1 & \epsilon \\ p_2 & 0 & 0 & 0 & 0 & \delta_2 \cdot \mathbf{1}_{\text{hasn't won } I_2} \end{array}$$

For example, player a obtains utility 10 if player p_1 wins item I_1. We see from the table that, at this auction, everyone except p_2 achieves their maximum value when p_1 wins the auction. Player p_2 has value for winning the auction only if he hasn't won I_2. In addition, since $\delta_2 > \delta_1$, if p_2 hasn't won I_2 then he can definitely outbid p_1 on I_1 and therefore p_1 has no chance of winning the auction of I_1. As we now show, this implies that there is a unique equilibrium of the auction conditioning on whether or not p_2 has won I_2:

Case 1: If p_2 has won I_2 then he has no value for I_1. There exists an equilibrium in undominated strategies where all players a, b, p_0, p_2 will bid 0, while p_1 bids 0^+. In fact this is in some sense the most natural equilibrium since it yields the highest utility for a and b. In this case the utility of the players from auctions I_1 and onward will be: $u(a) = 10$, $u(b) = 10$, $u(p_0) = \delta_1 - \epsilon$, $u(p_1) = \delta_1$, $u(p_2) = 0$.

Case 2: If p_2 has lost I_2, then he has value of $\delta_2 > \delta_1$ for I_1. Hence, p_1 has no chance of winning item I_1. Thus, the unique equilibrium that survives elimination of weakly dominated strategies in this case is for player a to bid 1^+, for player b to bid 1, for player p_0 to bid 0, for player p_1 to bid $\delta_1 - \epsilon$ and for player p_2 to bid δ_2. In this case the utility of the players from auctions I_1 and on will be: $u(a) = 2\epsilon$, $u(b) = \epsilon$, $u(p_0) = 0$, $u(p_1) = \epsilon$, $u(p_2) = 0$.

Using similar reasoning we deduce that player p_i can win I_i only if p_{i-1} has won I_{i-1}. If at any point some p_i does not win I_i then players a and b know that from that point onward no p_j can win auction I_j, and therefore they will get only utility ϵ from Z_1, Z_2. Thus there will be no reason for players a and b to allow unit-demand players to continue to win items, and thus the only equilibrium strategies from that point on will be for a to bid 1^+ on each of I_i and b to bid 1. This will lead to player a to get utility $O(\epsilon)$ from each auction for items $I_{i-1}, \ldots, I_2$, and player b to get no utility from these auctions. Thus, at any point in the auction, it is an equilibrium for players a and b to allow the unit demand player p_i to win auction I_i. This completes the proof of Theorem 5.3.1

Finally, as discussed throughout our analysis, the equilibrium described above survives iterated elimination of weakly dominated strategies. The reason is that, for every item k and bidder i, the proposed equilibrium strategy for bidder i does not require that he bids more than his value for item k less his utility in the continuation game subject to not winning item k. As discussed in Paes Leme et al. [56], this property guarantees that no player is playing a weakly dominated strategy.

6

WEAK SMOOTHNESS AND NO-OVERBIDDING

The second price auction is not a smooth mechanism based on the current definition of smoothness. In fact, the second price auction is not as robust as first price auctions: admits arbitrarily bad equilibria when players bid above their value. Moreover, Goeree [32] shows that signaling is bound to arise in a second price auction when bidders are strategising about future opportunities, and Paes Leme et al [56] show an example with unbounded inefficiency when running second price auctions sequentially and bidders are unit-demand.

The main difference of the second price auction and the previous two auctions is that it makes very loose connection between the bid a player needs to make to win and the price that was previously paid to the auctioneer. Several papers [15, 8, 50, 13] have used an assumption that players will not bid above their valuations to give good efficiency guarantees for second-price type of auctions. In this chapter, we extend our results to mechanisms that require such no-overbidding assumptions.

6.1 Efficiency under the No-Overbidding Assumption

We give a generalization of our framework to capture mechanisms that produce high efficiency under a no-overbidding refinement. In the context of a single-item second price auction, the no-overbidding refinement simply selects only equilibria where no player bids more than his value for the item. Observe that without this refinement, the second-price auction has very inefficient equilibria where low value players bid a huge amount, while high value players bid zero. Thereby such a refinement is necessary even in the single-item second price auction.

First, we give a definition of no-overbidding that generalizes the latter no-overbidding refinement to any mechanism design setting and any mechanism. In a single-item second-price auction the bid of a player is his maximum willingness to pay when he wins. The following defines the notion of a "bid" in the general mechanism design setting.

Definition 6.1.1 (Implicit Bid). *Given a mechanism $(\mathcal{A}, X, P)$, a player's implicit bid for an allocation x_i when using strategy a_i, is defined as the maximum he could ever pay conditional on allocation x_i:*

$$B_i(a_i, x_i) = \max_{a_{-i}:\ X_i(a)=x_i} P_i(a) \tag{6.1}$$

Definition 6.1.2 (Weakly Smooth Mechanism). *A mechanism is weakly (λ, μ_1, μ_2)-smooth for $\lambda, \mu_1, \mu_2 \geq 0$, if for any $v \in \mathcal{V}$ there exists a randomized action $\mathbf{a}_i^*(v)$ for each player i, such that for any action profile a:*

$$\sum_i U_i^{\mathcal{M}}(\mathbf{a}_i^*(v, a_i), a_{-i}; v_i) \geq \lambda \text{OPT}(v) - \mu_1 \mathcal{R}^{\mathcal{M}}(a) - \mu_2 \sum_i B_i(a_i, X_i(a))$$

Similarly, we can also define the relaxed notion of weak smoothness via swap deviations, with the corresponding implications on robustness of guarantees discussed in the previous sections.

Definition 6.1.3 (No-overbidding). *A randomized strategy profile* $\mathbf{a}$ *satisfies the no-overbidding assumption if:*

$$\mathbb{E}_{\mathbf{a}}\left[B_i(\mathbf{a}_i, X_i(\mathbf{a}))\right] \leq \mathbb{E}_{\mathbf{a}}\left[v_i(X_i(\mathbf{a}))\right] \tag{6.2}$$

i.e., at this strategy profile no player is bidding in a way that she could potentially pay more than her value subject to her expected allocation remaining the same.

Note that the smoothness property must be satisfied for any action profile and not only for non-overbidding action profiles. This is essential for the smoothness property to be composable. As an example, note that the single item second price auction is weakly $(1,0,1)$-smooth, but is not $(1,1)$-smooth (which would be true if smoothness was required only for non-overbidding strategies), and while a single item second price auction is optimal, this property is not maintained in composition. In a composition setting, local no-overbidding is not meaningful, since there is no clear induced valuation of a player at each local mechanism and is only meaningful on the overall mechanism.

Following arguments very similar to the proof of Theorem 3.2.1 together with the no-overbidding assumption it is easy to show the following efficiency theorem.

Theorem 6.1.4. *If a mechanism is weakly (λ, μ_1, μ_2)-smooth then any Bayes coarse correlated equilibrium that satisfies the no-overbidding assumption achieves efficiency at least $\frac{\lambda}{\mu_2+\max\{\mu_1,1\}}$ of the expected optimal.*

Composability of weak smoothness. Moreover, it is also easy to show analogous composability theorems, following the same arguments as in Theorems 4.2.2, 4.5.5 and 5.2.1. Specifically, the simultaneous composition of m weakly (λ, μ_1, μ_2)-smooth mechanisms is weakly (λ, μ_1, μ_2)-smooth if players have XOS $-$ $\mathcal{C}$ valuations and $(1 - k + \min\{\lambda, 1\} \cdot k, \mu_1, \mu_2)$ if players have MPH-k valuations across mechanisms. The sequential composition is weakly $(\lambda, \mu_1 + 1, \mu_2)$-smooth if players have unit-demand valuations across mechanisms.

6.2 On the No-Overbidding Assumption

Remark 1. In contrast to the smoothness used in [60] our definition of weakly smooth mechanisms allows us to prove efficiency under the weaker assumption of no-overbidding in expectation, rather than point-wise no-overbidding. The main difference is that we incorporate the willingness-to-pay inside the smoothness definition, while previous smoothness approaches would relate to value directly. The latter approach would require to use point-wise no-overbidding to relate bids to welfare in second-price auctions.

Remark 2. We use the non-overbidding assumption as an equilibrium refinement rather than as a strategy-space restriction. Several papers in the literature have used non-overbidding as a strategy space restriction (rather than as an equilibrium refinement). The two uses are equivalent in settings where the restricted strategy space always contains best-responses. Note that while overbidding is a dominated strategy in a single item auction, global no-overbidding is not dominated when running second price auctions simultaneously or sequentially. Overbidding equilibria that survive elimination of dominated strategies and that have non-constant inefficiency have been given both for the case of sequential [56] and simultaneous [25] second price auctions, even in the simplest scenario when bidders are unit-demand.

Restricting the strategy space to non-overbidding strategies, could potentially create artificial equilibria that were not equilibria of the original game, since this restricted strategy space does not always contain best-responses (see [25] for an example). On the other hand, the refined set of non-overbidding equilibria might be empty. Below we portray these differences via an example.

EXAMPLE 6.2.1. (No-overbidding Restriction vs. Refinement) Consider the following setting: there are two items A, B and two bidders a, b. Each item is sold separately and simultaneously via a second-price auction (breaking ties at random). Bidder a has value 1 for both items and 0 for any item individually (i.e. he is an AND bidder). Bidder b is a "Bayesian" additive player and we describe his additive values for each of the items: with probability $1/3$ he has value $2/3$ only for item a, with probability $1/3$ he has value $2/3$ only for item b and with probability $1/3$ he has a value of $2/3$ for each of the items.

Strategy space restriction. If we use the strategy space restriction of the no-overbidding assumption, then the following is an equilibrium: bidder a bids 0 on both items and bidder b bids "truthfully". Since bidder a can never overbid, it must be that the sum of his bids on the two items is at most 1. If he bids below $2/3$ on both items then he never wins both of them, thereby getting non-positive utility. Thus he must be bidding at least $2/3$ on one of the two items and in fact any such strategy is dominated by bidding more than $2/3$ on one of the two items and a non-zero amount on the other. In that case he wins both items with probability only $1/3$: it is the probability that player b has value only for the item that player a chose to bid more than $2/3$. Moreover, he is paying an expected price of $2/3 \cdot 2/3$. Thus his utility from any such deviation is negative and thereby his utility from any no-overbidding deviation is non-positive, making the current action an equilibrium within the restricted strategy space.

However, the above strategy profile is not a Bayes-Nash equilibrium of the original game without the restriction on the strategy space. Player a can deviate to bidding 1 on both items. In that case he wins the bundle deterministically and he pays an expected price of $2/3 \cdot 2/3 + 1/3 \cdot 4/3 = 8/9$. His utility from this deviation is $1/9$ and thereby profitable. Thus we see that using no-overbidding as a refinement can create artificial equilibria that were not existent in the original game. Most importantly, it disallows the players to make deviating strategies that can be profitable.

Refinement. If we use no-overbidding as a refinement, we are simply searching for a Bayes-Nash equilibrium of the original game, in which no-player bids in a way that his expected bid for the item that he wins is more than his expected value for these items. However, this set of Bayes-Nash equilibria can be potentially empty and in fact they are empty in the above instance.

First observe, that in any Bayes-Nash equilibrium, if player a bids below $2/3$ on some item, then he will lose the item if player b also has a value for it. Observe that because of the strategy of player a at equilibrium must satisfy the no-overbidding constraint, he cannot be winning the bundle when player b has value for both items, since for that to happen he must be bidding at least $2/3$ on both items (which would lead to overbidding). Thus he can only be winning the bundle when player b has value for only one item. Since player a doesn't know which item player b has value for, and since he cannot be bidding more than $2/3$ on both items, he can only be winning the bundle with probability at most $1/3$. Moreover, for any such no-overbidding bid that wins the bundle with probability $1/3$, the player has to be paying an expected payment of $4/3$. Leading to non-positive utility. Thus with any bid that satisfies the no-overbidding assumption, the player must be making 0 utility, which would then not render it an equilibrium, since he can always switch to bidding 1 on both items. Thus the refined set of non-overbidding equilibria is empty.

In this particular example, the main reason for the emptiness is that player a has complementary valuations across the items. It is an interesting open question whether Bayes-Nash equilibria with no-overbidding always exist in simultaneous second price auctions, when players have XOS valuations. This fact is known only in the complete information setting (see Christodoulou et al. [15]). ■

Some of our results carry over to the strategy-space restriction version. Specifically, when the smoothness deviations fall within the space of non-overbidding actions then the efficiency guarantees hold even when no-overbidding is used as a strategy-space restriction.

Moreover, it is easy to see that we can relax the no-overbidding refinement at the expense of efficiency, by saying that a player doesn't bid more than γ times his expected value. In that case the welfare guarantee of any weakly (λ, μ_1, μ_2)-smooth mechanism, will become $\frac{\lambda}{\gamma\mu_2+\max\{1,\mu_1\}}$.

6.3 Example: Simultaneous Second Price Auctions

We revisit the setting of simultaneous single-item auctions and now analyze the case of simultaneous second price auctions, under the lens of the results of this chapter. We portray how the price of anarchy results of Christodoulou et al [15] and Bhawalkar et al [8] on the Bayes-Nash price of anarchy of this setting, can be re-interpreted as a corollary of the weak smoothness property of a single item second price auction.

Lemma 6.3.1. *The second price auction is a weakly* $(1, 0, 1)$*-smooth mechanism.*

Proof. Consider valuation profile $v = (v_1, \ldots, v_n)$. The highest value player (wlog player 1) can deviate to submitting his true value v_1, while all non-highest value players should just deviate to bidding 0. No matter what the rest of the players are bidding, the utility of the highest bidder from the deviation is lower bounded as follows:

$$U_1^{\text{SPA}}(v_1, b_{-1}; v_1) = \left(v_1 - \max_{i \neq 1} b_i\right) \cdot \mathbf{1}\{v_1 > \max_{i \neq 1} b_i\} \geq \ v_1 - \max_{i \neq 1} b_i \geq v_1 - \max_i b_i$$

The result follows by observing that $\sum_i B_i(b_i, X_i(b)) = \max_i b_i$. ■

The weak $(1, 0, 1)$-smoothness property of the second price auction, combined with the composability theorems for weakly smooth mechanisms implies that simultaneous second price auctions is a weakly $(1, 0, 1)$-smooth mechanism when players have XOS valuations over the items. Hence, when valuations are drawn independently from arbitrary distributions every Bayes coarse correlated equilibrium that satisfies the no-overbidding assumption achieves at least $1/2$ of the expected optimal welfare.

Two-approximation for any monotone valuation. Quite surprisingly for MPH-k valuations, because $\lambda = 1$, we get that the simultaneous second price auction mechanism is weakly $(1, 0, 1)$-smooth independent of k. Thus every Bayes coarse correlated equilibrium that satisfies the no-overbidding assumption achieves at least $1/2$ of the expected optimal welfare for arbitrary monotone valuations (since MPH-m contains all monotone valuations). In other words, the price of anarchy bound for MPH-k valuations does not degrade with k.

While this result seems quite positive, we suggest that it can be interpreted as indicating the strength of the no-overbidding refinement, especially in the presence of complementarities. Indeed, the efficiency result is conditional on equilibrium existence, and the set of no-overbidding equilibria (i.e. equilibria where the expected sum of bids of a player for the items he won is at most his expected value for the items he won) might be empty.

7
BUDGET CONSTRAINTS

An important class of non-quasilinear preferences is when players have hard budget constraints on the payments they make, i.e.:

$$u_i(x_i, p_i; v_i, B_i) = \begin{cases} v_i(x_i) - p_i & \text{if } p_i \leq B_i \\ -\infty & \text{o.w.} \end{cases} \tag{7.1}$$

Studying the effect of budgets on the design of efficient mechanisms has received great attention in recent algorithmic game theory literature [18, 28, 31, 20] mostly in the realm of truthful mechanism design and assuming that the budgets are common knowledge. Little is known about the effect of budgets in the case of non-truthful mechanisms. For instance, only recently Huang et al. [37] analyzed efficiency in a two-player sequential first price auction game with budget constraints in the complete information setting.

In this chapter we examine the efficiency of smooth mechanisms in the presence of budget constraints and at a high level show the following informal theorem:

Informal Theorem 4. *If a mechanism is (λ, μ)-smooth and the smoothness deviations satisfy a minimal "conservative" assumption, then all the efficiency guarantees extend to the budget constraint setting but with respect to the optimal welfare achievable if every player's valuation is capped by his budget.*

7.1 The Effective Welfare Benchmark

Most of the literature has focused on producing pareto-optimal outcomes, i.e. a pair of allocation and prices such that there is no other pair that respects feasibility and budget constraints and such that all players have strictly higher utility and the auctioneer receives strictly higher revenue.

Instead, we analyze an orthogonal benchmark, which we call *Effective Welfare*, obtained by capping a player's value by his budget:

$$EW(x) = \sum_i \min\{v_i(x_i), B_i\} \tag{7.2}$$

We compare the social welfare of a mechanism to the maximum possible effective welfare. This benchmark reflects that we cannot expect players with low budgets to be effective at maximizing their own value.

The effective welfare benchmark was also analyzed in the context of truthful auctions by Dobzinski and Paes Leme [19]. However, [19] quantifies the even stronger ratio of the effective welfare (rather than social welfare) of the mechanism's outcome over the optimal effective welfare. Here we use effective welfare as a benchmark to analyze the welfare of non-truthful mechanisms at equilibrium.

7.2 Smoothness via Conservative Deviations

We show that a lot of our results carry over to the effective welfare benchmark, by introducing a strengthening of the smoothness property of mechanisms (a strengthening that is satisfied by almost all the applications we consider) and assuming that the valuation spaces $\mathcal{V}_i$ considered is closed under capping, i.e., for any valuation $v_i \in \mathcal{V}_i$ and any budget B_i, we also have $\min\{v_i(\cdot), B_i\} \in \mathcal{V}_i$.

We focus on smooth mechanisms, but all the results in this section extend to smooth mechanisms via swap deviations (with the exception that bounds extend to BAYES-CE rather than BAYES-CCE for reasons explained in Chapter 5) and to weakly smooth mechanisms assuming no-overbidding.

Definition 7.2.1 (Smooth Mechanism via Conservative Deviations). *A mechanism is conservatively (λ, μ)-smooth if it is (λ, μ)-smooth in the quasilinear utility setting and the smoothness deviations satisfy that for any a_i^* in the support of $\mathbf{a}_i^*(v)$, player i can never pay more than his maximum valued allocation:*

$$\max\{p : a_{-i} \in \mathcal{A}_{-i} \text{ and } p \in \text{SUPP}(P_i(a_i^*, a_{-i}))\} \leq \max_{x_i \in \mathcal{X}_i} v_i(x_i) \tag{7.3}$$

The next theorem shows that the expected social welfare at BAYES-CCE of conservatively smooth mechanisms is a good fraction of the optimal effective welfare, when players have budget constraints. Note that in the incomplete information setting, the private information of a player is his valuation and his budget. We will denote the valuation and budget pair as the type $t_i = (v_i, B_i)$ of player i and we will assume that it is distributed independently according to some distribution $\mathcal{F}_i$ on $\mathcal{V}_i \times \mathbb{R}_+$. Note that we allow the budget of a player to be correlated with his valuation. Under this notation we will write the utility of a player from mechanism $\mathcal{M}$ under budget constraint B_i as:

$$\hat{U}_i^{\mathcal{M}}(a; t_i) = \begin{cases} U_i^{\mathcal{M}}(a; v_i) & \text{if } \max\{p : p \in \text{SUPP}(P_i(a))\} \leq B_i \\ -\infty & \text{o.w.} \end{cases} \tag{7.4}$$

Theorem 7.2.2 (Efficiency with Budgets). *If a mechanism is conservatively (λ, μ)-smooth and its valuation space is closed under capping, then the social welfare at any*

Bayes coarse correlated equilibrium is at least $\frac{\lambda}{\max\{1,\mu\}}$ of the expected maximum effective welfare, even if players have private budget constraints.

Proof. We begin with the following observation: if for any action profile a in the support of a random action profile $\mathbf{a}$ it holds that $\max\{p : p \in \textsc{supp}(P_i(a_i^*, a_{-i}))\} \leq B_i$ then the expected utility of a player with a budget constraint B_i is the same as the expected utility of an unconstrained player with quasi-linear utility.

Consider a valuation and budget profile $t = (v, B)$. Let $\hat{v}$ be the corresponding capped valuation profile where each players valuation is replaced with $\hat{v}_i = \min\{v_i, B_i\}$.

Since we assumed that the mechanism is (λ, μ)-smooth via conservative deviations and the valuation space is closed under capping, for any strategy profile a there exists a randomized strategy $\mathbf{a}_i^*(\hat{v})$ for each player such that under the quasilinear utility setting:

$$\sum_i U_i^{\mathcal{M}}(\mathbf{a}_i^*(\hat{v}), a_{-i}; \hat{v}_i) \geq \lambda \textsc{Opt}(\hat{v}) - \mu \mathcal{R}^{\mathcal{M}}(a)$$

and such that for all a_i^* in the support of $\mathbf{a}_i^*(\hat{v})$:

$$\max\{p : \tilde{a}_{-i} \in \mathcal{A}_{-i} \text{ and } p \in \textsc{supp}(P_i(a_i^*, \tilde{a}_{-i}))\} \leq \max_{x_i \in \mathcal{X}_i} \hat{v}_i(x_i) \leq B_i$$

Suppose that in the budgeted setting each player i deviates to $\mathbf{a}_i^*(\hat{v})$. Then by the above conservativeness of this deviating strategy and the initial observation we know that the expected utility of a budgeted player under this deviation is the same as the expected utility of a player with quasi-linear utilities and value v_i. Subsequently the expected utility of a player with quasi-linear utilities and value v_i is at least the utility of a player with quasi-linear utilities and value $\hat{v}_i$. Thus we get:

$$\sum_i \hat{U}_i^{\mathcal{M}}(\mathbf{a}_i^*(\hat{v}), a_{-i}; t_i) = \sum_i U_i^{\mathcal{M}}(\mathbf{a}_i^*(\hat{v}), a_{-i}; v_i) \geq \sum_i U_i^{\mathcal{M}}(\mathbf{a}_i^*(\hat{v}), a_{-i}; \hat{v}_i)$$
$$\geq \lambda \textsc{Opt}(\hat{v}) - \mu \mathcal{R}^{\mathcal{M}}(a) \tag{7.5}$$

Using the above property we can now complete the proof of the Theorem similar to the proofs of Theorems 3.3.7 and 5.1.3. The only extra point we need to make is that due to the fact that a player can always drop out we know that at any equilibrium solution concept no player is going to ever be exceeding his budget at any action profile in the support of the solution concept since otherwise his utility would have been minus infinity. Hence, at any action profile in the support of an equilibrium the utility of a player will behave as if quasilinear. For completeness we present the proof.

Consider the following feasible deviation of player i under incomplete information: she random samples a type profile $\tau = (w, C) \sim \mathcal{F}$. Let $\hat{w}_i = \min\{w_i, C_i\}$ be the capped random sampled valuations. Then he plays $\mathbf{s}_i^*(t_i) = \mathbf{a}_i^*(\hat{v}_i, \hat{w}_{-i})$. By similar reasoning as in the proof of Theorem 3.2.1 we can show that for any strategy profile in the incomplete information game $s : \Sigma$:

$$\mathbb{E}_{\mathbf{t}}\left[\hat{U}_i^{\mathcal{M}}(\mathbf{s}_i^*(\mathbf{t}_i), s_{-i}(\mathbf{t}_{-i}); \mathbf{t}_i)\right] = \mathbb{E}_{\mathbf{t},\tau}\left[\hat{U}_i^{\mathcal{M}}(\mathbf{a}_i^*(\hat{\mathbf{w}}), s_{-i}(\mathbf{t}_{-i}); \tau_i)\right]$$

Summing over all players and using Equation (7.5):

$$\sum_{i\in[n]} \mathbb{E}_{\mathbf{t}}\left[\hat{U}_i^{\mathcal{M}}(\mathbf{s}_i^*(\mathbf{t}_i), s_{-i}(\mathbf{t}_{-i}); \mathbf{t}_i)\right] = \mathbb{E}_{\mathbf{t},\tau}\left[\sum_{i\in[n]} \hat{U}_i^{\mathcal{M}}(\mathbf{a}_i^*(\hat{\mathbf{w}}), s_{-i}(\mathbf{t}_{-i}); \tau_i)\right]$$
$$\geq \lambda \mathbb{E}_\tau\left[\textsc{Opt}(\hat{\mathbf{w}})\right] - \mu \mathbb{E}_{\mathbf{t}}\left[\mathcal{R}^{\mathcal{M}}(s(\mathbf{t}))\right]$$

If s $\in \Delta(\Sigma)$ is a BAYES-CCE, then each player doesn't want to deviate to the above s_i^*. Thus taking expectation over s of the above equation and by standard manipulations we get the theorem. ∎

7.3 Simultaneous Composability with Budget Constraints

Next, we show that efficiency guarantees for bidders with budget-constraints are simultaneously composable under the conservative smoothness property. Unfortunately, sequential composition doesn't carry over. In sequential mechanisms a good deviation may require that the player waits and plays according to equilibrium until his optimal mechanism arrives. While "waiting" he might exhaust his budget.

Theorem 7.3.1 (Simultaneous Composition with Budgets)**.** *Consider a simultaneous composition of m (λ, μ)-smooth mechanisms via conservative deviations and let $\mathcal{C}_i = \mathcal{V}_i^1 \times \ldots \times \mathcal{V}_i^m$. If $\mathcal{V}_i^j$ are closed under cappings and the valuation $v_i : \mathcal{X}_i \to \mathbb{R}^+$ of each player across mechanisms is $\text{XOS} - \mathcal{C}_i$, then the global mechanism is also (λ, μ)-smooth via conservative deviations and $\text{XOS} - \mathcal{C}_i$ is also closed under capping.*

Proof. The composability result is proved in a sequence of two lemmas: first we prove that the conservative smoothness property of a mechanism composes under XOS valuations and second we show that if the valuation space of each component mechanism is closed under capping then the corresponding valuation space of the composition mechanism is also closed under capping. The latter is shown by proving a structural property of XOS valuations: a valuation produced by capping an XOS valuation is also XOS and can be described by component valuations that are cappings of the component valuations of the XOS representation of the initial valuation.

Lemma 7.3.2. *The simultaneous composition of m (λ, μ)-smooth mechanisms via conservative deviations is also (λ, μ)-smooth via conservative deviations, when the valuation $v_i : \mathcal{X}_i \to \mathbb{R}^+$ of each player across mechanisms is* XOS-$\mathcal{C}_i$.

Proof. We need to show that the simultaneous composition is smooth in the quasi-linear setting and that the deviations used to show smoothness satisfy the property that every action a_i in their support satisfies Equation (7.3).

The fact that the composition is smooth just stems from Theorem 4.2.2, since each component mechanism is conservatively smooth and thereby smooth.

From the proof of Theorem 4.2.2 we know that for each action profile a the deviation that is used in the smoothness argument is a randomized deviation $\mathbf{a}_i^*(v)$ that consists of independent randomized deviations for each mechanism j following the distribution of $\mathbf{a}_i^{j,*}(v^{j,*})$, where $v^{j,*}$ is the valuation profile for mechanism j where each player has valuation $v_i^{j,*}$ on $\mathcal{X}_i^j$ and where v_i^* is the additively separable valuation that corresponds to allocation x_i^* according to the XOS-$\mathcal{C}_i$ definition, i.e., $v_i(x_i^*) = \sum_j v_i^{j,*}(x_{ij}^*)$ and for all $x_i \in \mathcal{X}_i$: $v_i(x_i) \geq \sum_j v_i^{j,*}(x_i^j)$.

Now by conservative smoothness of each component mechanism $\mathcal{M}_j$ we know that for any action $a_i^{j,*}$ in the support of $\mathbf{a}_i^{j,*}(v_j^*)$:

$$\max\{p_i^j : a_{-i}^j \in \mathcal{A}_{-i}^j \text{ and } p_i^j \in \textsc{supp}(P_i^j(a_i^{j,*}, a_{-i}^j))\} \leq \max_{x_i^j \in \mathcal{X}_i^j} v_i^{j,*}(x_i^j) \tag{7.6}$$

For each mechanism $\mathcal{M}_j$ let $\hat{x}_i^j$ be the allocation that corresponds to the maximizer on the right hand side of the above inequality.

By the fact that action spaces are independent across mechanisms it is easy to see that for any action a_i:

$$\max\{\sum_{j\in[m]} p_i^j : a_{-i}^j \in \mathcal{A}_{-i}^j \text{ and } p_i^j \in \text{SUPP}(P_i^j(a_i^{j,*}, a_{-i}^j))\} = \sum_{j\in[m]} \max\{p_i^j : a_{-i}^j \in \mathcal{A}_{-i}^j \text{ and } p_i^j \in \text{SUPP}(P_i^j(a_i^{j,*}, a_{-i}^j))\} \quad (7.7)$$

Any action a_i in the support of the randomized deviation $\mathbf{a}_i^*(v)$ is going to consist of actions a_i^j in the support of the $\mathbf{a}_i^{j,*}(v^{j,*})$. By Equations (7.6) and (7.7) and by the fact that v_i^* is part of the XOS representation of v_i we get that for any a_i in the support of the smoothness deviation $\mathbf{a}_i^*(v)$:

$$\begin{aligned}\max\{\sum_{j\in[m]} p_i^j : a_{-i}^j \in \mathcal{A}_{-i}^j \text{ and } p_i^j \in \text{SUPP}(P_i^j(a_i^{j,*}, a_{-i}^j))\} &\leq \sum_j v_i^{j,*}(\hat{x}_i^j) \\ &\leq v_i(\hat{x}_i^1, \ldots, \hat{x}_i^m) \\ &\leq \max_{x_i \in \mathcal{X}_i} v_i(x_i)\end{aligned}$$

The latter completes the proof. ■

To complete the proof we show a property of capped XOS valuations across mechanism outcomes:

Lemma 7.3.3. *Suppose that a valuation $v : \mathcal{X}_i \to \mathbb{R}^+$ across m mechanisms is* XOS *and can be represented by a set of additively separable valuations $\mathcal{L}$, i.e. $v_i(x_i) = \max_{\ell\in\mathcal{L}} \sum_{j\in[m]} v_i^{j,\ell}(x_i^j)$. Then the capped valuation $\hat{v}(x_i) = \min\{v(x_i), B_i\}$ is also* XOS *and can be expressed using valuations $\hat{v}_i^{j,\ell} : \mathcal{X}_i^j \to \mathbb{R}^+$ that are cappings of the induced valuations $v_i^{j,\ell} : \mathcal{X}_i^j \to \mathbb{R}^+$ of the initial valuation v_i.*

Proof. Since we focus on the valuation of a specific player i, we drop i from indexing throughout the proof. Let $\mathcal{L}$ be the set of additively separable valuations that are used to express valuation v, i.e. $\forall x \in \mathcal{X} : v(x) = \max_{\ell \in \mathcal{L}} \sum_j v_j^\ell(x_j)$. Let $\ell(x)$ be the additive valuation corresponding to outcome x, i.e.:

$$\ell(x) = \arg\max_{\ell} \sum_j v_j^\ell(x_j)$$

Consider the set of additive valuations $\hat{\mathcal{L}}$ that contains $\ell(x)$ for each $x \in \mathcal{X}$, potentially having multiple copies of the same additive valuation. The additive valuation associated with each $\ell(x) \in \hat{\mathcal{L}}$ is now defined as follows:

$$\hat{v}_j^{\ell(x)}(\tilde{x}_j) = \min\left\{ v_j^{\ell(x)}(\tilde{x}_j), v_j^{\ell(x)}(x_j), \left(B_i - \sum_{k<j} v_k^{\ell(x)}(x_k)\right)^+ \right\} \tag{7.8}$$

Consider an outcome $\tilde{x}$. We will show that

$$\hat{v}(\tilde{x}) = \max_{\ell(x) \in \hat{\mathcal{L}}} \sum_j \hat{v}_j^{\ell(x)}(\tilde{x}_j)$$

Case 1: If $v(\tilde{x}) \leq B_i$ then:

$$\hat{v}(\tilde{x}) = v(\tilde{x}) = \sum_j v_j^{\ell(\tilde{x})}(\tilde{x}_j) \leq B_i \tag{7.9}$$

The last inequality implies that for all j:

$$\sum_{k \leq j} v_k^{\ell(\tilde{x})}(\tilde{x}_k) \leq B_i$$

which in turn implies that

$$v_j^{\ell(\tilde{x})}(\tilde{x}_j) \leq B_i - \sum_{k<j} v_k^{\ell(\tilde{x})}(\tilde{x}_k)$$

Therefore, for all j:

$$v_j^{\ell(\tilde{x})}(\tilde{x}_j) = \hat{v}_j^{\ell(\tilde{x})}(\tilde{x}_j)$$

Combining with Equation (7.9) we get:

$$\hat{v}(\tilde{x}) = \sum_j \hat{v}_j^{\ell(\tilde{x})}(\tilde{x}_j) \tag{7.10}$$

Now we need to show that for all $\ell(x) \in \hat{\mathcal{L}}$: $v(\tilde{x}) \geq \sum_j \hat{v}_j^{\ell(x)}(\tilde{x}_j)$. Using the XOS property of the initial representation of the uncapped value we have:

$$\hat{v}(\tilde{x}) = v(\tilde{x}) \geq \sum_j v_j^{\ell(x)}(\tilde{x}_j) \geq \sum_j \hat{v}_j^{\ell(x)}(\tilde{x}_j) \tag{7.11}$$

By Equations (7.10) and (7.11) we get that:

$$\hat{v}(\tilde{x}) = \max_{\ell(x)\in\hat{\mathcal{L}}} \sum_j \hat{v}_j^{\ell(x)}(\tilde{x}_j)$$

Case 2: If $v(\tilde{x}) > B_i$ then $\hat{v}(\tilde{x}) = B_i$. First we observe that by the definition of the capped additive valuations $\hat{v}_j^{\ell(x)}$ we know that for any $\ell(x) \in \hat{\mathcal{L}}$:

$$\begin{aligned}
\sum_j \hat{v}_j^{\ell(x)}(\tilde{x}_j) = &\ \min\left\{ v_j^{\ell(x)}(\tilde{x}_j), v_j^{\ell(x)}(x_j), \left(B_i - \sum_{k<j} v_k^{\ell(x)}(x_k)\right)^+ \right\} \\
\leq &\ \sum_j \min\left\{ v_j^{\ell(x)}(x_j), \left(B_i - \sum_{k<j} v_k^{\ell(x)}(x_k)\right)^+ \right\} \\
\leq &\ B_i = \hat{v}(\tilde{x})
\end{aligned}$$

In addition, since $v(\tilde{x}) = \sum_j v_j^{\ell(\tilde{x})}(\tilde{x}_j) > B_i$, we know that:

$$\begin{aligned}
\sum_j \hat{v}_j^{\ell(\tilde{x})}(\tilde{x}_j) = &\ \sum_j \min\left\{ v_j^{\ell(\tilde{x})}(\tilde{x}_j), v_j^{\ell(\tilde{x})}(\tilde{x}_j), \left(B_i - \sum_{k<j} v_k^{\ell(\tilde{x})}(\tilde{x}_k)\right)^+ \right\} \\
= &\ \sum_j \min\left\{ v_j^{\ell(\tilde{x})}(\tilde{x}_j), \left(B_i - \sum_{k<j} v_k^{\ell(\tilde{x})}(\tilde{x}_k)\right)^+ \right\} \\
= &\ B_i = \hat{v}(\tilde{x})
\end{aligned}$$

By the above two sets of equations we get again that:

$$\hat{v}(\tilde{x}) = \max_{\ell(x)\in\hat{\mathcal{L}}} \sum_j \hat{v}_j^{\ell(x)}(\tilde{x}_j)$$

Thus we conclude that for any $\tilde{x}$, the above property holds and therefore $\hat{v}_j^{\ell(x)}$ for all $\ell(x) \in \hat{\mathcal{L}}$ is an XOS representation of $\hat{v}$ that uses only capped induced valuations of the initial representation of v. ∎

The above two lemmas complete the proof of the theorem. ■

Combining Theorems 7.3.1 and 7.2.2 we get robust efficiency guarantees for budget constrained bidders in the global mechanism.

7.4 Example: Simultaneous Item Auctions with Budgets

We revisit the game defined by simultaneous first or second price auctions with bidders having XOS valuations over the items. Note the valuation space for which smoothness of a first or a second price auction holds is any valuation of the form $v_i(x_i) = v_i \cdot x_i$ for $x_i \in \{0, 1\}$ and for any $v_i \geq 0$. Such a valuation space is obviously closed under capping.

Thus the results of this section imply that when the first price auction is used, even if each player has a global private budget constraint on the payments that she makes across items, every BAYES-CCE achieves at least $1 - \frac{1}{e}$ of the optimal effective welfare. If second price auctions are used then every BAYES-CCE that satisfies the no-overbidding assumption achieves $\frac{1}{2}$ of the optimal effective welfare.

8

ALGORITHMIC CHARACTERIZATIONS OF SMOOTHNESS

The definition of a smooth mechanism is a descriptive one: i.e. it states that a mechanism is smooth if each player has some "special" strategies that guarantee him a good fraction of his contribution to the optimal welfare and at a reasonable price. However, it does not imply some prescriptive definition: i.e.

What type of mechanisms are smooth for constant λ and μ?

In this chapter we address this question from an algorithmic point of view. We give conditions on the algorithm that decides the allocation profile, such that if paired with some appropriate pricing rule the resulting mechanism is smooth, for some constant λ and μ. We show a strong connection to greedy algorithms. The main result of this chapter is the following informal theorem:

Informal Theorem 5. *If the algorithm that decides the allocation profile can be viewed as a greedy optimization over a matroid feasibility constraint, then by coupling this allocation with a first-price payment scheme, the resulting mechanism is $\left(\frac{1}{3}, 1\right)$-smooth, implying a* BAYES-CCE-POA *of at most* 3.

We also provide results for the case of matroid intersections, showing that a greedy algorithm subject to a matching constraint (intersection of two partition matroids) can yield a $\left(\frac{1}{2}, 4\right)$-smooth mechanism, while the optimal algorithm can yield a $\left(\frac{1}{k+1}, 1\right)$-smooth mechanism.

8.1 Combinatorial Allocation Spaces and Greedy Mechanisms

Consider the following setting: the allocation space $\mathcal{X}_i$ of each player i consists of the power set of a finite set of elements $\mathcal{E}_i$ which we will refer to as the ground elements of player i. Moreover, we denote with $\mathcal{E} = \mathcal{E}_1 \cup \ldots \cup \mathcal{E}_n$ the set of all ground elements. The outcome of the mechanism is a subset $S \subseteq \mathcal{E}$ of this ground set and thereby the allocation of each player is $S_i = S \cup \mathcal{E}_i$.

We assume that there is some feasibility constraint $F \subseteq 2^{\mathcal{E}}$ defined on $\mathcal{E}$, which defines which subsets $S \subseteq \mathcal{E}$ are feasible. The outcome of the mechanism is restricted to fall within F. Each bidder i has a value w_t for each element $t \in \mathcal{E}_i$. The interpretation of these values is that player receives value w_t if element t is in the outcome set of the mechanism, and the overall valuation of a player is additive: $v_i(S_i) = \sum_{t \in S_i} w_t$. We will also denote with $w_i = (w_t)_{t \in \mathcal{E}_i}$ the vector of weights of elements of a player i and with $w = (w_t)_{t \in \mathcal{E}}$ a value profile on all the elements. The designer's goal is to pick the feasible set with the highest total (social) value.

Greedy Mechanism on Reported Bids. We consider the following mechanism: from each player i, the auctioneer solicits bids b_t for each $t \in \mathcal{E}_i$, i.e. $\mathcal{A}_i = \mathbb{R}_+^{|\mathcal{E}_i|}$. We will denote with $a_i = (b_t)_{t \in \mathcal{E}_i}$ an action of player i and with $b = (b_t)_{t \in \mathcal{E}}$ a bid profile on all elements. The auctioneer runs the greedy matroid algorithm on the reported bid profile to decide which elements of $\mathcal{E}$ are to be picked, i.e. elements are considered in decreasing ordered of bids and each element is added to the outcome as long as it is feasible. Each player is asked to pay his bid for each of his elements in his allocation.

8.2 Smoothness for Matroid Feasibility Constraints

We now proceed to the main result of this chapter. Consider the case where the feasibility constraint F, is the collection $\mathcal{I}$ of independent sets of a matroid $\mathrm{M} = (\mathcal{E}, \mathcal{I})$ (see Schrijver [63] for an extensive exposition of matroids), defined on the ground set. We show that the greedy mechanism on reported bids is a $\left(\frac{1}{3}, 1\right)$-smooth mechanism.

To show the smoothness property we will heavily use an exchange property of matroid feasibility constraints, proved by Lee et al. [44].

Lemma 8.2.1 (Generalized Rotta Exchange [44]). *Let $\mathrm{M} = (\mathcal{E}, I)$ be a matroid and $A, B \in \mathcal{I}$. Let $A_1, ..., A_n$ be subsets of A such that each element of A appears in exactly q of them. Then there are sets $B_1, ..., B_m \subseteq B$ such that each element of B appears in at most q of them, and for each i, $A_i \cup (B/\ B_i) \in \mathcal{I}$.*

Theorem 8.2.2. *The greedy mechanism on reported bids is a $(\frac{1}{3}, 1)$-smooth mechanism when the valuation of each player on his ground set is additive. Hence, it has* BAYES-CCE-POA *at most* 3.

Proof. Consider a valuation profile v. Suppose that each player i deviates to $a_i^* = \left(\frac{w_t}{\alpha}\right)_{t \in \mathcal{E}_i}$. Let S^* be the optimal base for valuation profile v and $S_i^* = S^* \cap \mathcal{E}_i$, be player i's allocation in the optimal base.

Consider an action profile a, where $a_i = (b_t)_{t \in \mathcal{E}_i}$, and let S, be the selected set under action profile a. Let $a' = (a_i^*, a_{-i})$, be the induced action profile and S' be the set allocated after the deviation and $S_i' = S' \cap \mathcal{E}_i$.

We denote with $W(S, a) = \sum_{t \in S} b_t$, the. By Lemma 8.2.1 for $q = 1$, we have that there exist disjoint sets $T_1, \ldots, T_n$ of S, such that $Q = S_i^* \cup (S - T_i) \in \mathcal{I}$.

By optimality of the greedy algorithm on the reported bid profile and since Q is feasible, we have:

$$\begin{aligned}\sum_{t\in S_i'}\frac{w_t}{\alpha}+\sum_{t\in S'-S_i'}b_t \geq & \sum_{t\in S_i^*}\frac{w_t}{\alpha}+\sum_{t\in S-T_i-\mathcal{E}_i}b_t+\sum_{t\in(S\cap\mathcal{E}_i)-T_i}\frac{w_t}{\alpha}\\ \geq & \sum_{t\in S_i^*}\frac{w_t}{\alpha}+\sum_{t\in S-T_i-\mathcal{E}_i}b_t\end{aligned}$$

By optimality of the algorithm on the initial bid profile we have:

$$\sum_{t\in S}b_t=W(S,a)\geq W(S',a)=\sum_{t\in S'}b_t\geq\sum_{t\in S'-S_i'}b_t$$

Combining we get:

$$\begin{aligned}\sum_{t\in S_i'}\frac{w_t}{\alpha}\geq & \sum_{t\in S_i^*}\frac{w_t}{\alpha}+\sum_{t\in S-T_i-\mathcal{E}_i}b_t-\sum_{t\in S}b_t\\ = & \sum_{t\in S_i^*}\frac{w_t}{\alpha}-\sum_{t\in S\cap(T_i\cup\mathcal{E}_i)}b_t\\ \geq & \sum_{t\in S_i^*}\frac{w_t}{\alpha}-\sum_{t\in S\cap T_i}b_t-\sum_{t\in S\cap\mathcal{E}_i}b_t\end{aligned}$$

Observe that by definition the utility of the player under the deviation is: $U_i^{\mathcal{M}}(a_i^*,a_{-i};v_i)=\left(1-\frac{1}{\alpha}\right)\sum_{t\in S_i'}w_t$. Using the previous inequalities we can lower bound his utility as follows:

$$\begin{aligned}U_i^{\mathcal{M}}(a_i^*,a_{-i};v_i)= & \left(1-\frac{1}{\alpha}\right)\sum_{t\in S_i'}w_t\\ \geq & \left(1-\frac{1}{\alpha}\right)\sum_{t\in S_i^*}w_t-\left(1-\frac{1}{\alpha}\right)\cdot\alpha\cdot\left(\sum_{t\in S\cap T_i}b_t+\sum_{t\in S\cap\mathcal{E}_i}b_t\right)\end{aligned}$$

Summing over all players:

$$\begin{aligned}\sum_i U_i^{\mathcal{M}}(a_i^*,a_{-i};v_i)\geq & \left(1-\frac{1}{\alpha}\right)\textsc{Opt}(v)-\left(1-\frac{1}{\alpha}\right)\cdot\alpha\cdot\left(\sum_i\sum_{t\in S\cap T_i}b_t+\sum_i\sum_{t\in S\cap\mathcal{E}_i}b_t\right)\\ \geq & \left(1-\frac{1}{\alpha}\right)\textsc{Opt}(v)-(\alpha-1)\cdot 2\sum_{t\in S}b_t\end{aligned}$$

where the last inequality follows, since T_i are disjoint sets and thereby $\sum_i\sum_{t\in S\cap T_i}b_t\leq\sum_{t\in S}b_t$. By setting $\alpha=\frac{1}{2}+1=\frac{3}{2}$, yields the result. ∎

By Lemma 4.2.3, we also get that the mechanism is $\left(\frac{1}{3}, 1\right)$-smooth even when players have XOS valuations over their ground elements and not additive.

Corollary 8.2.3. *The greedy mechanism on reported bids is a* $(\frac{1}{3}, 1)$*-smooth mechanism when the valuation of each player on his ground set is* XOS.

8.2.1 Action Space Restrictions and Extension to Polymatroids

We examine the generalization of the above theorem to the setting of polymatroids. Specifically, in a polymatroid setting, each element $t \in \mathcal{E}_i$ corresponds to a divisible good. Subsequently the allocation $\mathcal{X}_i = \mathbb{R}_+^{|\mathcal{E}_i|}$ of a player is the vector of allocated units from each element $t \in \mathcal{E}_i$: $x_i = (x_t)_{t \in \mathcal{E}_i}$. The mechanism decides on a vector of allocated units of each element: $x = (x_t)_{t \in \mathcal{E}} \in \mathbb{R}_+^{|\mathcal{E}|}$ and this vector has to satisfy a polymatroid constraint, i.e. for any $S \subseteq \mathcal{E}$, $\sum_{t \in S} x_t \leq f(S)$, and $f(\cdot)$ is a submodular function with $f(\emptyset) = 0$. We will assume that $f(\cdot)$ is a rational function.

The valuation of a player is linear across elements and is homogeneous for each element, i.e. $v_i(x_i) = \sum_{t \in \mathcal{E}_i} w_t \cdot x_t$. In this setting, we analyze the following mechanism:

Mechanism 3: Polymatroid mechanism.

From each player i solicit bids b_t for each $t \in \mathcal{E}_i$. Denote with $a_i = (b_t)_{t \in \mathcal{E}_i}$ and $b = (b_t)_{t \in \mathcal{E}}$

Run the greedy polymatroid algorithm with weights b to decide the final allocation x, i.e. at each iteration pick element t from remaining with maximum b_t and increase x_t until some polymatroid constraint becomes tight. Then remove t from consideration.

Charge each player i, $\sum_{t \in \mathcal{E}_i} b_t \cdot x_t$.

We show that the polymatroid mechanism is also $\left(\frac{1}{3}, 1\right)$-smooth.

Theorem 8.2.4. *The polymatroid mechanism is $\left(\frac{1}{3}, 1\right)$-smooth.*

Proof. It is well-known (see e.g. Bikhchandani et al. [10] or Schrijver [63]) that for a sufficiently small discretization of the allocation space in δ units, if we consider the extended ground set where each element t, is duplicated $\frac{f(\{t\})}{\delta}$ times, then the feasibility constraint implied by the polymatroid on these extended element set is a matroid. Moreover, if we denote with (t, k) the k-th copy of element t, then if we assign a weight of $b_t \cdot \delta$ to each element (t, k), then as the discretization goes to zero, the greedy algorithm on the matroid corresponds to the greedy algorithm on the polymatroid. Subsequently the payment of the polymatroid mechanism coincides with the payment of the extended matroid mechanism, when run on bids $b_t \cdot \delta$ for each copy of t. Last if we denote with $w_t' = w_t \cdot \delta$, then the value of a player for an allocation of discretized units, corresponds to the value of a player in the discretized matroid that has value w_t' for each copy of element t.

Thus we can view the polymatroid mechanism as the limit of a matroid mechanism where the players are restricted to submit the same bid on all copies of the same element. If we show smoothness of this restricted bid mechanism, then the smoothness of the polymatroid mechanism will follow by taking the limit of the discretization to zero.

In order to show smoothness, the only thing we need to observe is that the deviations used in the smoothness proof of the matroid mechanism (Theorem 8.2.2) are simply scaled versions of the valuation of a player. Thus it is easy to see, that if such scaled versions are allowed in the restricted bid space, then the same proof shows smoothness of the matroid mechanism under the restricted bid space. This is formalized in the following observation.

Observation 8.2.5. *Suppose that a mechanism $\mathcal{M}$ is (λ, μ)-smooth. Consider the mechanism $\mathcal{M}'$, which is identical to $\mathcal{M}$, with the exception that the action space of each player is restricted to some subset $\mathcal{A}_i' \subset \mathcal{A}_i$. If every action in the support of the deviations $\mathbf{a}_i^*(v)$ used to show smoothness of $\mathcal{M}$, fall into action space $\mathcal{A}_i'$, then $\mathcal{M}'$ is also (λ, μ)-smooth.*

Since we assumed that the value of a player is additive and homogeneous, observe that the weight w_t' of a player for each element of the discretized matroid is identical for all copies of element t. Thus the smoothness deviations of Theorem 8.2.2 would correspond to bidding $\frac{w_t'}{\alpha}$ for each element of t, which is an action that belongs to the restricted strategy space. Thus the matroid mechanism is smooth even under this restricted space, as long as the value of a player for all copies of an element is identical. Hence, the theorem follows. ■

Submodular valuations. Suppose that instead of additive and homogeneous valuations, each player has a value $v_i(x_i)$ that is monotone submodular on the euclidean lattice defined on $\mathbb{R}_+^{|\mathcal{E}_i|}$. Then by Theorem 4.3.16, we know that it can be expressed as a maximum over additively separable valuations such that each valuation is a capped marginal valuation and hence is a concave valuation on $\mathbb{R}_+$: i.e.

$$v_i(x_i) = \sup_{\ell \in \mathcal{L}} \sum_{t \in \mathcal{E}_i} v_t(x_t)$$

with $v_t(\cdot)$ being an increasing concave function. Moreover, it is easy to see that every increasing concave function can be expressed as the supremum of functions that are linear up to a point and then constant: i.e. $v_t(x_t) =$

$\sup_{\ell \in \mathcal{L}} w_t^\ell \min\{x_t, q_t^\ell\}$. Thus we can conclude that any submodular valuation can be written as:

$$v_i(x_i) = \sup_{\ell \in \mathcal{L}} \sum_{t \in \mathcal{E}_i} w_t^\ell \cdot \min\{x_t, q_t^\ell\}$$

for some index set $\mathcal{L}$. Thus in order to prove smoothness of the polymatroid mechanism, by Lemma 4.2.3 it suffices to show smoothness for the following much simpler class of valuations:

$$v_i(x_i) = \sum_{t \in \mathcal{E}_i} w_t \cdot \min\{x_t, q_t\} \tag{8.1}$$

However, we readily observe that if we consider an arbitrarily small discretization of the polymatroid then the valuations of the players for the copies of an element t, will not be identical. Instead, if we consider some arbitrary order of the copies, then the player will have a value of $w_t \cdot \delta$ for the first $\frac{q_t}{\delta}$ copies and zero value for subsequent copies. Thus to render the polymatroid mechanism smooth, we need to allow for the player to express such valuations for the copies of the same element. To achieve this we introduce the following modification of the polymatroid mechanism.

Mechanism 4: Polymatroid mechanism with capacities.

1 From each player i solicit a bid b_t and a capacity q_t for each $t \in \mathcal{E}_i$. Denote with $a_i = (b_t)_{t \in \mathcal{E}_i}$ and $b = (b_t)_{t \in \mathcal{E}}$ and $q = (q_t)_{t \in \mathcal{E}}$
2 Run the greedy polymatroid algorithm with weights b and capacities q_t to decide the final allocation x, i.e. at each iteration pick element t from remaining with maximum b_t and increase x_t until some polymatroid constraint becomes tight or x_t reaches q_t. Then remove t from consideration.
3 Charge each player i, $\sum_{t \in \mathcal{E}_i} b_t \cdot x_t$.

Under this mechanism the algorithm allows the player to submit capacities to the mechanism. By doing so, if we consider an arbitrarily small discretization of the polymatroid, then the player can submit a valuation that is $\frac{w_t \cdot \delta}{\alpha}$, for the first $\frac{q_t}{\delta}$ copies of element t and zero for the remaining, by simply submitting a weight of $\frac{w_t}{\alpha}$ and a capacity of q_t, to the polymatroid mechanism with capacities.

8.3 Smoothness for Matroid Intersections

We now turn to more complex feasibility constraints. We start by showing that the greedy algorithm leads to a smooth mechanism via swap deviations, for constant λ and μ, even when the feasibility constraint corresponds to a matching constraint: i.e. each ground element $t \in \mathcal{E}$ corresponds to an edge (u_t, v_t) in a bipartite graph $(U, V, \mathcal{E})$. This is a special case of an intersection of two partition matroids (see [63]).

Then we turn to the intersection of arbitrary k matroids and we show that running the optimal algorithm (which in general is an NP-hard problem) rather than the greedy algorithm yields a smooth mechanism, whose parameters λ and μ degrade with k. We conjecture that similar behavior holds for the greedy algorithm, but we leave it as an open question for future research.

Conjecture 8.3.1. BAYES-CCE-POA *is $O(k)$ for the greedy mechanism when the feasibility constraint is the intersection of k matroids.*

8.3.1 Matchings and Greedy Allocation

We present the smoothness theorem for matching feasibility constraints. Unlike the matroid setting where we used the existing machinery of Generalized Rota exchanges to create a charging argument, in this setting there is no analogous machinery for the greedy algorithm. Hence, we construct a charging argument that allows us to show that from a deviation either a player gets high utility or some part of the revenue at equilibrium is high. Moreover, each part of the equilibrium revenue is not charged more than twice by our charging scheme.

Theorem 8.3.2. *The greedy mechanism is $\left(\frac{1}{2}, 4\right)$-smooth via swap deviations when valuations are additive and the feasibility constraint is a matching constraint.*

Proof. Consider a valuation profile v and an action profile a with corresponding bid profile b and let S be the set output by the greedy algorithm on action profile a. Suppose that each player i deviates to bidding the pointwise maximum of a_i and $a_i' = \left(\frac{w_t}{2}\right)_{t \in S_i^*}$. This is a valid swap deviation according to definition 5.1.1. Let $a' = (a_i', a_{-i})$ be the action profile when player i deviates and b' the corresponding bid profile. Let S' be the set allocated after the deviation and let $S_i' = S' \cap \mathcal{E}_i$.

Write $A_i = S_i^* \cap S_i'$ and $U_i = S_i^* - S_i'$. (So A_i is the set of optimal items for agent i that are allocated under S', and U_i are those that are unallocated). We have:

$$U_i^{\mathcal{M}}(a_i', a_{-i}; v_i) \geq \sum_{t \in A_i} \left(w_t - \max\left\{\frac{w_t}{2}, b_t\right\}\right) \geq \frac{1}{2} \sum_{t \in A_i} w_t - \sum_{t \in A_i} b_t \geq \frac{1}{2} \sum_{t \in A_i} w_t - \sum_{t \in S_i^*} b_t.$$

So

$$\sum_i U_i^{\mathcal{M}}(a_i', a_{-i}; v_i) \geq \frac{1}{2} \sum_i \sum_{t \in A_i} w_t - \sum_i \sum_{t \in S_i^*} b_t = \frac{1}{2} \sum_i \sum_{t \in A_i} w_t - \sum_{t \in S^*} b_t.$$

Moreover, since S^* is a feasible allocation and S is the outcome of the greedy algorithm on matchings, then by the well known 2-approximation guarantee of the algorithm: $\sum_{t \in S^*} b_t \leq 2 \sum_{t \in S} b_t$. Yielding:

$$\sum_i U_i^{\mathcal{M}}(a_i', a_{-i}; v_i) \geq \frac{1}{2} \sum_i \sum_{t \in A_i} w_t - 2 \sum_{t \in S} b_t$$

We claim (proved below) that there exists a mapping $\phi\colon (\cup_i U_i) \to S$ such that $b_{\phi(t)} \geq \frac{1}{2}w_t$ for all $t \in (\cup_i U_i)$, and moreover $|\phi^{-1}(x)| \leq 2$ for all $x \in S^g$. This will imply that $2\sum_{x\in S} b_x \geq \frac{1}{2}\sum_i \sum_{t\in U_i} w_t$. We will then conclude that

$$\begin{aligned}\sum_i U_i^{\mathcal{M}}(a_i', a_{-i}; v_i) \geq & \ \frac{1}{2}\sum_i \sum_{t\in A_i} w_t - 2\sum_{x\in S} b_x + \frac{1}{2}\sum_i \sum_{t\in U_i} w_t - 2\sum_{x\in S} b_x \\ = & \ \frac{1}{2}\textsc{Opt}(v) - 4\sum_{x\in S} b_x = \frac{1}{2}\textsc{Opt}(v) - 4\mathcal{R}^{\mathcal{M}}(a),\end{aligned}$$

establishing smoothness.

Construction of charging mapping. It remains to construct the promised mapping ϕ. We first introduce the notion of an exchange graph for sets that are in the intersection of two matroids $\mathrm{M}_1 = (\mathcal{E}, I_1)$ and $\mathrm{M}_2 = (\mathcal{E}, I_2)$:

Definition 8.3.3 (Exchange Graph). *For a set $S \subset I_1 \cap I_2$, the exchange graph for S is a directed bipartite graph $G(S)$ with node sets S and $\mathcal{E} - S$ such that: for $v \in S$ and $u \in \mathcal{E}\backslash S$, we have edge (v, u) if $S - v + u \in I_1$, and edge (u, v) if $S - v + u \in I_2$.*

In a matching, a set of elements is in an independent set of matroid M_1 if no two elements have the same left endpoint in the bipartite graph $(U, V, \mathcal{E})$, while it is an independent set of matroid M_2, if no two elements have the same right endpoint. Then a feasible set of the mechanism can be viewed as the intersection of these two matroids. We now provide some extra properties of the exchange graph in the special case of a matching feasibility constraint.

Observation 8.3.4. *Given $S \in I_1 \cap I_2$ and $t \in \mathcal{E} - S$, there is at most one $s \in S$ such that $(s, t) \in G(S)$ and at most one $s' \in S$ such that $(t, s') \in G(S)$.*

Proof. Since we are in a matching setting, $(s, t) \in G(S)$ means s and t share an endpoint (say a left endpoint). Since S is a matching, t cannot share a left endpoint with multiple elements of S. Similarly, t can share a right endpoint with at most one $s' \in S$, hence can have an edge (t, s') with at most one $s' \in S$. ■

Observation 8.3.5. *Given $S \in I_1 \cap I_2$, $T \in I_1 \cap I_2$, and $s \in S$, there is at most one $t \in T - S$ such that $(s, t) \in G(S)$ and at most one $t' \in T - S$ such that $(t', s) \in G(S)$.*

Proof. Similar to the previous observation. ■

Next we argue about the structure of S', by way of $G(S)$. We remind that S is the greedy outcome under bid profile b and S' is the greedy outcome under bid profile b' produced by action profile $a' = (a'_i, a_{-i})$. We also remind that the bids in b' are the same as the bids in b with only some elements $T \in I_1 \cap I_2$ having an increased bid.

Lemma 8.3.6. *There exist paths $\pi_1, \ldots, \pi_\ell$ in $G(S)$ such that:*

1. *in each path, the bids b'_t on the nodes are either monotonically increasing or monotonically decreasing,*
2. *in each path, the node with maximum bid from b' is the unique element of $T - S$ on the path,*
3. *the paths are disjoint, except that each $t \in T$ could be the maximum-b' element for at most one increasing (in b') path and one decreasing (in b') path, and*
4. *S' is precisely S with all nodes from $S \cap (\cup_i \pi_i)$ removed and all nodes from $S - (\cup_i \pi_i)$ added.*

Proof. Let H_k be the top k elements from $\mathcal{E}$ with respect to bids B' (breaking ties in the same manner as the greedy algorithm). We will prove the stronger result that, for each k, our lemma holds for the sets $S \cap H_k$ and $S' \cap H_k$ (in item number 4). The proof will be by induction on k. Taking $k = |\mathcal{E}|$ will then give the stated lemma.

For the base case $k = 1$, let x be the single element in H_1. If $x \notin T$ then $b_x = b'_x$ and hence $S \cap H_1 = \{x\} = S' \cap H_1$ so the result holds trivially. If $x \in T$ and $x \in S$ then the result again holds trivially. If $x \in T$ but $x \notin S$, then we have $S \cap H_1 = \emptyset$ but $S' \cap H_1 = \{x\}$. In this case, take path P_1 to be the singleton node $\{x\}$ to get the desired result.

Now consider $k > 1$. By induction, there are paths $\pi_1, \ldots, \pi_\ell$ with the required properties for $S \cap H_{k-1}$ and $S' \cap H_{k-1}$. Let x be the single element in $H_k - H_{k-1}$. If x is in both S and S', or in neither, then paths $\pi_1, \ldots, \pi_\ell$ satisfy the required properties.

Suppose $x \in S$ but $x \notin S'$. Then, since $x \notin S'$, there exists some element $y \in S' - S$ such that $b'_y \geq b'_x$ and either (x, y) or (y, x) is in $G(S)$. Assume $(x, y) \in G(S)$, as the other case is symmetric. Since y is considered before x by the greedy algorithm on a', we have $y \in H_{k-1}$. So y must lie on a path π_i. From our earlier observation, there can be no $x' \in S$, $x' \neq x$, such that $(x', y) \in G(S)$. Thus, either y is the maximum-a' element of π_i, π_i is decreasing and no increasing path ending at y, or else y is the endpoint of π_i with lowest bid a'. In either case, extending π_i by appending x retains the required properties of our paths.

Finally, suppose $x \notin S$ but $x \in S'$. If $x \in T$ then create a new path containing only the singleton x, and we are done. Otherwise, there must be some element $y \in S - S'$ such that $b_y \geq b_x$ (and hence $b'_y \geq b'_x$, since $x \notin T$) and either (x, y)

or (y, x) is in $G(S)$. Assume $(x, y) \in G(S)$, as the other case is symmetric. This case is now similar to the previous case. Since y is considered before x by the greedy algorithm on a, and hence on a', we have $y \in H_{k-1}$. So y must lie on a path π_i. From our earlier observation, there can be no $x' \in S'$, $x' \neq x$, such that $(x', y) \in G(S)$ (since S' is in $I_1 \cap I_2$). Thus y is an endpoint of a path π_i. Since it cannot be the maximum-a' endpoint of the path, it is the endpoint of π_i with lowest-a' bid. In either case, extending π_i by appending x retains the required properties of our paths.

So, in all cases, the required paths exist for this value of k. The result follows by induction. ∎

Finally, we argue about properties of elements not allocated by the greedy algorithm when their bids are increased. As before, fix a bid profile b and greedy outcome S for b, and suppose bid profile b' is b with increased bids on some set of elements $T \in I_1 \cap I_2$. Let S' be the greedy outcome for b'.

Lemma 8.3.7. *Suppose $t \in T - S'$. Then t is adjacent (in $G(S)$) to some $x \in S$ such that either*

1. *$b_x \geq b'_t$, or*
2. *$x \notin S'$ and x lies on a path π_i (from the previous lemma) with a neighbor $y \notin T$ such that $b'_y \geq b'_t$, or*
3. *$x \notin S'$ and x has a neighbor $y \notin T$ in $G(S)$ such that y is the endpoint of a path π_i (from the previous lemma) and $b'_y \geq b'_t$. Moreover, the path π_i is increasing if $(x, t) \in G(S)$ and decreasing if $(t, x) \in G(S)$.*

Proof. If $t = (\alpha, \beta) \notin S'$, there must exist some $y \in S'$ with $b'_z \geq b'_t$ and y conflicting (i.e., shares a vertex with) with t (wlog let it be vertex α). If $y \in S$ then take $x = y$ and we're done. Let's assume that this is not the case and Condition 1 is not true.

Then, we have $y \in S' - S$, so our previous lemma states that z lies on a path π_i with the appropriate properties. Since y shares vertex α with t, we can take x to be the element of S that also shares this vertex. Then (x, y) and (x, t) are in $G(S)$, and moreover $x \notin S'$. If y is the endpoint of path π_i then we are done, since Condition 3 is satisfied.

Otherwise, it must be that the path π_i continues. We need to argue that the path is increasing and that (x, y) is part of the path. If the path was decreasing then it means that there exists some edge (x', y) such that $b'_{x'} \geq b'_y$. Thus x' shares vertex α with y and hence with t. Moreover, for this reason $x' \notin T$ and therefore $b'_x = b_x$. Thus condition 1 is satisfied, with $x = x'$, a contradiction to our first assumption. So it must be that the path is increasing and hence an edge (r, y) is part of the path. But from Observation 8.3.4 the only such r is x. Thus Condition 2 is satisfied. ■

We're now ready to define our promised mapping $\phi : \cup_i U_i \to S$. Take $T = S_i^*$ in the above Lemmas. Note $U_i = T - S'$. For each $t \in T - S'$, if Condition 1 of Lemma 8.3.7 is satisfied then take $\phi(t) = x$, for the x in the condition.

If Conditions 2 or 3 of Lemma 8.3.7 are satisfied then we first create a temporary association, based on which we subsequently construct the mapping. If Condition 2 of Lemma 8.3.7 is satisfied then we associate t with the x in the condition, whilst if Condition 3 is satisfied then we associate t with the endpoint y. Then each $x \in S$ is associated with only one t, since each such x lies on a unique path, and the neighbor from T with which it's associated is determined by the direction of that path. Moreover, each endpoint of a path y is associated with only one other note t, since y lies on a unique path and the node t with which it is associated is determined by the monotonicity of the path.

Now define $\phi(t)$ as follows: starting from the node x associated with t, follow the path in the direction of increasing b' until reaching some $x' \in S$ that is either (a) associated with some other node $t' \in T$, or (b) the last element of S along the path. In either case, we will set $\phi(t) = x'$. Note that $b'_{x'} \geq b'_t$, since by construction $b'_{x'} \geq b'_y \geq b'_t$. Moreover, $x' \notin T$ by the observation that each such x' is adjacent with some $x \in T$ in graph $G(S)$ and thereby is conflicting with some $x \in T$. Thus $b_x = b'_x \geq b'_t$ as required.

We must argue that this mapping ϕ satisfies the required properties. We already have that $b_{\phi(t)} \geq b'_t \geq \frac{w_t}{2}$ for each t. We next argue that $|\phi^{-1}(x)| \leq 2$ for each $x \in S$. This follows because each x is mapped-to at most once for each of its (two) adjacent elements $t \in S^*$. If b_x is greater than b'_t, then x is mapped-to directly from t. If b_x is less than b'_t then x can potentially be mapped to via an association with t along the (at most one) path containing x, but only by one other element t' (i.e., corresponding to the next-lowest element along that path that has an association).

The mapping ϕ therefore satisfies the required properties, completing the proof of Theorem 8.3.2. ■

8.3.2 Intersections of Matroids and Optimal Algorithm

We now analyze the mechanism where instead of running the greedy algorithm over the reported bid profile, we run the optimal algorithm so as to decide the outcome set. Then each player is charged his bid for his allocated set. We refer to this mechanism as the optimal mechanism with first prices. We state the theorem for additive valuations, but it is easy to observe that by Lemma 4.2.3 the theorem extends to XOS valuations on the elements.

Theorem 8.3.8. *The optimal mechanism with first prices is a $\left(\frac{1}{k+2}, 1\right)$-smooth mechanism when valuations are additive and the feasibility constraint on the ground set is the intersection of k matroids. Thus the* BAYES-CCE-POA *is at most $k+2$.*

Proof. Consider a valuation profile v. Suppose that each player i deviates to $a_i^* = \left(\frac{w_t}{\alpha}\right)_{t\in\mathcal{E}_i}$. Let S^* be the optimal base for valuation profile v and $S_i^* = S^* \cap \mathcal{E}_i$, be player i's allocation in the optimal base.

Consider an action profile a, where $a_i = (b_t)_{t\in\mathcal{E}_i}$, and let S, be the selected set under action profile a. Let $a' = (a_i^*, a_{-i})$, be the induced action profile and S' be the set allocated after the deviation and $S_i' = S' \cap \mathcal{E}_i$.

Suppose that the feasibility constraint on the elements of the ground set is the intersection of k matroid constraints, $\mathrm{M}_1, \ldots, \mathrm{M}_k$. By Lemma 8.2.1 applied to every matroid $\mathrm{M}_t = (\mathcal{E}, \mathcal{I}^t)$, we have that there exist disjoint sets $T_1^t, \ldots, T_n^t$ such that $S_i^* \cup (S - T_i^t) \in \mathcal{I}^t$. Thus it is easy to see that: $Q = S_i^* \cup \left(S - \cup_{t=1}^k T_i^t\right) \in \cap_{t=1}^k \mathcal{I}^t$ is a feasible set. Let $T_i = \cup_{t=1}^k T_i^t$, observe that since for each $t \in [k]$, $T_1^t, \ldots, T_n^t$, are disjoint sets, an element appears in at most k of the sets $T_1, \ldots, T_n$.

The rest of the proof follows along similar lines as in theorem 8.2.2. By the optimality of the algorithm on the reported bid profile and since $Q = S_i^* \cup (S - T_i)$ is feasible, we have:

$$\begin{aligned}
\sum_{t\in S_i'} \frac{w_t}{\alpha} + \sum_{t\in S'-S_i'} b_t \geq & \sum_{t\in S_i^*} \frac{w_t}{\alpha} + \sum_{t\in S-T_i-\mathcal{E}_i} b_t + \sum_{t\in (S\cap \mathcal{E}_i)-T_i} \frac{w_t}{\alpha} \\
\geq & \sum_{t\in S_i^*} \frac{w_t}{\alpha} + \sum_{t\in S-T_i-\mathcal{E}_i} b_t
\end{aligned}$$

By optimality of the algorithm on the initial bid profile we have:

$$\sum_{t\in S} b_t = W(S,a) \geq W(S',a) = \sum_{t\in S'} b_t \geq \sum_{t\in S'-S_i'} b_t$$

Combining we get:

$$\begin{aligned}
\sum_{t\in S_i'} \frac{w_t}{\alpha} \geq & \sum_{t\in S_i^*} \frac{w_t}{\alpha} + \sum_{t\in S-T_i-\mathcal{E}_i} b_t - \sum_{t\in S} b_t \\
= & \sum_{t\in S_i^*} \frac{w_t}{\alpha} - \sum_{t\in S\cap (T_i \cup \mathcal{E}_i)} b_t \\
\geq & \sum_{t\in S_i^*} \frac{w_t}{\alpha} - \sum_{t\in S\cap T_i} b_t - \sum_{t\in S\cap \mathcal{E}_i} b_t
\end{aligned}$$

Observe that by definition the utility of the player under the deviation is: $U_i^{\mathcal{M}}(a_i^*, a_{-i}; v_i) = \left(1 - \frac{1}{\alpha}\right) \sum_{t\in S_i'} w_t$. Using the previous inequalities we can lower bound his utility as follows:

$$\begin{aligned}
U_i^{\mathcal{M}}(a_i^*, a_{-i}; v_i) = & \left(1 - \frac{1}{\alpha}\right) \sum_{t\in S_i'} w_t \\
\geq & \left(1 - \frac{1}{\alpha}\right) \sum_{t\in S_i^*} w_t - \left(1 - \frac{1}{\alpha}\right) \cdot \alpha \cdot \left(\sum_{t\in S\cap T_i} b_t + \sum_{t\in S\cap \mathcal{E}_i} b_t \right)
\end{aligned}$$

Summing over all players:

$$\begin{aligned}
\sum_i U_i^{\mathcal{M}}(a_i^*, a_{-i}; v_i) \geq & \left(1 - \frac{1}{\alpha}\right) \textsc{Opt}(v) - \left(1 - \frac{1}{\alpha}\right) \cdot \alpha \cdot \left(\sum_i \sum_{t\in S\cap T_i} b_t + \sum_i \sum_{t\in S\cap \mathcal{E}_i} b_t \right) \\
\geq & \left(1 - \frac{1}{\alpha}\right) \textsc{Opt}(v) - (\alpha - 1) \cdot (k+1) \sum_{t\in S} b_t
\end{aligned}$$

where the last inequality follows, since an element $t \in \mathcal{E}$ appears in at most k of the sets $T_1, \ldots, T_n$ and thereby $\sum_i \sum_{t \in S \cap T_i} b_t \leq k \sum_{t \in S} b_t$. By setting $\alpha = \frac{1}{k+1} + 1 = \frac{k+2}{k+1}$, yields the result. ■

9

SMOOTHNESS IN LARGE MARKETS

... while there may be asymmetry in agents' information, it is relatively unimportant for the problem at hand because any single agent has only a small amount of information not known by the other agents.

– Gul and Postlewaite, 1992, p. 1273

As a market grows larger in size, the effect of each individual player is negligible. Thereby it is natural to expect that the effects of strategic manipulations by each individual player will be alleviated and the worst-case inefficiency of equilibrium outcomes will improve.

In this chapter we propose a formalization of the above intuition via the smoothness framework. We define the notion of smoothness in the limit: *a mechanism might not be (λ, μ)-smooth for any finite size of the market but becomes (λ, μ)-smooth in the limit of an infinitely large market.*

The limit smoothness framework also provides worst-case efficiency guarantees for any finite market size, i.e. if a mechanism is (λ, μ)-smooth in the limit, then for any ϵ, there exists a finite market size, such that the mechanism is $\left(\frac{\lambda}{1+\epsilon}, \mu\right)$-smooth and therefore for sufficiently large markets the BAYES-CCE-POA is at most $(1+\epsilon)\frac{\max\{1,\mu\}}{\lambda}$.

Via the limit smoothness framework we show a very general full efficiency result for simultaneous uniform price auctions in a market with many goods and combinatorial valuations. Specifically, we show that if the number of players and the number of units of each good grows and each player doesn't arrive in the market with some small negligible probability and has finite demand from each good, then the resulting mechanism is $(1,1)$-smooth in the limit, and therefore any Bayes-CCE converges to a fully efficient allocation. Our results heavily generalize on previous work on efficiency in large markets [64], that only considers a single good and decreasing marginal valuations, whilst not providing any approximation guarantee for finite market sizes.

9.1 Smoothness in the Limit

To define the notion of smoothness in the limit, we need to consider a sequence of mechanisms $\mathcal{M}^n = (\mathcal{A}^n, X^n, P^n)$ defined on a sequence of mechanism design settings $(n, \mathcal{X}^n, \mathcal{V}^n)$ and a sequence of independent valuation distributions $\mathcal{F}^n$. Normally this would be the same mechanism adapted to a growing number of players or a growing number of resources. For instance it could be a single item first price auction where only the number of players n is changing. However, the number of items can also be growing as a function of the number of players n. All these variances are captured in the above general formulation of a sequence of mechanisms for each market size n.

For shorter notation we will use U_i^n and $\mathcal{R}^n$ instead of $U_i^{\mathcal{M}^n}$ and $\mathcal{R}^{\mathcal{M}^n}$ for the utility and the revenue under mechanism $\mathcal{M}^n$.

Definition 9.1.1 (Smooth Mechanism in the Limit). *A sequence of mechanisms $\mathcal{M}^n$ is (λ, μ)-smooth in the limit if for each market size n and for each valuation profile v^n, there exists a deviation sequence $a_i^{*,n}(v^n)$ for each $i \in [n]$, such that for any action profile sequence $a^n \in \mathcal{A}^n$:*

$$\limsup_{n\to\infty} \frac{\lambda \textsc{Opt}^n(v^n)}{\sum_i U_i^n(a_i^{*,n}, a_{-i}^n; v_i) + \mu \mathcal{R}^n(a^n)} \leq 1 \tag{9.1}$$

If a sequence of mechanisms is (λ, μ)-smooth in the limit, then the price of anarchy must converge to the limit price of anarchy as is formalized below. Moreover, it implies that for any sufficiently large but finite market size the price of anarchy of all BAYES-CCE is at most $1+\epsilon$ away from the limit price of anarchy.

Theorem 9.1.2. *If a mechanism is (λ, μ)-smooth in the limit then*

$$\limsup_{n\to\infty} \text{BAYES-CCE-POA}^n \leq \frac{\max\{\mu, 1\}}{\lambda},$$

i.e. for any ϵ there exists a market size $n(\epsilon)$ such that for any $n \geq n(\epsilon)$, every Bayes coarse correlated equilibrium of the mechanism $\mathcal{M}^n$ with value distributions $\mathcal{F}^n$ achieves at least $(1+\epsilon)\frac{\lambda}{\max\{1,\mu\}}$ of the expected optimal welfare.

Proof. By (λ, μ)-smoothness in the limit, we have that for any ϵ there exists a market size $n(\epsilon)$ such that for any $n \geq n(\epsilon)$:

$$\frac{\lambda \text{OPT}^n(v^n)}{\sum_i U_i^n(a_i^{*,n}, a_{-i}^n; v_i) + \mu \mathcal{R}^n(a^n)} \leq 1+\epsilon$$

Since, the action sequence a^n is arbitrary, the above holds for any action profile a^n. Thus by rearranging, we get that for any such n, mechanism $\mathcal{M}^n$ is a $\left(\frac{\lambda}{1+\epsilon}, \mu\right)$-smooth mechanism. Therefore by Theorem 3.1.2, BAYES-CCE-POA$^n \leq (1+\epsilon)\frac{\max\{1,\mu\}}{\lambda}$.

■

9.2 Simultaneous Uniform Price Auctions with Noisy Arrival

Consider a setting with n bidders and m goods. There are $k_j(n)$ units of good $j \in [m]$. Each player $i \in [n]$ has a value $v_i : \mathbb{N}^m \to [\rho, H]$, that assigns a value for each possible allocation of units of each good. We will assume that the values are bounded away from zero and are bounded. We will also assume that the players have value for up to some number of r units of each good, and r remains constant independent of n as the market grows. More formally, the value of a player satisfies the following condition: for any allocation vector $x_i = (x_{i1}, \ldots, x_{im})$, where x_{ij} denotes the units of good j allocated:

$$v_i(x_i) = v_i(\min\{x_i, r\}), \tag{9.2}$$

where $\min\{x_i, r\}$ is the coordinate-wise minimum of x_{ij} and r. We will refer to such functions as r-demand valuations.

Simultaneous Uniform Price. The units of each good $j \in [m]$ are simultaneously and independently sold via the means of a uniform price auction: bidders submit bids $b_{i,1}^j \geq \ldots \geq b_{i,r}^j$. Then the bids of good j are ordered in decreasing order and the first $k_j(n)$ bids, each wins a unit. Thus for a player to win x_j units of good j it has to be that his highest x_j bids are in the top $k_j(n))$ bids among all players. Then each player is charged the highest losing bid for each unit of good j that he is allocated, i.e. if we denote with $\theta_t^j(b^j)$ the t-th highest bid after we order the bids of all players, then each player pays $\theta_{k_j(n)+1}^j(b^j)$ for each unit.

We will denote the above mechanism with $\mathcal{M}^n = (\mathcal{A}^n, X^n, P^n)$ and the above mechanism design setting with $(n, \mathcal{X}^n, \mathcal{V}^n)$.

Random Arrival. We will also assume that each player doesn't arrive in the market with probability δ. In that case we can think of the player as submitting an all zero bid vector on all the auctions and having zero valuation for any allocated units. Moreover, when a player decides his bid he does not know which player arrived. We will even assume (wlog) that he doesn't know whether he will arrive too when deciding his bids and in the end he submits his predetermined bids conditional on arriving.

Endogenous Noise Mechanism. We will view the simultaneous uniform price auction with random arrivals as an ex-ante mechanism $\mathcal{M}_\delta^n$, where the noise is endogenized in the rules of the mechanism and then we will show that mechanism $\mathcal{M}_\delta^n$ is $(1,1)$-smooth in the limit. We will refer to this mechanism as *simultaneous uniform price auction with endogenous δ-noisy demand*.

Let $\mathbf{z}_i$ be an indicator random variable that designates whether a player arrived in the auction, i.e. $\mathbf{z}_i$ is an independent Bernoulli trial with success probability $1-\delta$. The action space of mechanism $\mathcal{M}_\delta^n$ is the same as mechanism $\mathcal{M}$, i.e. $\mathcal{A}_i^n \triangleq \mathbb{R}_+^{m\cdot r}$. The allocation space of each player is the space of functions from an arrival vector z to an allocation of the simultaneous uniform price auction, i.e. $\mathcal{X}_i^{\delta,n} \triangleq \{0,1\}^n \rightarrow \mathcal{X}_i^n$. Let $x_i^n(b,z) = X_i^n(b\cdot z)$, where $b\cdot z = (b_1\cdot z_1,\ldots,b_n\cdot z_n)$. Then the allocation and payment function of mechanism $\mathcal{M}^{\delta,n}$ is:

$$X_i^{\delta,n}(b) = x_i^n(b,\cdot) \tag{9.3}$$

$$P_i^{\delta,n}(b) = \mathbb{E}_{\mathbf{z}}\left[P_i^n(b\cdot\mathbf{z})\right] \tag{9.4}$$

The value of a player for an allocation $x_i^\delta \in \mathcal{X}_i^{\delta,n}$ is the expected value under the random arrival vector:

$$v_i^\delta(x_i^\delta) = \mathbb{E}_{\mathbf{z}}\left[\mathbf{z}_i \cdot v_i\left(x_i^\delta(\mathbf{z})\right)\right] \tag{9.5}$$

The utility of each player from the mechanism $\mathcal{M}_\delta^n$, denoted $U_i^{\delta,n}$ for conciseness, under some action profile b is:

$$U_i^{\delta,n}(b;v_i) = v_i^\delta(x_i^\delta) - P_i^{\delta,n}(b) = \mathbb{E}_{\mathbf{z}}\left[\mathbf{z}_i\cdot v_i\left(X_i^n(b\cdot\mathbf{z})\right) - P_i^n(b\cdot\mathbf{z})\right] \tag{9.6}$$

Under the above formulation, the optimal allocation of the ex-ante mechanism is the expected ex-post optimal allocation over the random arrivals:

$$\textsc{Opt}^{\delta,n}(v^n) = \mathbb{E}_z\left[\textsc{Opt}^n(v^n \cdot z)\right],$$

where $v \cdot z = (v_1 \cdot z_1, \ldots, v_n \cdot z_n)$.

Thus mechanism $\mathcal{M}_\delta^n$ falls into our framework and in the next section we will show that it is $(1,1)$-smooth in the limit.

9.2.1 Full Efficiency in the Limit

Theorem 9.2.1. *Simultaneous uniform price auctions with endogenous δ-noisy demand are $(1,1)$-smooth in the limit when bidders have monotone r-demand combinatorial valuation over goods, such that $v_i(\cdot) \geq \rho > 0$, for some ρ and when the supply of each good $k_j(n)$ goes to infinity as the market grows.*

We will prove the theorem in a sequence of three Lemmas. First we show that for any market size n, mechanism $\mathcal{M}_\delta^n$ satisfies an almost $(1,1)$-smoothness property where instead of revenue on the right hand side of the smoothness definition we have replaced it with a sum of *threshold bids*. Then in the second Lemma we use the results of Swinkels [64] to show that these threshold bids will converge to the revenue of the auction in the limit. Together with our minor technical assumptions on bidder valuations, this will allow us to conclude the theorem.

Lemma 9.2.2. *For a fixed market size n, each player i, for any monotone valuation v_i and for any allocation $x_i \in [1,..,r]^m$, player i has a deviation $b_i'(x_i)$ such that for any* *bid profile b^n:*

$$U_i^{\delta,n}(b_i'(x_i), b_{-i}^n; v_i) \geq \mathbb{E}_z \left[z_i \cdot \left(v_i(x_i) - \sum_{j\in[m]} x_{ij} \cdot \theta^j_{k_j(n)-r+1}(b_{-i}^{j,n} \cdot z_{-i}) \right) \right] \tag{9.7}$$

Proof. Since the market size is fixed, we will drop the index on n. Thus we have k_j units of each good and the bid profile b.

Consider the deviation $b_i'(x_i)$ where player i deviates to bidding $v_i(x_i)$ on the first x_{ij} marginal bids on each uniform price auction $j \in [m]$.

Let $\theta_{j,t}(b \cdot z)$ be the t-th highest arriving bid at uniform price auction j. We will show that the utility of a player under the deviation, conditional on any instance of z is at least: $z_i \cdot v_i(x_i) - \sum_{j\in[m]} x_{ij} \cdot \theta^j_{k_j-r+1}(b_{-i}^j \cdot z_{-i})$. The theorem then follows by taking expectation over z. If $z_i = 0$, it follows trivially, since the utility is zero.

If $z_i = 1$, observe that a player wins x_i if for all $j \in [m] : \theta_{j,k_j-x_{ij}+1}(b^j \cdot z) < v_i(x_i)$. In that case, he gets utility at least: $z_i \cdot v_i(x_i) - \sum_{j\in[m]} x_{ij} \cdot \theta_{j,k_j-x_{ij}+1}(b_{-i}^j \cdot z_{-i})$. If a player doesn't win allocation x_i, then it must be that there exists some item $q \in [m]$, such that $\theta_{q,k_q-x_{iq}+1}(b_{-i}^j \cdot z_{-i}) \geq v_i(x_i)$ and such that player i is winning at most $x_{iq} - 1$ units of item q. Thus player i's utility from the deviation in this case is at least:

$$-\sum_{j\in[m]} x_{ij} \cdot \theta_{j,k_j-x_{ij}+1}(b_{-i}^j \cdot z_{-i}) + \theta_{j,k_q-x_{iq}+1}(b_{-i}^j \cdot z_{-i}) \geq v_i(x_i) - \sum_{j\in[m]} x_{ij} \cdot \theta_{j,k_j-x_{ij}+1}(b_{-i}^j \cdot z_{-i})$$

Thus in any case we get that the utility conditional on z is lower bounded by the desired amount. Since $x_{ij} \leq r$, the theorem follows. ∎

Lemma 9.2.3. *For a fixed market size n, for any valuation profile v^n satisfying our assumptions, there exists deviations $b_i^{*,n}(v^n)$ for each player i such that for any bid* *profile b^n:*

$$\sum_i U_i^{\delta,n}(b_i^{*,n}(v^n), b_{-i}^n; v_i) \geq \textsc{Opt}^{\delta,n}(v^n) - \sum_{j\in[m]} k_j(n)\cdot \mathbb{E}_z\left[\theta^j_{k_j(n)-r+1}(b^{j,n}\cdot z)\right] \tag{9.8}$$

Proof. Since for this lemma, we have fixed a market size n, we will drop it from any indexing, i.e. k_j is the available units of each item j, b is a bid profile, v is a valuation profile. Since the valuation profile v is fixed, we will denote with $x_i^*(z)$ the optimal allocation of player i under valuation profile v and arrival profile z.

Consider the following deviation b_i^*: random sample an arrival profile $\tilde{\mathbf{z}}_{-i}$. Then $b_i^*(v) = b_i'(x_i^*(1,\tilde{\mathbf{z}}_{-i}))$ designated by Lemma 9.2.2. By Lemma 9.2.2 for any bid profile b:

$$\begin{aligned}
U_i^\delta(b_i^*(v), b_{-i}; v_i) \geq\ & \mathbb{E}_{\mathbf{z},\tilde{\mathbf{z}}}\left[\mathbf{z}_i\cdot v_i(x_i^*(1,\tilde{\mathbf{z}}_{-i})) - \mathbf{z}_i\sum_{j\in[m]} x_{ij}^*(1,\tilde{\mathbf{z}}_{-i})\cdot\theta^j_{k_j-r+1}(b^j_{-i}\cdot\mathbf{z}_{-i})\right]\\
=\ & \mathbb{E}_{\mathbf{z},\tilde{\mathbf{z}}}\left[\mathbf{z}_i\cdot v_i(x_i^*(\mathbf{z}_i,\tilde{\mathbf{z}}_{-i})) - \mathbf{z}_i\sum_{j\in[m]} x_{ij}^*(\mathbf{z}_i,\tilde{\mathbf{z}}_{-i})\cdot\theta^j_{k_j-r+1}(b^j_{-i}\cdot\mathbf{z}_{-i})\right]\\
=\ & \mathbb{E}_{\mathbf{z},\tilde{\mathbf{z}}}\left[\tilde{\mathbf{z}}_i\cdot v_i(x_i^*(\tilde{\mathbf{z}})) - \tilde{\mathbf{z}}_i\sum_{j\in[m]} x_{ij}^*(\tilde{\mathbf{z}})\cdot\theta^j_{k_j-r+1}(b^j_{-i}\cdot\mathbf{z}_{-i})\right]\\
\geq\ & \mathbb{E}_{\mathbf{z},\tilde{\mathbf{z}}}\left[\tilde{\mathbf{z}}_i\cdot v_i(x_i^*(\tilde{\mathbf{z}})) - \sum_{j\in[m]} x_{ij}^*(\tilde{\mathbf{z}})\cdot\theta^j_{k_j-r+1}(b^j\cdot\mathbf{z})\right]
\end{aligned}$$

Summing over all players:

$$\begin{aligned}
\sum_i U_i^\delta(b_i^*(v), b_{-i}; v_i) \geq\ & \mathbb{E}_{\mathbf{z},\tilde{\mathbf{z}}}\left[\sum_i \tilde{\mathbf{z}}_i\cdot v_i(x_i^*(\tilde{\mathbf{z}})) - \sum_{j\in[m]}\theta^j_{k_j-r+1}(b^j\cdot\mathbf{z})\sum_i x_{ij}^*(\tilde{\mathbf{z}})\right]\\
\geq\ & \mathbb{E}_{\mathbf{z},\tilde{\mathbf{z}}}\left[\sum_i \tilde{\mathbf{z}}_i\cdot v_i(x_i^*(\tilde{\mathbf{z}})) - \sum_{j\in[m]} k_j\cdot\theta^j_{k_j-r+1}(b^j\cdot\mathbf{z})\right]\\
=\ & \textsc{Opt}^\delta(v) - \sum_{j\in[m]} k_j\cdot\mathbb{E}_{\mathbf{z}}\left[\theta^j_{k_j-r+1}(b^j\cdot\mathbf{z})\right]
\end{aligned}$$

■

At this point we are ready to use a re-interpretation of the results of Swinkels [64], whose relevant to our analysis conclusion we present here:

Lemma 9.2.4 (Swinkels [64])**.** *For any sequence of bid profiles b^n and for any constant r, if $k_j(n) \to \infty$ then:*

$$\lim_{n\to\infty} \mathbb{E}_{\mathbf{z}}\left[\theta^j_{k_j(n)-r+1}(b^{j,n}\cdot \mathbf{z}))\right] - \mathbb{E}_{\mathbf{z}}\left[\theta^j_{k_j(n)+1}(b^{j,n}\cdot \mathbf{z}))\right] = 0 \tag{9.9}$$

Proof of Theorem 9.2.1 : By Lemma 9.2.4 for any sequence of bid profiles b^n and for any ϵ there exists $n(\epsilon)$ such that for any $n \geq n(\epsilon)$:

$$\mathbb{E}_{\mathbf{z}}\left[\theta^j_{k_j(n)-r+1}(b^{j,n}\cdot \mathbf{z}))\right] \leq \mathbb{E}_{\mathbf{z}}\left[\theta^j_{k_j(n)+1}(b^{j,n}\cdot \mathbf{z}))\right] + \epsilon$$

and therefore by Lemma 9.2.3:

$$\sum_i U_i^{\delta,n}(b_i^{*,n}(v^n), b^n_{-i}; v_i) \geq \\ \textsc{Opt}^{\delta,n}(v^n) - \sum_{j\in[m]} k_j(n)\cdot \mathbb{E}_{\mathbf{z}}\left[\theta^j_{k_j(n)+1}(b^{j,n}\cdot \mathbf{z})\right] - \sum_{j\in[m]} k_j(n)\cdot\epsilon \tag{9.10}$$

Since for any player i, $v_i(\cdot) \geq \rho$, we have that $\textsc{Opt}^{\delta,n}(v^n) \geq (1-\delta)\frac{\rho}{r\cdot m}\sum_{j\in[m]} k_j(n)$, by the existence of the following allocation: order arriving players arbitrarily at each good and sequentially allocate one unit of the good to next player in the order (wrapping around if necessary) until the units run out or all players are satisfied. If player i got x_{ij} units of good j, then the welfare of the allocation is: $\sum_{i\in[n]} z_i\cdot v_i(x) \geq \sum_{i\in[n]} \frac{\rho}{r\cdot m}\sum_{j\in[m]} z_i\cdot x_{ij} = \frac{\rho}{r\cdot m}\sum_{j\in[m]}\sum_{i\in[n]} z_i\cdot x_{ij}$. Observe that $\sum_{j\in[n]} z_i\cdot x_{ij} = \min\{r\cdot\sum_{i\in[n]} z_i, k_j(n)\}$. In expectation over $\mathbf{z}$, this quantity is at least the minimum of the two expectations. Since, $\mathbb{E}_{\mathbf{z}}\left[\sum_{i\in[n]}\mathbf{z}_i\right] = (1-\delta)\cdot n$ and since $r\cdot n \geq k_j(n)$, we have that the allocation achieves welfare at least: $\frac{\rho}{r\cdot m}\sum_{j\in[m]}(1-\delta)k_j(n)$.

Hence, we can transform the additive error in Equation (9.10) into a multiplicative:

$$\sum_i U_i^{\delta,n}(b_i^{*,n}(v^n), b_{-i}^n; v_i) + \sum_{j\in[m]} k_j(n) \cdot \mathbb{E}_{\mathbf{z}}\left[\theta^j_{k_j(n)+1}(b^{j,n}\cdot \mathbf{z})\right]$$
$$\geq \left(1 - \frac{r\cdot m}{\rho(1-\delta)}\epsilon\right)\textsc{Opt}^{\delta,n}(v^n)$$

By observing that:

$$\begin{aligned}\mathcal{R}^{\mathcal{M}_\delta^n}(b^n) = \sum_i P_i^{\delta,n}(b^n) = &\sum_{j\in[m]} \mathbb{E}_{\mathbf{z}}\left[\theta^j_{k_j(n)+1}(b^n\cdot z)\cdot \sum_i x_{i,j}^n(b^{j,n},\mathbf{z})\right]\\ = &\sum_{j\in[m]} \mathbb{E}_{\mathbf{z}}\left[k_j(n)\cdot\theta^j_{k_j(n)+1}(b^{j,n}\cdot\mathbf{z}))\right]\end{aligned}$$

where the last inequality follows by the fact that if $\sum_i x_{i,j}^n(b^{j,n}, z) < k_j(n)$, then it must be that $\theta^j_{k_j(n)+1}(b^{j,n}\cdot z) = 0$, and otherwise $\sum_i x_{i,j}^n(b^{j,n}, z) = k_j(n)$. Combining with the multiplicative error equation we get:

$$\frac{\textsc{Opt}^{\delta,n}(v^n)}{\sum_i U_i^{\delta,n}(b_i^{*,n}(v^n), b_{-i}^n; v_i) + \mathcal{R}^{\mathcal{M}_\delta^n}(b^n)} \leq \frac{\rho\cdot(1-\delta)}{\rho\cdot(1-\delta)-\epsilon\cdot r\cdot m} \tag{9.11}$$

For ϵ small enough, the above ratio can be made smaller than any $1+\epsilon'$, which yields the theorem. ■

Part III

Applications

10
SINGLE ITEM AUCTIONS

In this section we analyze the smoothness of single-item auctions. Throughout the thesis we presented extensive results for first and second price auctions. Here we analyze two other standard auction formats: the all-pay auction and the hybrid auction where the winner pays a mixture of his bid and the second highest bid and give a complete list of our results.

10.1 All-Pay Auction

Simultaneous and sequential all-pay auctions have not been studied in the literature and could prove useful in capturing simultaneous or sequential all-pay contests, which is a natural model for several online crowd-sourcing environments.

Lemma 10.1.1. *The all-pay auction is a* $(1/2, 1)$*-smooth mechanism via conservative deviations.*

Proof. To see why smoothness holds, note that under any valuation profile $v = (v_1, \ldots, v_n)$, the highest value player (wlog player 1) can deviate to submitting a randomized bid b_1^* drawn drawn uniformly at random from $[0, v_1]$, while all non-highest value players should just deviate to bidding 0. No matter what the rest of the players are bidding, the utility of the highest bidder from the deviation is:

$$U_1^{\mathrm{APA}}(\mathbf{b}_1^*, b_{-1}; v_1) \geq \int_{\max_{i\neq 1} b_i}^{v_1} v_1 f(x)dx - \mathbb{E}[\mathbf{b}_1^*]$$
$$\geq \frac{1}{2}v_1 - \max_i b_i \geq \frac{1}{2}v_1 - \sum_i b_i = \frac{1}{2}\textsc{Opt}(v) - \mathcal{R}^{\mathcal{M}}(b)$$

■

Therefore we get an efficiency guarantee of $1/2$ for the simultaneous composition of m all-pay auctions and an efficiency guarantee of $1/4$ for the sequential composition.

Corollary 10.1.2 (Simultaneous with Budgets)**.** *If we run m simultaneous all-pay auctions and bidders have budgets and fractionally subadditive valuations then the expected effective welfare at every* BAYES-CCE *is at least $1/2$ of the expected optimal effective welfare.*

Corollary 10.1.3 (Sequential)**.** *If we run m sequential all-pay auctions with unit-demand bidders then* BAYES-CE-POA ≤ 4.

In Appendices A.4 and A.5, we present almost matching lower bounds on the inefficiency of the all-pay auction.

10.2 Hybrid Auction

In the γ-hybrid auction the winner pays his bid with probability γ and the second highest bid with probability $(1-\gamma)$.

Lemma 10.2.1. *The γ-hybrid auction is weakly $(\gamma(1-\frac{1}{e})+(1-\gamma)^2, 1, (1-\gamma)^2)$-smooth.*

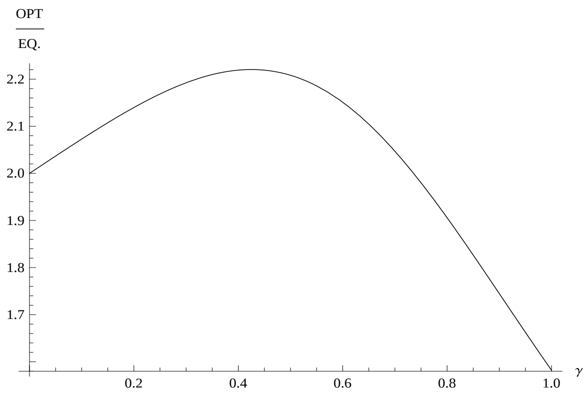

Figure 10.1: Efficiency bound for the hybrid auction.

Proof. Consider a valuation profile v and a bid profile b. Let $b_{\max}$ the highest bid and $v_{\max}$ the highest value. The non highest value bidders deviation is bidding 0. The highest value bidder's deviation is bidding as follows: With probability γ he submits a bid according to distribution with density $f(t) = \frac{1}{v_{\max}-t}$ and support $[0, (1-\frac{1}{e})v_{\max}]$. With probability $(1-\gamma)$ he submits his true value.

In the first case the utility of the bidder is at least:

$$\int_{b_{\max}}^{(1-\frac{1}{e})v_{\max}} (v_{\max} - \gamma t - (1-\gamma)b_{\max}) f(t)dt \geq \int_{b_{\max}}^{(1-\frac{1}{e})v_{\max}} (v_{\max} - t) f(t)dt = \left(1 - \frac{1}{e}\right) v_{\max} - b_{\max}$$

In the case when he submits his true value then when $b_{\max} < v_{\max}$ he wins and gets utility

$$v_{\max} - \gamma v_{\max} - (1-\gamma)b_{\max} = (1-\gamma)(v_{\max} - b_{\max})$$

When $b_{\max} \geq v_{\max}$ he either loses or ties and in any case gets non-negative utility and thereby utility at least $(1-\gamma)(v_{\max} - b_{\max})$.

Thus overall the expected utility from the deviation is at least:

$$\gamma\left(1-\frac{1}{e}\right)v_{\max}-\gamma b_{\max}+(1-\gamma)^2(v_{\max}-b_{\max})$$

The lemma follows by just observing that the payment under bid profile b is at least $\gamma b_{\max}$ ■

Corollary 10.2.2 (Simultaneous with Budgets). *If we run m simultaneous γ-hybrid auctions and bidders have budgets and fractionally subadditive valuations then every* BAYES-CCE *that satisfies the no-overbidding assumption achieves $\frac{\gamma(1-\frac{1}{e})+(1-\gamma)^2}{1+(1-\gamma)^2}$ of the expected optimal social welfare.*

Corollary 10.2.3 (Sequential). *If we run m sequential γ-hybrid auctions with unit-demand bidders then every* BAYES-CE *that satisfies the no-overbidding assumption achieves $\frac{\gamma(1-\frac{1}{e})+(1-\gamma)^2}{2+(1-\gamma)^2}$ of the expected optimal social welfare.*

11
POSITION AUCTIONS

In this section we consider a generalized version of position auctions introduced by Abrams et al [2], which allows us to extend analysis of ad auctions to simultaneous and sequential composition as well as to the case where players have budget constraints. It also allows us to capture settings where bidders have values not only per-click but also per-impression, which is considered an interesting direction from a practical perspective since many companies on the web strive mainly for impressions rather than clicks. We also propose new simple mechanisms that are approximately efficient and robust in terms of simultaneous and sequential composition and in terms of budget constraints.

11.1 General Monotone Valuations

Consider a setting where the outcome space is an allocation of n positions to n agents. Each agent has a valuation v_{ij} for being allocated position j and such that the valuations of all the agents are monotone decreasing: if $j \leq j'$ then $v_{ij} \geq v_{ij'}$. The value v_{ij} could be thought of as the value of player i for appearing at position j. This value could consist of a per-click part and a per-impression part. For instance, if the bidder thinks that his click-through-rate at position j is a_{ij} and knows that his value per-click is v_i^{c}, while he also has a value v_{ij}^{im} for appearing at position j, then his valuation for position j is: $v_{ij} = a_{ij}v_i^{\mathrm{c}} + v_{ij}^{\mathrm{im}}$. We just assume that the above total value is monotone in position.

Observe that in our framework notation the allocation space $\mathcal{X}$ consists of vectors $x = (j_1, \ldots, j_n)$ such that $j_i \neq j_{i'}$ for all $i \neq i'$. In addition the allocation space of each player $\mathcal{X}_i = \{1, \ldots, n\}$ is totally ordered from his perspective (in that any outcome where he gets a higher position is greater than any outcome where he gets a lower one) and the value of a player is monotone with respect to his own ordering of allocations.

A position mechanism M defines the action space of the players. We will consider here mechanisms where players submit only a single bid b_i (interpreted differently by the different mechanisms that we consider). Given a bid profile, the allocation function of a position mechanism is to assign a position $j_i(b)$ to each player i and the payment of each player is some function of the bid profile $P_i(b)$.

We show that for the class of position-monotone valuations a greedy first price pay-per-impression mechanism (Mechanism 5) is $(\frac{1}{2}, 1)$-smooth and its second price analog is weakly $(\frac{1}{2}, 0, 1)$-smooth. In Appendix A.6 we show that the standard Generalized Second Price auction that doesn't take into account per-impression values has high inefficiency, and therefore our modification is necessary for constant efficiency guarantees.

MECHANISM 5: Greedy first price pay-per-impression mechanism.

1. Solicit a single bid b_i from each player i;
2. Order the players according to bids;
3. Allocate positions to players in the order of the bids (i.e. the highest bidder gets the first position, etc.);
4. Charge each player his bid b_i

Lemma 11.1.1. *Mechanism 5 is $\left(\frac{1}{2}, 1\right)$-smooth via conservative deviations when valuations are monotone in the position.*

Proof. Consider a valuation profile v and any bid profile b and let j_i^* be the optimal position of player i under valuation profile v. Suppose that player i deviates to bidding according to the uniform distribution $U[0, v_{ij_i^*}]$. For a given bid profile b, let $\pi(j)$ be the player allocated at position j. If the bid that the player submits is greater than $b_{\pi(j_i^*)}$ then he is allocated position at least as high as j_i^*. By monotonicity of valuations with respect to position we get that his utility from the deviation is at least:

$$\begin{aligned} U_i^{\mathcal{M}}(\mathbf{b}_i^*, b_{-i}; v_i) &\geq \int_0^{v_{ij_i^*}} \left(v_{ij_i^*} \cdot \mathbf{1}_{\{b_{\pi(j_i^*)} < t\}} - t \right) f(t)dt \\ &\geq \int_{b_{\pi(j_i^*)}}^{v_{ij_i^*}} v_{ij_i^*} f(t)dt - \int_0^{v_{ij_i^*}} tf(t)dt \\ &= v_{ij_i^*} - b_{\pi(j_i^*)} - \frac{v_{ij_i^*}}{2} = \frac{v_{ij_i^*}}{2} - b_{\pi(j_i^*)} \end{aligned}$$

where $f(x) = 1/v_{ij_i^*}$ is the density function. Summing over all players we get the theorem. ■

The fact that the above mechanism is smooth for any monotone valuation allows us to invoke Corollary 4.3.13 and get composability results. In addition the fact that the class of monotone valuations is closed under cappings allows us to invoke our budget constraint results.

Corollary 11.1.2 (Simultaneous with Budgets)**.** *If we run m greedy first price pay-per-impression mechanisms simultaneously and bidders have monotone fractionally subadditive valuations and budget constraints then any* BAYES-CCE *achieves at least* $1/2$ *of the expected optimal effective welfare.*

Corollary 11.1.3 (Sequential)**.** *If we run m greedy first price pay-per-impression mechanisms sequentially and bidders have unit-demand valuations then* BAYES-CE-POA ≤ 4.

Note that in the last theorem, unit-demand valuations in our terminology, means that a players value for getting several impressions at different position mechanisms is of the form:

$$v_i(j_i^1, j_i^2, \ldots, j_i^m) = \max_{k \in [m]} v_i^k(j_i^k),$$

where the induced valuations $v_i^k(j_i^k)$ are monotone in the position j_i^k allocated at position mechanism $\mathcal{M}_k$.

Threshold-Price Mechanism 5. We also consider the variation of Mechanism 5 studied in Abrams et al. [2], where each player is charged the bid of the player in the position beneath him. We show that such a mechanism is conservatively and weakly $\left(\frac{1}{2}, 0, 1\right)$-smooth, implying a bound of $1/4$ in isolation, when composed simultaneously under budget constraints and when composed sequentially.

In [2] it was shown that in the full information setting there will always exist a Pure Nash Equilibrium of this mechanism that achieves optimal social welfare, thereby generalizing the result of Edelman et al [21] where only valuations per-click where considered. However, no price of anarchy analysis exists for this mechanism and the Bayesian setting or solution concepts that use randomization have not been studied.

First, we clarify our no-overbidding assumption for the mechanism of [2]. In this mechanism when a player is allocated position j with a bid b_i then his maximum willingness-to-pay is b_i, since in the strategy profile where the player in position $j+1$ bids b_i too, he is charged b_i. Thus under Definition 6.1.1 of the implicit bid we have:

$$B_i(b_i, j) = b_i$$

Using a proof identical to that of Lemma 11.1.1 we can show the weak and conservative smoothness of this mechanism.

Lemma 11.1.4. *The threshold price version of Mechanism 5 where each player is charged the bid in the position beneath him is weakly* $\left(\frac{1}{2}, 0, 1\right)$*-smooth via conservative deviations.*

Our no-overbidding assumption states that in expectation no player is bidding more than his value for the expected position he gets. A randomized bid profile $\mathbf{b}$ satisfies the no-overbidding assumption if:

$$\mathbb{E}_{\mathbf{b}_i}[\mathbf{b}_i] \leq \mathbb{E}_{\mathbf{b}}[v_{ij_i(\mathbf{b})}]$$

If a player participates in many position auctions his strategy is to submit a bid b_i^k at each position auction $\mathcal{M}_k$. Let $b_i = (b_i^k)_{k\in[m]}$ and $b^k = (b_i^k)_{i\in[n]}$. Then the no-overbidding assumption generalizes to:

$$\mathbb{E}_{\mathbf{b}_i}\left[\sum_{k\in[m]} \mathbf{b}_i^k\right] \leq \mathbb{E}_{\mathbf{b}}\left[v_i(j_i^1(\mathbf{b}^1), \ldots, j_i^m(\mathbf{b}^m))\right]$$

Under this no-overbidding assumption, our framework gives the following results.

Corollary 11.1.5 (Simultaneous with Budgets)**.** *If we run* m *greedy threshold price pay-per-impression mechanisms simultaneously and bidders have monotone fractionally subadditive valuations and budget constraints then any* BAYES-CCE *that satisfies the no-overbidding assumption, achieves at least* $1/4$ *of the expected optimal effective welfare.*

Corollary 11.1.6 (Sequential)**.** *If we run* m *greedy threshold price pay-per-impression mechanisms sequentially and bidders have unit-demand valuations then every* BAYES-CE *that satisfies the no-overbidding assumption, achieves at least* $1/4$ *of the expected optimal welfare.*

11.2 Per-Click Valuations

To draw a stronger connection with existing position auction literature we now examine the case when bidders have only valuations per-click and not per impression. We will consider two special cases of bidder valuations:

1. $v_{ij} = a_{ij}\tilde{v}_{ij}$: click-through-rates of player i at position j depend on both i and j in a non-separable way and players have position specific per-click valuations.
2. $v_{ij} = a_{ij}\tilde{v}_i$: per-click valuations are position independent

While this class of valuations neglects effects captured by the more general valuation models, special cases of this model are widely used in the literature. The latter case contains the separable model that has been long studied in the algorithmic game theory literature and has become the standard [21, 13, 47].

Definition 11.2.1. *We say that the click through rates are separable, when $a_{ij} = \alpha_j\gamma_i$ for all i and j, that is, the click through rate is the product of a factor depending on the slot and a factor depending on the advertiser.*

We use our smoothness framework to strengthen results in the literature. We give a simple smooth mechanism for the first case, which is equivalent to the standard form of the Generalized First Price (GFP) auction when specialized to the case of separable click-through rates $a_{ij} = \alpha_j\gamma_i$, showing an $1/2$ efficiency bound on GFP and its generalization to the first case above, and an $1/4$ efficiency bound for the corresponding second price analog.

For the second case, we show that the above auction is $(1 - 1/e, 1)$ smooth when using first price and weakly $(1 - 1/e, 0, 1)$ smooth when using second price. This result generalizes the efficiency bound of $\frac{1}{2}(1-\frac{1}{e})$ of Caragiannis et al [13] that considered only the separable case when $a_{ij} = \alpha_j \gamma_i$.

Note, however, that this class of valuations is not closed under capping, so our results do not extend to the case with budgets. The smoothness results we provide in the remainder of the section do imply efficiency guarantees in isolation and for special cases of complement-free valuations (e.g. bidders have value only for the k highest impressions they got and their value per impression is of the form for which smoothness is proved).

11.2.1 Variable Click Value

First we consider the case when each bidder i has a click-through-rate $a_{ij} \in [0,1]$ when he occupies position j and a value $\tilde{v}_{ij}$ when he receives a click at position j. We assume that a_{ij} and $\tilde{v}_{ij}$ are both decreasing in j. In addition, players submit per click bids and thereby are charged $a_{ij}b_i$. Hence, a player's utility at bid profile b is:

$$U_i^{\mathcal{M}}(b; v_i) = a_{ij_i(b)}(\tilde{v}_{ij_i(b)} - b_i) \tag{11.1}$$

The utility of a player is quasi-linear with value $v_{ij} = a_{ij}\tilde{v}_{ij}$ and payment scheme $P_i(b) = a_{ij_i(b)}b_i$.

MECHANISM 6: Greedy first price pay-per-click position mechanism for non-separable click-through-rates.

1 Solicit a single bid b_i from each player i;
2 Allocate position 1 to the bidder $i_1 = \arg\max_i a_{i1}b_i$. Allocate position 2 to bidder $i_2 = \arg\max_{i \neq i_1} a_{i2}b_i$, etc.;
3 If player i is allocated to position j charge him $a_{ij}b_i$.

Lemma 11.2.2. *Mechanism 6 is $\left(\frac{1}{2}, 1\right)$-smooth when click-through-rates and valuations per click are monotone in the position.*

Proof. Consider a valuation profile v and a bid profile b and let j_i^* be the optimal position of player i under valuation profile v. Suppose that player i deviates to bidding according to the uniform distribution $U[0, \tilde{v}_{ij^*(i)}]$. Then his utility from the deviation is:

$$U_i^{\mathcal{M}}(\mathbf{b}_i', b_{-i}; v_i) = \int_0^{\tilde{v}_{ij_i^*}} a_{ij_i(t,b_{-i})}(\tilde{v}_{ij_i(t,b_{-i})} - t)f(t)dt$$

where $f(x) = 1/\tilde{v}_{ij^*(i)}$ is the density function. Let $\pi(j_i^*)$ be the player allocated at position j_i^* under bid profile b. Observe that if $a_{ij_i^*}t \geq a_{\pi(j_i^*)j_i^*}b_{\pi(j_i^*)}$ then if player i hasn't already been allocated a higher position he will definitely win position j_i^*. Thereby, for t in the above range we know that player i will get a position higher than or equal to j_i^*.

Since $\tilde{v}_{ij}$ is decreasing in j and $t \in [0, \tilde{v}_{ij_i^*}]$ for

$$t \geq \frac{a_{\pi(j_i^*)j_i^*}}{a_{ij_i^*}} b_{\pi(j_i^*)} = \tau_i$$

we have that $\tilde{v}_{ij_i(t,b_{-i})} - t$ is positive. By the monotonicity of a_{ij} with respect to j we have:

$$\begin{aligned}
U_i^{\mathcal{M}}(\mathbf{b}_i', b_{-i}; v_i) &\geq \int_{\tau_i}^{\tilde{v}_{ij_i^*}} a_{ij_i^*}(\tilde{v}_{ij_i^*} - t)f(t)dt - \int_0^{\tau_i} a_{ij_i(t,b_{-i})}tf(t)dt \\
&\geq \int_{\tau_i}^{\tilde{v}_{ij_i^*}} a_{ij_i^*}(\tilde{v}_{ij_i^*} - t)f(t)dt - \int_0^{\tau_i} a_{ij_i^*}tf(t)dt \\
&= \int_{\tau_i}^{\tilde{v}_{ij_i^*}} a_{ij_i^*}\tilde{v}_{ij_i^*}f(t)dt - \int_0^{v_{ij_i^*}} a_{ij_i^*}tf(t)dt \\
&= a_{ij_i^*}\tilde{v}_{ij_i^*} - a_{ij_i^*}\tau_i - \int_0^{\tilde{v}_{ij_i^*}} a_{ij_i^*}tf(t)dt \\
&= a_{ij_i^*}\tilde{v}_{ij_i^*} - a_{\pi(j_i^*)j_i^*}b_{\pi(j_i^*)} - a_{ij_i^*}\frac{\tilde{v}_{ij_i^*}}{2} \\
&= \frac{a_{ij_i^*}\tilde{v}_{ij_i^*}}{2} - a_{\pi(j_i^*)j_i^*}b_{\pi(j_i^*)}
\end{aligned}$$

By summing over all players we get the theorem. ■

Separable CTRs. Observe that when the click-through-rates are separable, then Mechanism 6 takes the standard form of the Generalized First Price (GFP) auction that has been studied in the literature. Specifically, the allocation function of Mechanism 6 can be concisely described as: weight each players bid by his quality factor and allocate positions in order of the weighted bid. Each player is then charged his bid, per-click: $a_{j_i(b)}\gamma_i b_i$. Hence, the utility of a player at some bid profile is:

$$U_i^{\mathcal{M}}(b; v_i) = a_{j_i(b)}(\tilde{v}_{ij_i(b)} - b_i) \tag{11.2}$$

The specialization of Lemma 11.2.2 for separable click-through-rates gives a $\frac{1}{2}$ efficiency bound for the Generalized First Price auction even when the per-click valuations of the players.

Corollary 11.2.3. *The Generalized First Price Auction is* $\left(\frac{1}{2}, 1\right)$*-smooth when click-through-rates are separable and valuations per-click are dependent on the position.*

11.2.2 Position Independent Value-per-Click

Better smoothness properties can be derived if the value per-click of a player is the same for all positions (denoted by $\tilde{v}_i$), even when click-through-rates are not separable.

Lemma 11.2.4. *Mechanism 6 is* $\left(1 - \frac{1}{e}, 1\right)$*-smooth when per-click valuations are position independent even if click-through-rates are not separable.*

Proof. Suppose that player i deviates to bidding according to distribution with density function $f(t) = \frac{1}{\tilde{v}_i - t}$ and support $[0, (1 - \frac{1}{e})\tilde{v}_i]$. As stated in the proof of Lemma 11.2.2 if

$$t \geq \frac{a_{\pi(j_i^*)j_i^*}}{a_{ij_i^*}} b_{\pi(j_i^*)} = \tau_i$$

(where $\pi(j)$ is the player at position j in the current bid profile) then player i is assigned a position at least as high as j_i^*. Thus his utility from the deviation is:

$$\begin{aligned} U_i^{\mathcal{M}}(\mathbf{b}_i', b_{-i}; v_i) = & \int_0^{(1-\frac{1}{e})\tilde{v}_i} a_{ij_i(t,b_{-i})}(\tilde{v}_i - t) f(t) dt \\ \geq & \int_{\tau_i}^{(1-\frac{1}{e})\tilde{v}_i} a_{ij_i^*}(\tilde{v}_i - t) f(t) dt \\ = & \left(1 - \frac{1}{e}\right) a_{j_i^*}\tilde{v}_i - a_{\pi(j_i^*)} b_{\pi(j_i^*)} \end{aligned}$$

By summing over all players we get the theorem. ■

If the click-through-rates are separable, then this brings us to the standard model studied in the literature, where the valuation of a player i from being assigned at position j is: $a_j \gamma_i v_i$ and thereby the utility of a player at some bid profile is:

$$U_i^{\mathcal{M}}(b; v_i) = a_{j_i(b)} \gamma_i (\tilde{v}_i - b_i) \tag{11.3}$$

The specialization of Lemma 11.2.4, gives a better smoothness property for the Generalized First Price Auction.

Corollary 11.2.5. *The GFP auction is $\left(1 - \frac{1}{e}, 1\right)$-smooth when per-click valuations are position independent and click-through-rates are separable.*

Threshold Price Mechanism 6. We consider a threshold-price version of Mechanism 6 where a player is charged, per-click, the minimum bid he had to make to get his position. To draw a strong connection with existing position auction literature we will analyze this mechanism only on the special case where click-through-rates are separable and valuations are position independent: $v_{ij} = a_j \gamma_i \tilde{v}_i$.

Under this valuation model the threshold-price version of Mechanism 6 becomes the standard Generalized Second Price (GSP) auction introduced by Edelman et al [21] and extensively studied from the price of anarchy perspective [47, 13]. In this mechanism, under strategy profile b, each player i is charged $\frac{\gamma_{\pi(j_i(b)+1)} b_{j_i(b)+1}}{\gamma_i}$ per-click, where $\pi(j)$ is the player that got position j under bid profile b, and thus his utility at some bid profile is:

$$U_i^{\mathcal{M}}(b) = a_j \gamma_i \left(\tilde{v}_i - \frac{\gamma_{\pi(j_i(b)+1)} b_{j_i(b)+1}}{\gamma_i} \right) \tag{11.4}$$

In this mechanism a player's willingness-to-pay for a position j is simply $a_j \gamma_i b_i$ since in the special case where the player beneath him was bidding $\frac{\gamma_i}{\gamma_{\pi(j_i(b)+1)}} b_i$, player i is charged an expected total payment of $a_j \gamma_i b_i$. Thus from Definition 6.1.1 of willingness-to-pay we have that:

$$B_i(b_i, j) = a_j \gamma_i b_i$$

Our no-overbidding assumption will then become:

$$\mathbb{E}_{\mathbf{b}}[a_{j_i(b)} \mathbf{b}_i] \leq \mathbb{E}_{\mathbf{b}}[a_{j_i(b)} \tilde{v}_i]$$

Caragiannis et al [13] use a point-wise no-overbidding assumption that for any bid in the support of a player's strategy $b_i \leq \tilde{v}_i$ and prove that such an overbidding is weakly dominated. That assumption would imply our weaker in expectation assumption.

Under the no-overbidding assumption and using the same proof as in Lemma 11.2.4 specialized for separable click-through-rates would give that the

Generalized Second Price auction is weakly $\left(1-\frac{1}{e},0,1\right)$-smooth. This implies the efficiency result of $\frac{1}{2}\left(1-\frac{1}{e}\right)$ that was given in Caragiannis et al [13] and the proof of Lemma 11.2.4 is a generalization of their analysis.

Corollary 11.2.6. *The Generalized Second Price auction is weakly* $\left(1-\frac{1}{e},0,1\right)$*-smooth when per-click valuations are position independent and click-through-rates are separable.*

In fact applying the same proof of Lemma 11.2.4 we get a generalization of this result for non-separable click-through-rates.

Corollary 11.2.7. *The threshold-price version of Mechanism 6 is weakly* $\left(1-\frac{1}{e},0,1\right)$*-smooth when per-click valuations are position independent, even if the click-through-rates are not separable.*

This result implies an efficiency bound of $\frac{1}{2}\left(1-\frac{1}{e}\right)$ under the same no-overbidding assumption that was used by Caragiannis et al [13].

12

DIRECT MECHANISMS

In this section we will focus on mechanisms where players directly report their values. These mechanisms show an interesting use of threshold bids (critical also in truthful mechanism design) in *smooth mechanism design*. In a direct mechanism the player declares a value $\tilde{v}_i(x_i)$ for any allocation outcome x_i. Thus the action space $\mathcal{A}_i$ of each player defined by the mechanism is equal to the set of valuations $\mathcal{V}_i$.

We will focus on direct mechanisms that are also ex-post individually rational. We apply the individual rationality constraint to non-truthful mechanisms in the sense of requiring that if reporting valuations truthfully the resulting utility of any agent is non-negative point-wise over every randomness of the mechanism. Note that this doesn't imply that the player do reports truthfully. It just places the the restriction on the allocation function $X : \mathcal{A} \to \mathcal{X}$ and the payment function $P : \mathcal{A} \to \mathbb{R}_+^n$ of the mechanism that $P_i(\tilde{v}) \leq \tilde{v}_i(X_i(\tilde{v}))$.

12.1 Critical Payments and Smooth Direct Mechanisms

To motivate the connection to truthful mechanism design, we first describe a single-dimensional service-based mechanism design setting. Consider a setting where the allocation space is just a feasibility set of players that can be served. In other words the outcome space from the perspective of each player is binary $\mathcal{X}_i = \{0, 1\}$ and the outcome space of the mechanism is some feasible subset of the product space. In addition, suppose that the value of a player was just a number v_i for being served. The utility of an agent would then be of the form $u_i(x,p) = v_i x_i - p_i$, where x_i is either 0 or 1. It is a well known result that an efficient truthful direct mechanism needs to charge player i, the minimum value he needs to have to still be allocated: $P_i(\tilde{v}) = \tau_i(\tilde{v}_{-i}) = \inf\{z : X_i(z, \tilde{v}_{-i}) = 1\}$ and 0 if he is not allocated.

Is there a similar characterization for the more general setting? For more general settings the standard mechanism that is efficient and guarantees non-negative prices and individual rationality is the VCG mechanism. Unfortunately, the VCG mechanism doesn't have a similar simple "threshold bid" interpretation. Despite this fact, one could still define threshold bids in the more general quasi-linear setting as follows:

Definition 12.1.1. *Given a direct mechanism M and a bid profile $\tilde{v}$, we say that the threshold bid $\tau_i(x_i, \tilde{v}_{-i})$ of player i for allocation x_i is the minimal value that player i has to single mindedly declare for allocation x_i such that he is allocated x_i whenever $\tilde{v}_i(x_i) \geq \tau_i(x_i, \tilde{v}_{-i})$.*

To make a mechanism truthful in a single parameter setting remember that one had to strongly tie together the threshold bids of the players with their actual payments. In what follows we show that even in smooth mechanism design in order to get approximately efficient smooth mechanisms one needs to approximately tie threshold bids to the payments.

Definition 12.1.2. *A direct mechanism is c-threshold approximate, for some $c \geq 0$, if for any feasible allocation $x \in \mathcal{X}$ and any reported valuation profile $\tilde{v}$:*

$$\sum_i \tau_i(x_i, \tilde{v}_{-i}) \leq c \sum_i P_i(\tilde{v}) \tag{12.1}$$

The above relation between threshold bids and payments is in the essence of the analysis of Lucier and Borodin [46] as described in the next section.

As an example, consider a first price single item auction setting. Consider a bid profile $\tilde{v}$ and a feasible allocation where player i wins. The threshold payment for player i to win the auction is exactly equal to the payment that the winner was paying under bid profile $\tilde{v}$. The rest of the players are not allocated hence their threshold payment is 0. Hence, we observe that a single item first price auction is a 1-threshold approximate mechanism.

First, we show that if a mechanism is c-threshold approximate then this implies a good efficiency guarantee on the Bayes-Nash Equilibria of the game that it induces.

Theorem 12.1.3. *If a direct mechanism is c-threshold approximate and individually rational then it is $\left(\beta\left(1-e^{-1/\beta}\right), \beta c\right)$-smooth via conservative deviations for any $\beta \geq 0$.*

Proof. Consider an instance of the valuation profile v and a bid profile $\tilde{v}$. Suppose that bidder i submits a single-minded value θ for his optimal allocation $x_i^*(v)$. If $\tau_i(x_i^*(v), \tilde{v}_{-i}) \geq \theta$ the agent doesn't get allocated $x_i^*(v)$. Otherwise he is allocated and pays $P_i(t, \tilde{v}_i)$. Since the mechanism satisfies ex post individual rationality this payment cannot be more than θ. Otherwise a player with true value t for $x_i^*(v)$ would be getting negative utility.

Thus the utility of a player from this deviation is at least:

$$U_i^{\mathcal{M}}(\theta, \tilde{v}_{-i}; v_i) \geq (v_i(x_i^*(v)) - \theta)\mathbf{1}_{\theta > \tau_i(x_i^*(v), \tilde{v}_{-i})}$$

Now a player by using a randomized Θ that follows a distribution with density $f(\theta) = \frac{\beta}{v_i(x_i^*(v))-\theta}$ for $\theta \in [0, v_i(x_i^*(v))(1-e^{-1/\beta})]$ he will get expected utility at least:

$$\begin{aligned} u_i(\Theta, \tilde{v}_{-i}) \geq & \int_{\tau_i(x_i^*(v), \tilde{v}_{-i})}^{v_i(x_i^*(v))(1-e^{-1/\beta})} (v_i(x_i^*(v)) - t) f(\theta) d\theta \\ \geq & \ \beta \left(1 - e^{-1/\beta}\right) v_i(x_i^*(v)) - \beta \tau_i(x_i^*(v), \tilde{v}_{-i}) \end{aligned}$$

Adding over all players and using the c-threshold payment approximate property we get the theorem. ■

A second price auction on the other hand is not threshold approximate for any μ. The reason is the following: consider a type profile v and a bid profile $\tilde{v}$ where a player with a very small value bids a huge number H and the rest of the players bid truthfully. Then the threshold payment for the highest value player is H, while the payment that the auction receives is of the order of the values of the rest of the players. Thus payments and threshold payments are unrelated. The latter is the crucial difference between first price and second price payment rules and the reason why we need to employ no-overbidding assumptions to give efficiency guarantees for second-price payment schemes. Similar to how c-threshold approximate direct mechanisms are connected to smoothness, the following property of weak c-threshold approximate mechanisms is connected to weak smoothness:

Definition 12.1.4. *A direct mechanism is weakly (c_1, c_2)-threshold approximate, for some $c_1, c_2 \geq 0$, if for any feasible allocation $x \in \mathcal{X}$ and any reported valuation profile $\tilde{v}$:*

$$\sum_i \tau_i(x_i, \tilde{v}_{-i}) \leq c_1 \sum_i P_i(\tilde{v}) + c_2 \sum_i B_i(\tilde{v}_i, X_i(\tilde{v})) \tag{12.2}$$

Theorem 12.1.5. *If a direct mechanism is (c_1, c_2)-threshold approximate and individually rational then it is weakly and conservatively*

$$\left(\beta \left(1 - e^{-1/\beta}\right), \beta c_1, \beta c_2\right) \textit{-smooth}$$

for any $\beta \geq 0$.

The proof is similar to that of Theorem 12.1.3 and is omitted.

12.2 Greedy Direct Combinatorial Auctions

A very interesting instance of c-threshold approximate mechanisms in the literature is that of Greedy Direct Auctions introduced by Lucier and Borodin [46]. In the terminology that we introduced in the previous section, Lucier and Borodin [46] proved that in any direct combinatorial auction setting if the allocation is decided by a greedy c-approximate mechanism then coupling the mechanism with a first price payment rule we get a c-threshold approximate mechanism. These proofs don't assume anything about the valuation of a player and allow for complements.

We note that not all greedy algorithms adhere to the framework defined by Lucier and Borodin [46]. For instance, the greedy matroid mechanisms we studied in Chapter 8 do not fall into their framework. The greedy mechanisms studied in [46] are as follows.

1. Solicit valuation reports $\tilde{v}$.
2. At each iteration pick an (agent, set) pair (i, S) that maximizes a ranking function $r(i, S, \tilde{v}_i(S))$, and allocate S to i.
3. Remove both i and S from consideration and repeat until all items are allocated.

The ranking function is monotone in S (by inclusion) and $v_i(S)$ and could potentially be adaptive with respect to the existing allocation. For the case of general combinatorial auctions a $\sqrt{k}$-approximate greedy such algorithm exists, where k is the number of items.

Hence, for the setting of greedy first price c-approximate combinatorial auctions our framework implies:

Corollary 12.2.1. *Any* BAYES-CCE *of a greedy c-approximate first price combinatorial auction achieves at least $\frac{1-e^{-c}}{c}$ of the expected optimal social welfare. If bidders have budgets then it achieves the same fraction of the optimal effective welfare.*

Lucier and Borodin [46] give a bound of $\frac{1}{c+O(\log(c))}$ for the efficiency of such a greedy auction. More specifically the bound is $\frac{c-1-\log(c)}{c(c+1+\log(c))}$. Our bound is asymptotically same, but is strictly better. Our bound decreases as $\frac{1}{c+O(ce^{-c})}$ rather than $\frac{1}{c+O(\log(c))}$. When $c = 1$ the bound coincides with the bound of $1 - \frac{1}{e}$ for the first price single item auction and our bound is always larger than $\frac{1}{c+\frac{1}{e-1}} \approx \frac{1}{c+0.58}$ and thereby decreases exactly linearly with c.

Our composability framework gives new results for the case when several greedy combinatorial auctions are run simultaneously or sequentially.

Corollary 12.2.2 (Simultaneous with Budgets). *If we run m greedy c-approximate first price combinatorial auctions simultaneously and bidders have budgets and monotone fractionally subadditive valuations across mechanisms, then any* BAYES-CCE *achieves at least $\frac{1-e^{-c}}{c}$ of the expected optimal effective welfare.*

Corollary 12.2.3 (Sequential). *If we run m greedy c-approximate first price combinatorial auctions sequentially and bidders have unit-demand valuations across mecha-*

nisms, then any BAYES-CE *achieves at least* $\frac{1-e^{-c}}{c+1}$ *of the expected optimal social welfare.*

Recall that fractional subadditivity across mechanisms does not impose any assumption on the valuations within each greedy combinatorial auction. Hence, the valuations of the bidders could have complements within the items sold in each greedy auction, as long as they don't have complements across items sold in different auctions.

Lucier and Borodin [46] also examine a second-price type of payment scheme where each agent is charged his threshold bid for the allocation that he is awarded. In such a mechanism the willingness-to-pay for an allocation is exactly a player's bid for that allocation. They show that the mechanism is weakly $(0, c)$-threshold approximate. By applying theorem 12.1.5 for $\beta = 1$ we get an efficiency guarantee $\frac{\beta(1-e^{-1/\beta})}{\beta c+1} = \frac{1-1/e}{c+1}$ for an individual greedy mechanism and for simultaneous and sequential composition. Instead of using the generic smoothness result of Theorem 12.1.5, by reinterpreting the techniques of Lucier and Borodin [46] we can easily show that this mechanism is actually weakly $(1, 0, c)$-smooth, by considering the deviation where each player switches to single-mindedly bidding his true valuation on his optimal set [1]. This gives the slightly better efficiency guarantee of $\frac{1}{c+1}$.

[1] Since the mechanism charges threshold bids, the utility of a player from this deviation is at least $v_i(x_i^*(v)) - \tau_i(x_i^*(v), \tilde{v}_{-i})$. Summing over all players and using the fact that the greedy mechanism is $(0, c)$-threshold approximate we get the weak $(1, 0, c)$-smoothness result.

13
BANDWIDTH ALLOCATION

In this section we consider the single-link version of the bandwidth allocation setting of Johari and Tsitsiklis [38]. In this setting a bandwidth of C is to be split among n bidders.

13.1 Kelly's Proportional Bandwidth Allocation Mechanism

The bidders submit a bid b_i which they have to pay no matter how much bandwidth they receive. Given the bid profile each player is allocated a bandwidth proportional to his bid:

MECHANISM 7: Proportional bandwidth allocation mechanism.

Solicit a single bid b_i from each player i;
Allocate to player i bandwidth $x_i(b) = \frac{b_i C}{\sum_{j \in N} b_j}$;
Charge each player his bid b_i

Each player has a concave value function $v_i(x_i)$ for getting a share of bandwidth x_i, with $v_i(0) = 0$, and his utility is quasi-linear with respect to payments:

$$U_i^{\mathcal{M}}(b; v_i) = v_i(x_i(b)) - b_i \tag{13.1}$$

As one can easily observe the latter mechanism falls into our general definition of a mechanism with quasi-linear preferences. We will show that such a mechanism is $(2-\sqrt{3}, 1)$-smooth. This will imply efficiency guarantees of approximately $1/4$ for any CE and BNE as well as for simultaneous compositions and sequential composition of bandwidth allocation mechanisms. For the simultaneous setting it also implies such a bound even when players have budget constraints. Johari and Tsitsiklis [38] give an efficiency bound of $3/4$ but their efficiency guarantee is proved only for the case of pure nash equilibria and only in the complete information setting. Hence, though our bound is slightly worse, it is a bound that extends to a plethora of relaxations and extensions.

Lemma 13.1.1. *The proportional bandwidth allocation mechanism is* $(2-\sqrt{3},1)$*-smooth via conservative deviations when value functions* $v_i : [0, C] \to \mathbb{R}^+$ *are concave and* $v(0) = 0$*.*

Proof. Given a valuation profile v for each player, let $x_i^*(v)$ be the bandwidth allocated to player i in the optimal allocation. For simplicity we will denote it with x_i^* for the remainder of the proof since we focus on a specific valuation profile.

Consider the deviation where player i deviates to bidding uniformly at random $\mathbf{b}_i^* \sim U[0, \lambda v_i(x_i^*)]$, for some constant λ that will be determined later on. Then his expected utility for any bid profile b_{-i} is as follows:

$$U_i^{\mathcal{M}}(\mathbf{b}_i^*, b_{-i}; v_i) = \int_0^{\lambda v_i(x_i^*)} \frac{v_i(x_i(t, b_{-i}))}{\lambda v_i(x_i^*)} dt - \frac{1}{2}\lambda v_i(x_i^*)$$

Given the bids of the rest of the players b_{-i}, if player i bids above $\frac{x\sum_{j\neq i} b_j}{C-x}$ then he is given a bandwidth share of at least x for any x. Thus for all the $t \geq \frac{x_i^*\sum_{j\neq i} b_j}{\mu C - x_i^*}$ player i is allocated a bandwidth of at least x_i^*/μ.

Thus by monotonicity of v_i his utility from the deviation is at least:

$$U_i^{\mathcal{M}}(\mathbf{b}_i^*, b_{-i}; v_i) \geq \int_{\frac{x_i^*\sum_{j\neq i} b_j}{\mu C - x_i^*}}^{\lambda v_i(x_i^*)} \frac{v_i(x_i^*/\mu)}{\lambda v_i(x_i^*)} dt - \frac{1}{2}\lambda v_i(x_i^*)$$

By concavity and the fact that $v_i(0) = 0$ we know that $v_i\left(\frac{x_i^*}{\mu}\right) \geq \frac{v_i(x_i^*)}{\mu}$ for any

$\mu \geq 1$. Thus:

$$\begin{aligned} U_i^{\mathcal{M}}(\mathbf{b}_i^*, b_{-i}; v_i) \geq & \int_{\frac{x_i^* \sum_{j \neq i} b_j}{\mu C - x_i^*}}^{\lambda v_i(x_i^*)} \frac{v_i(x_i^*)}{\mu \lambda v_i(x_i^*)} dt - \frac{1}{2} \lambda v_i(x_i^*) \\ = & \int_{\frac{x_i^* \sum_{j \neq i} b_j}{\mu C - x_i^*}}^{\lambda v_i(x_i^*)} \frac{1}{\mu \lambda} dt - \frac{1}{2} \lambda v_i(x_i^*) \\ = & \frac{1}{\mu} v_i(x_i^*) - \frac{1}{\lambda \mu} \frac{x_i^* \sum_{j \neq i} b_j}{\mu C - x_i^*} - \frac{1}{2} \lambda v_i(x_i^*) \end{aligned}$$

Since $x_i^* \leq C$ and $\sum_{j \neq i} b_j \leq \sum_j b_j$ we get:

$$u_i(B_i, b_{-i}) \geq \frac{1}{\mu} v_i(x_i^*) - \frac{1}{\lambda \mu} \frac{x_i^* \sum_j b_j}{(\mu - 1)C} - \frac{1}{2} \lambda v_i(x_i^*)$$

Summing over all players we get:

$$\begin{aligned} \sum_i u_i(B_i, b_{-i}) \geq & \left(\frac{1}{\mu} - \frac{\lambda}{2} \right) \sum_i v_i(x_i^*) - \sum_i \frac{1}{\lambda \mu} \frac{x_i^* \sum_j b_j}{(\mu - 1)C} \\ = & \left(\frac{1}{\mu} - \frac{\lambda}{2} \right) \sum_i v_i(x_i^*) - \frac{1}{\lambda \mu (\mu - 1)} \sum_j b_j \end{aligned}$$

By setting $\lambda = \frac{1}{\mu(\mu-1)}$ we get that the mechanism is $(\frac{1}{\mu} - \frac{\lambda}{2}, 1)$-smooth. By optimizing over μ we get that the best bound is implied by $\mu = \frac{1}{2}(3 + \sqrt{3})$ for which we get that the mechanism is $(2 - \sqrt{3}, 1)$-smooth. ■

Now observe that the valuation space for which smoothness is proved is the set of all concave valuations on $\mathbb{R}^+$, with $v(0) = 0$. Clearly, $\mathbb{R}^+$ is a distributive lattice. Submodularity with respect to $\mathbb{R}^+$ simply means concavity. By theorem 4.3.16 we know that if several such bandwidth allocation mechanisms happen simultaneously and the valuation of the player is submodular on the product lattice, then we can express such a valuation with induced valuations that are capped marginals:

$$v_j(x_i^j) = v(x_i^j \wedge \tilde{x}_i^j, \tilde{x}_i^{-j}) - v(0, \tilde{x}_i^{-j})$$

Those function are concave with respect to $\mathbb{R}^+$ and have $v(0) = 0$. Therefore we can apply our simultaneous composition theorem for any submodular valuation with respect to the lattice $\mathbb{R}^m_+$. If the valuations are continuously differentiable then submodularity on $\mathbb{R}^m_+$ simply means that: $\frac{\partial^2 v(x_i)}{(\partial x_i^j)^2} \leq 0$ and $\frac{\partial^2 v(x_i)}{\partial x_i^j \partial x_i^{j'}} \leq 0$. Thus the function is concave coordinate-wise and has decreasing differences.

In addition observe that even if we cap a submodular valuation on $\mathbb{R}^+$ then it remains submodular. Thus we can also invoke our budget constraint theorems.

Corollary 13.1.2 (Simultaneous with Budgets). *If we run m simultaneous bandwidth allocation mechanisms and the valuations are submodular on $\mathbb{R}^m_+$ and bidders have budgets then any* BAYES-CCE *achieves at least $\frac{1}{2-\sqrt{3}} \approx \frac{1}{3.73}$ of the expected optimal effective welfare.*

For sequential composition we get our theorem for the case where the valuation of the bidder is the maximum among his valuations on different links: $v_i(x_i) = \max_j v_{ij}(x_{ij})$.

Corollary 13.1.3 (Sequential). *If we rum m sequential bandwidth allocation mechanisms and the valuations are unit-demand then any* BAYES-CE *achieves at least $\frac{1}{2(2-\sqrt{3})}$ of the expected optimal social welfare.*

14
MULTI-UNIT AUCTIONS

Consider the following setting: An auctioneer wants to sell k units of a good. A bidder's valuation is an increasing concave function $v_i(j)$ of the amount j of goods he gets. We will consider two types of auctions.

In the first section, we consider auctions where players report their marginal values for each extra unit of the good. Then the auctioneer allocates units greedily in decreasing order of marginal bids. We analyze both discriminatory payments rules, where a player pays his marginal bid for each unit allocated and uniform payment rules, where the price of each allocated unit is the same and equals the highest un-allocated marginal bid.

In the second section, we consider an auction with a simplified bidding language, where the players simply report a per-unit bid b_i and a desired quantity q_i. Essentially they declare that the are willing to pay b_i for each unit, up to q_i units. The mechanism again allocates units greedily and charges a uniform payment.

We show that both types of mechanisms are smooth or weakly smooth for any concave valuation function.

14.1 Marginal Bid Multi-Unit Auctions

We consider the following auction:

We will denote with $k_i(b)$ the units allocated to bidder i under bid profile b.

MECHANISM 8: Greedy First-Price Multi-Unit Auction with concave values.

1 Solicit bids $b_{i1}, \ldots, b_{ik}$ for marginal values from each player i which are restricted to be decreasing;
2 At each iteration allocate the extra unit to the bidder that has the maximum marginal bid for getting it, conditional on the items he has already been allocated;
3 Repeat until all units are allocated or until no player has value for an extra unit;
4 If player i is allocated k_i units then charge him $\sum_{j=1}^{k_i} b_{ij}$.

We will also denote with $p_j(b)$ to be the $j-th$ lowest price for which a unit was sold by the algorithm, i.e. the bid of the j-th from the end unit that was sold. The utility of a bidder is still quasi-linear with money:

$$U_i^{\mathcal{M}}(b; v_i) = v_i(k_i(b)) - \sum_{j=1}^{k_i(b)} b_{ij} \tag{14.1}$$

We show that the greedy multi-unit auction is $(\frac{1}{2}\left(1-\frac{1}{e}\right), 1)$-smooth thereby implying an efficiency guarantee of $\frac{e-1}{2e} \approx 1/3.16$ when studied in isolation.

Lemma 14.1.1. *Mechanism 8 is* $\left(\frac{1}{2}\left(1-\frac{1}{e}\right), 1\right)$*-smooth via conservative deviations when bidders' valuations* $v_i : \mathbb{N} \to \mathbb{R}^+$ *are concave with* $v_i(0) = 0$.

Proof. Suppose that bidder i deviates to stating that his k_i^* highest marginal valuations are all t for some randomly drawn t according to the distribution with probability density function $f(t) = \frac{1}{\frac{v_i(k_i^*)}{k_i^*} - t}$ and support $[0, \frac{v_i(k_i^*)}{k_i^*}\left(1-\frac{1}{e}\right)]$. For his remaining marginal valuations he bids 0. Then the utility of player i from this deviation is:

$$\begin{aligned} U_i^{\mathcal{M}}(\mathbf{b}_i^*, b_{-i}; v_i) &= \int_0^{\frac{v_i(k_i^*)}{k_i^*}\left(1-\frac{1}{e}\right)} (v_i(k_i(t, b_{-i})) - k_i(t, b_{-i})t) f(t) dt \\ &= \int_0^{\frac{v_i(k_i^*)}{k_i^*}\left(1-\frac{1}{e}\right)} k_i(t, b_{-i}) \left(\frac{v_i(k_i(t, b_{-i}))}{k_i(t, b_{-i})} - t\right) f(t) dt \end{aligned}$$

Since bidder i bids positive only for his k_i^* highest marginals we know that he is allocated at most k_i^* units. Hence, $k_i(t, b_{-i}) \leq k_i^*$ for all t. In addition by concavity we know that for any $k \leq k_i^* : \frac{v(k)}{k} \geq \frac{v(k_i^*)}{k_i^*}$. Hence:

$$\begin{aligned} U_i^{\mathcal{M}}(\mathbf{b}_i^*, b_{-i}; v_i) &\geq \int_0^{\frac{v_i(k_i^*)}{k_i^*}\left(1-\frac{1}{e}\right)} k_i(t, b_{-i}) \left(\frac{v_i(k_i^*)}{k_i^*} - t\right) f(t)dt \\ &= \int_0^{\frac{v_i(k_i^*)}{k_i^*}\left(1-\frac{1}{e}\right)} k_i(t, b_{-i})dt \end{aligned}$$

For any $j \in [1, k_i^*]$, if $t > p_j(b)$ then $k_i(t, b_{-i}) \geq j$. Hence, we have:

$$\begin{aligned} U_i^{\mathcal{M}}(\mathbf{b}_i^*, b_{-i}; v_i) &\geq \int_{p_j(b)}^{\frac{v_i(k_i^*)}{k_i^*}\left(1-\frac{1}{e}\right)} j dt \\ &= \frac{j}{k_i^*}\left(1 - \frac{1}{e}\right) v_i(k_i^*) - jp_j(b) \end{aligned} \tag{14.2}$$

Now we need to find the right pick of j in our analysis, such that when adding the above inequality for all players then the negative part on the right hand side will be the total revenue of the auction at bid profile b.

Since prices $p_j(b)$ are increasing in j we know that:

$$jp_j(b) \leq \sum_{t=0}^{j-1} p_{j+t}(b)$$

If we choose a j such that $2j - 1 \leq k_i^*$ then:

$$jp_j(b) \leq \sum_{t=0}^{j-1} p_{j+t}(b) \leq \sum_{t=j}^{k_i^*} p_t(b) \leq \sum_{t=1}^{k_i^*} p_t(b)$$

Observe that since k_i^* are integers, if we choose $j = \lceil \frac{k_i^*}{2} \rceil$ then we know that $j \leq \frac{k_i^*+1}{2}$ and therefore $2j - 1 \leq k_i^*$. Thus if we apply Inequality (14.2) for $j = \lceil \frac{k_i^*}{2} \rceil$ we get:

$$U_i^{\mathcal{M}}(\mathbf{b}_i^*, b_{-i}; v_i) \geq \frac{1}{2}\left(1 - \frac{1}{e}\right) v_i(k_i^*) - \sum_{t=1}^{k_i^*} p_t(b)$$

Last observe that since $\sum_i k_i^* = k$ and prices $p_t(b)$ are increasing in t: $\sum_i \sum_{t=1}^{k_i^*} p_t \leq \sum_{t=1}^{k} p_t(b) = \sum_i P_i(b)$. Hence, by summing over all players and using the latter inequality we get the theorem. ■

Now similar to the bandwidth allocation setting we can apply our simultaneous composability theorem even under budgets when players have submodular valuations on the product lattice on $\mathbb{N}^m$. Submodularity on $\mathbb{N}^m$ means that the functions must be concave coordinate-wise and must satisfy the decreasing differences.

Corollary 14.1.2 (Simultaneous with Budgets)**.** *If we run m greedy multi-unit auctions and bidders have submodular valuations on $\mathbb{N}^m$ and budget constrains then every* BAYES-CCE *achieves at least $\frac{e-1}{2e} \approx \frac{1}{3.16}$ of the expected optimal effective welfare.*

For sequential composition we require that the bidders are unit-demand over mechanisms: e.g. they have mechanism specific concave value functions and that their utility is the maximum over all mechanisms of the utility they get from each mechanism, $v_i(k_i) = \max_j v_{ij}(k_{ij})$. Observe that such valuations are a generalization of the standard notion unit-demand valuations where players just want one unit. We could simulate unit-demand valuations with unit-demand over mechanisms by just saying that $v_{ij}(k_{ij}) = \hat{v}_{ij}$ if $k_{ij} >= 1$. Our notion of unit-demand valuations over mechanisms just says that you should pick the mechanism that gave you the maximum value for the units it gave you.

Corollary 14.1.3 (Sequential)**.** *If we run m greedy multi-unit auctions sequentially and bidders have unit-demand valuations over mechanisms then every* BAYES-CE *achieves at least $\frac{e-1}{4e} \approx \frac{1}{6.32}$ of the expected optimal social welfare.*

One could also think of running a second-price equivalent of Mechanism 8 which is described in Mechanism 9.

Markakis and Telelis [50] studies exactly this auction and uses a no-overbidding assumption, where the willingness-to-pay of a player is the sum of

MECHANISM 9: Greedy Multi-Unit Threshold Price Auction with concave values.

1 Solicit marginal bids $b_{i1}, \ldots, b_{ik}$ from each player i which are restricted to be decreasing;
2 At each iteration allocate the extra unit to the bidder that has the maximum marginal value for getting it conditional on the items he has already been allocated;
3 For each unit that a player receives charge him the highest marginal bid that the mechanism didn't allocate to

his k highest marginal bids if he is allocated k units. Under this no-overbidding assumption and using similar analysis as in Theorem 14.1.1 we can prove that this auction is weakly $\left(\frac{1}{2}\left(1-\frac{1}{e}\right), 0, 1\right)$-smooth, thereby implying an efficiency guarantee of $\frac{1}{4}\left(1-\frac{1}{e}\right)$. This largely improves upon the results of Markakis et al. [50] where only a logarithmic bound in the number of units $O(\log(k))$ was proved for the case of mixed and Bayes-Nash equilibria. Our bound also has implications for budgets and simultaneous and sequential composition.

14.2 Uniform Bid Multi-Unit Auction

Another multi-unit auction that has been widely used in practice is the Uniform Price Auction. In the uniform price auction every bidder is asked to report a pair (q_i, b_i) where $q_i \in \mathbb{N}$ is a quantity and b_i is a per-unit bid. The auction then orders the bids in decreasing order and serves the units until reaching capacity.

Uniform price auctions are frequently used in practice because they have the advantage that no-matter what the players bid, everyone pays the same price for the allocated items. Hence, they give a fairness feeling and also avoid any friction when someone was allocated the same unit at a different price.

MECHANISM 10: Uniform Price Auction with concave values.

1. Solicit quantity, bid pairs (q_i, b_i) from each player i;
2. Let Q_t be the total units allocated until iteration t;
3. At each iteration t pick unallocated player with highest b_i and allocate him $\min\{q_i, k - Q_t\}$, until all units are sold ;
4. Charge everyone the highest losing bid, i.e. the bid of the last player that was partially satisfied or if the last player was completely satisfied then the bid of highest player that was unallocated.

In such an auction the willingness-to-pay of a player that received k_i units and bid b_i per unit is exactly $k_i b_i$ since in the worst case the highest losing bid could be just below your bid.

Lemma 14.2.1. *The Uniform Price Auction is weakly* $\left(\frac{1}{2}\left(1-\frac{1}{e}\right), 0, 1\right)$*-smooth via conservative deviations when bidders valuations* $v_i : \mathbb{N} \to \mathbb{R}^+$ *are concave with* $v_i(0) = 0$.

Proof. Consider a strategy profile $a = (k, b)$ and a bid profile $v : \mathbb{N} \to \mathbb{R}^+$. Suppose that bidder i deviates to bidding $a_i' = (k_i^*, t)$ where t is a number drawn randomly according to the distribution with probability density function $f(t) = \frac{\beta}{\frac{v_i(k_i^*)}{k_i^*} - t}$ and support $[0, \frac{v_i(k_i^*)}{k_i^*}\left(1 - \frac{1}{e^{1/\beta}}\right)]$.

Denote by B_t the bid of the t-th last unit sold. Similar to Mechanism 8 it holds that if $t > B_j$ then the player is allocated at least j units. Thereby using similar analysis as in the proof of Lemma 14.1.1 we can show that the above deviation yields utility at least:

$$U_i^{\mathcal{M}}(\mathbf{b}_i^*, b_{-i}; v_i) \geq \frac{\beta}{2}\left(1 - \frac{1}{e^{1/\beta}}\right) v_i(k_i^*) - \beta \sum_{t=1}^{k_i^*} B_t$$

Then by summing among all players we can derive:

$$\sum_i U_i^{\mathcal{M}}(\mathbf{b}_i^*, b_{-i}; v_i) \geq \frac{\beta}{2}\left(1 - \frac{1}{e^{1/\beta}}\right) v_i(k_i^*) - \beta \sum_i k_i(a) b_i$$
$$= \frac{\beta}{2}\left(1 - \frac{1}{e^{1/\beta}}\right) v_i(k_i^*) - \beta \sum_i B_i(a_i, k_i(a))$$

Setting $\beta = 1$ we get the theorem. ■

We also get the same composability guarantees as Mechanism 8:

Corollary 14.2.2 (Simultaneous with Budgets). *If we run m uniform price auctions and bidders have submodular valuations on $\mathbb{N}^m$ and budget constrains then every* BAYES-CCE *that satisfies the no-overbidding assumption achieves at least $\frac{e-1}{4e} \approx \frac{1}{6.32}$ of the expected optimal effective welfare.*

Corollary 14.2.3 (Sequential). *If we run m uniform price auctions sequentially and bidders have unit-demand valuations over mechanisms then every* BAYES-CE *achieves at least $\frac{e-1}{4e} \approx \frac{1}{6.32}$ of the expected optimal social welfare.*

15

COMBINATORIAL PUBLIC PROJECTS

In this section we consider a first price auction for choosing a set of public projects and show its smoothness properties under different assumptions on the valuations of the players over the projects.

We then give an application of a simultaneous public good auction where the number of participants at each one is small. In most settings the smoothness of the mechanism implies a price of anarchy that grows with the number of players or more specifically the maximum number of players interested in any given project.

15.1 Item Bidding Mechanism for Public Projects

We consider the following formal setting: there are n bidders and m public projects. The mechanism wants to choose a set S of k projects to implement and each player i has a value $v_i : 2^{[m]} \to \mathbb{R}_+$ on the projects. We consider the following mechanism:

MECHANISM 11: Item-Bidding Mechanism for Combinatorial Public Projects.

Solicit bids b_{ij} from each player i for each project j;
For a project $j \in [m]$, let $B_j = \sum_{i \in [n]} b_{ij}$;
Pick the k projects with the highest B_j and let $S(b)$ be this set of projects;
Charge each player his sum of bids for the chosen projects $\sum_{j \in S(b)} b_{ij}$

We provide three smoothness theorems for the project bidding mechanism, according to the allowable class of bidder valuations.

Theorem 15.1.1. *For agents with arbitrary monotone valuations the Item-Bidding Mechanism is* $\left(\frac{1}{2}, n \cdot k\right)$*-smooth.*

Proof. Consider a valuation profile v and a bid profile b. Let $\textsc{Opt}(v)$ be the optimal set of projects for valuation profile v. Let $p_1(b)$ be the total bid of the highest valued project under bid profile b. Suppose that agent i switches to b_i' in which he draws a random bid t uniformly at random from $[0, \frac{v_i(\textsc{Opt}(v))}{k}]$ (i.e. with density $f(t) = \frac{k}{v_i(\textsc{Opt}(v))}$) and submits this random bid t on all the projects in $\textsc{Opt}(v)$. If $p_1(b) < t$ then the player gets all of the projects in $\textsc{Opt}(v)$ selected and hence gets a value of $v_i(\textsc{Opt}(v))$. The expected payment that he pays is at most his expected total bid, which is $k\frac{v_i(\textsc{Opt}(v))}{2k}$. By the quasi-linearity of utilities and the linearity of expectation, his expected utility under this deviation is at least:

$$\begin{aligned} U_i^{\mathcal{M}}(b_i', b_{-i}; v_i) &\geq \int_{p_1(b)}^{\frac{v_i(\textsc{Opt}(v))}{k}} v_i(\textsc{Opt}(v)) f(t) dt - \frac{v_i(\textsc{Opt}(v))}{2} \\ &= \int_{p_1(b)}^{\frac{v_i(\textsc{Opt}(v))}{k}} k \cdot dt - \frac{v_i(\textsc{Opt}(v))}{2} \\ &= \frac{v_i(\textsc{Opt}(v))}{2} - k \cdot p_1(b) \end{aligned}$$

By summing over all players and using the trivial fact that $p_1(b) \leq \sum_i P_i(b)$ we get the theorem. ∎

Theorem 15.1.2. *For agents with fractionally subadditive monotone valuations the Item-Bidding Mechanism is* $\left(\frac{1}{2}\left(1 - \frac{1}{e}\right), n\right)$*-smooth.*

Proof. To simplify the notation in the proof we will assume that k, the number of chosen projects, is even. Consider a valuation profile v and a bid profile b. Let $\textsc{Opt}(v)$ be the optimal set of projects for valuation profile v. Let $(v_{ij}^*)_{j\in[m]}$ be the representative additive for player i for set $\textsc{Opt}(v)$, i.e.

$$v_i(\textsc{Opt}(v)) = \sum_{j\in\textsc{Opt}(v)} v_{ij}^* = \max_{\ell\in L} \sum_{j\in\textsc{Opt}(v)} v_{ij}^\ell$$

Assume that items are reordered such that projects 1 to $k/2$ are the ones with the highest v_{ij}^* in the above representative additive valuation. Hence, by definition, $\sum_{j=1}^{k/2} v_{ij}^* \geq \frac{1}{2} v_i(\textsc{Opt}(v))$.

Suppose that agent i switches to b_i' in which for each project $j \in [1,\ldots,k/2]$ he draws an independent random number $\mathbf{t}_j$ with density $f_j(t_j) = \frac{1}{v_{ij}^*-t_j}$ and support $[0, v_{ij}^*(1-e^{-1})]$ and submits this random bid $\mathbf{t}_j$ on project $j \in \textsc{Opt}(v)$. He submits a 0 on any other project. Let $X(\mathbf{t}) \subseteq [m]$ be the set of projects that are chosen for some random draw $\mathbf{t} = (\mathbf{t}_j)_{j\in[m]}$ of the deviating bids of player i. Hence, a player's utility from the deviation is:

$$U_i^{\mathcal{M}}(b_i', b_{-i}; v_i) = \mathbb{E}_{\mathbf{t}}\left[v_i(X(\mathbf{t})) - \sum_{j\in X(t)} \mathbf{t}_j\right]$$

Using the fractionally subadditive property of the valuation we know that $v_i(X(\mathbf{t})) \geq \sum_{j\in X(\mathbf{t})} v_{ij}^*$. Thus:

$$U_i^{\mathcal{M}}(b_i', b_{-i}; v_i) \geq \mathbb{E}_{\mathbf{t}}\left[\sum_{j\in X(\mathbf{t})} v_{ij}^* - \mathbf{t}_j\right] = \sum_{j\in[m]} \mathbb{E}_{\mathbf{t}_j}\left[(v_{ij}^* - \mathbf{t}_j)\cdot \mathbf{1}_{j\in X(\mathbf{t})}\right]$$

Now, observe that for all $j \in [m]$, $\mathbf{t}_j \leq v_{ij}^*$, by the definition of the deviating bids. Thus, each term in the above sum is non-negative. Thus:

$$U_i^{\mathcal{M}}(b_i', b_{-i}; v_i) \geq \sum_{j=1}^{k/2} \mathbb{E}_{\mathbf{t}_j}\left[(v_{ij}^* - \mathbf{t}_j)\cdot \mathbf{1}_{j\in X(\mathbf{t})}\right]$$

Let $p_t(b)$ denote the t-th highest total bid under the initial bid profile b. For any $j \in \textsc{Opt}(v)$, if $\mathbf{t}_j > p_{k/2+1}(b)$ then project j is definitely selected, since player i is bidding non-positive on only $k/2$ projects and we know that the bids of the rest of the players exceed $p_{k/2+1}$ on at most $k/2$ projects. Thus we can lower bound each term in the above sum as follows:

$$\begin{aligned}
U_i^{\mathcal{M}}(b_i', b_{-i}; v_i) &\geq \sum_{j=1}^{k/2} \mathbb{E}_{\mathbf{t}_j}\left[(v_{ij}^* - \mathbf{t}_j) \cdot \mathbf{1}_{\mathbf{t}_j > p_{\frac{k}{2}+1}(b)}\right] \geq \sum_{j=1}^{k/2} \int_{p_{\frac{k}{2}+1}(b)}^{(1-e^{-1})v} (v_{ij}^* - t_j) f(t_j) dt_j \\
&= \left(1 - \frac{1}{e}\right) \sum_{j=1}^{k/2} v_{ij}^* - \sum_{j=1}^{k/2} p_{\frac{k}{2}+1}(b) \\
&\geq \left(1 - \frac{1}{e}\right) \frac{1}{2} v_i(\textsc{Opt}(v)) - \frac{k}{2} p_{\frac{k}{2}+1}(b) \\
&\geq \left(1 - \frac{1}{e}\right) \frac{1}{2} v_i(\textsc{Opt}(v)) - \sum_{j=1}^{k/2} p_t(b) \\
&\geq \left(1 - \frac{1}{e}\right) \frac{1}{2} v_i(\textsc{Opt}(v)) - \sum_{i \in [n]} P_i(b)
\end{aligned}$$

Summing over all agents we get the theorem. ■

Theorem 15.1.3. *When bidder valuations are unit-demand then the Item-Bidding Mechanism is $(1 - e^{-1}, \frac{n}{k})$-smooth.*

Proof. Consider a valuation profile v and a bid profile b. Let v_{ij}^* be the maximum valued project of each player: i.e. $v_{ij}^* = \max_{j \in [m]} v_{ij}$. Suppose that each player switches to the following randomized bid: he draws a random bid $\mathbf{t}$ from distribution with density $f(t) = \frac{1}{v_{ij}^* - t}$ and support $[0, (1 - e^{-1})v_{ij}^*]$. Then he submits this random bid on project $j^* = \arg\max_{j \in [m]} v_{ij}$ and submits a 0 on all other projects. Let $p_r(b)$ be the total bid of the r-th highest valued project under the initial bid profile b. For any $j \in \textsc{Opt}(v)$, if $\mathbf{t}_j > p_k(b)$ then project j is definitely selected. Thus a bidders utility from the deviation is at least:

$$U_i^{\mathcal{M}}(b_i', b_{-i}; v_i) \geq \int_{p_k(b)}^{(1-e^{-1})v_{ij}^*} (v_{ij}^* - t) f(t) dt = (1 - e^{-1})v_{ij}^* - p_k(b)$$

By summing up over all players and using the facts that $\sum_i v_{ij}^* \geq \sum_i v_i(\textsc{Opt}(v))$ and $p_k(b) \leq \frac{1}{k} \sum_{t=1}^{k} p_t(b) = \frac{1}{k} \sum_i P_i(b)$ (by the definition of prices) we get the theorem. ■

15.2 Local Public Good Auctions in Networks

Consider a social network setting where players bid for facilities to be placed on nodes in a social network. Each node is an agent and when a facility is placed on a node then all of the neighbors of the node can use it. There exists a set of facilities F_u that can be placed on each node u (let F_u contain also the empty facility for the case where no facility is built). Now we assume that auctioneers run a public good auction on each node to decide which facility they are going to place. Specifically, he asks from the node and its neighbors to submit a bid for each possible facility. Then he is going to choose the facility that received the highest sum of bids and charge each player his bid for the chosen facility.

By Theorem 15.1.3, each mechanism is $(1-e^{-1}, d_i)$-smooth where d_i is the degree of the node that is auctioned. Now our framework shows that if we run simultaneous such auctions and the valuation of a player is a fractionally subadditve valuation over the facilities placed on his neighboring nodes, then the overall social welfare of this game will be at least $(1-e^{-1})\frac{1}{D}$ of the optimal, where $D = \max_i d_i$, i.e. the price of anarchy is at most $\frac{D}{1-e^{-1}}$. Similarly, one could imagine of a setting where facilities are not placed on nodes of the graph but rather on edges of it or on hyper-edges in a hyper-graph that tries to model groups of interested agents. In such settings our framework implies that the above simultaneous local public good mechanism has price of anarchy at most $\frac{k}{1-e^{-1}}$, where k is the size of the hyper-edge.

BIBLIOGRAPHY

[1] Ittai Abraham, Moshe Babaioff, Shaddin Dughmi, and Tim Roughgarden. Combinatorial auctions with restricted complements. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 3–16, New York, NY, USA, 2012. ACM.

[2] Zoë Abrams, Arpita Ghosh, and Erik Vee. Cost of conciseness in sponsored search auctions. WINE'07, pages 326–334, 2007.

[3] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, FOCS '95, pages 322–, Washington, DC, USA, 1995. IEEE Computer Society.

[4] Robert J. Aumann. Mixed and behavior strategies in infinite extensive games. *Advances in Game Theory, Annals of Mathematics Study 52*, pages 627–650, 1964.

[5] Moshe Babaioff, Brendan Lucier, Noam Nisan, and Renato Paes Leme. On the efficiency of the walrasian mechanism. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 783–800, New York, NY, USA, 2014. ACM.

[6] Junjik Bae, Eyal Beigman, Randall Berry, Michael L. Honig, and Rakesh Vohra. On the efficiency of sequential auctions for spectrum sharing. In *Proceedings of the First ICST international conference on Game Theory for Networks*, GameNets'09, pages 199–205, Piscataway, NJ, USA, 2009. IEEE Press.

[7] Dirk Bergemann and Stephen Morris. Correlated Equilibrium in Games with Incomplete Information. Cowles Foundation Discussion Papers 1822, Cowles Foundation for Research in Economics, Yale University, October 2011.

[8] 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*, SODA '11, pages 700–709. SIAM, 2011.

[9] Sushil Bikhchandani. Auctions of heterogeneous objects. *Games and Economic Behavior*, 26(2):193 – 220, 1999.

[10] Sushil Bikhchandani, Sven de Vries, James Schummer, and Rakesh V. Vohra. An ascending vickrey auction for selling bases of a matroid. *Operations Research*, 59(2):400–413, 2011.

[11] A. Blum and Y. Mansour. Learning, regret minimization, and equilibria. In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*, chapter 4, pages 4–30. Cambridge University Press, 2007.

[12] 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*, STOC '08, pages 373–382, New York, NY, USA, 2008. ACM.

[13] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, Maria Kyropoulou, Brendan Lucier, Renato Paes Leme, and Éva Tardos. Bounding the inefficiency of outcomes in generalized second price auctions. *Journal of Economic Theory*, (0):–, 2014.

[14] Ioannis Caragiannis and Alexandros Voudouris. Welfare guarantees for proportional allocations. 2014.

[15] George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian combinatorial auctions. In *Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part I*, ICALP '08, pages 820–832, Berlin, Heidelberg, 2008. Springer-Verlag.

[16] EdwardH. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.

[17] Bart de Keijzer, Evangelos Markakis, Guido Schfer, and Orestis Telelis. Inefficiency of standard multi-unit auctions. In HansL. Bodlaender and GiuseppeF. Italiano, editors, *Algorithms ESA 2013*, volume 8125 of *Lecture Notes in Computer Science*, pages 385–396. Springer Berlin Heidelberg, 2013.

[18] Shahar Dobzinski, Ron Lavi, and Noam Nisan. Multi-unit auctions with budget limits. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '08, pages 260–269, Washington, DC, USA, 2008. IEEE Computer Society.

[19] Shahar Dobzinski and Renato Paes Leme. Efficiency guarantees in auctions with budgets. *CoRR*, abs/1304.7048, 2013.

[20] P. Duetting, M. Henzinger, and M. Starnberger. Auctions with Heterogeneous Items and Budget Limits. *ArXiv e-prints*, September 2012.

[21] Benjamin Edelman. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, March 2007.

[22] Richard Engelbrecht-Wiggans and Robert J. Weber. An example of a multi-object auction game. *Management Science*, 25(12):pp. 1272–1277, 1979.

[23] Uriel Feige. On maximizing welfare when utility functions are subadditive. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, STOC '06, pages 41–50, New York, NY, USA, 2006. ACM.

[24] Uriel Feige, Michal Feldman, Nicole Immorlica, Rani Izsak, Brendan Lucier, and Vasilis Syrgkanis. A unifying hierarchy of valuations with complements and substitutes. Working paper, 2014.

[25] 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*, STOC '13, pages 201–210, New York, NY, USA, 2013. ACM.

[26] Michal Feldman, Nicole Immorlica, Brendan Lucier, Tim Roughgarden, and Vasilis Syrgkanis. The price of anarchy in large games. Working paper, 2014.

[27] Michal Feldman, Brendan Lucier, and Vasilis Syrgkanis. Limits of efficiency in sequential auctions. In Yiling Chen and Nicole Immorlica, editors, *Web and Internet Economics*, volume 8289 of *Lecture Notes in Computer Science*, pages 160–173. Springer Berlin Heidelberg, 2013.

[28] Amos Fiat, Stefano Leonardi, Jared Saia, and Piotr Sankowski. Single valued combinatorial auctions with budgets. In *Proceedings of the 12th ACM conference on Electronic commerce*, EC '11, pages 223–232, New York, NY, USA, 2011. ACM.

[29] Franoise Forges. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision*, 35(3):277–310, 1993.

[30] Ian L. Gale and Mark Stegeman. Sequential auctions of endogenously valued objects. *Games and Economic Behavior*, 36(1):74–103, July 2001.

[31] Gagan Goel, Vahab Mirrokni, and Renato Paes Leme. Polyhedral clinching auctions and the adwords polytope. In *Proceedings of the 44th symposium on Theory of Computing*, STOC '12, pages 107–122, New York, NY, USA, 2012. ACM.

[32] Jacob K. Goeree. Bidding for the future: signaling in auctions with an aftermarket. *Journal of Economic Theory*, 108(2):345 – 364, 2003.

[33] Theodore Groves. Incentives in teams. *Econometrica*, 41(4):pp. 617–631, 1973.

[34] John C. Harsanyi. Games with Incomplete Information Played by 'Bayesian' Players. *Management Science*, 14:159–182, 320–324, 486–502, 1967-68.

[35] Sergiu Hart and Andreu Mas-Colell. *Simple Adaptive Strategies: From Regret-matching to Uncoupled Dynamics*. World Scientific Publishing Co., Inc., River Edge, NJ, USA, 2012.

[36] A. Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. Non-price equilibria in markets of discrete goods. In *Proceedings of the 12th ACM conference on Electronic commerce*, EC '11, pages 295–296, New York, NY, USA, 2011. ACM.

[37] Z. Huang, N. R. Devanur, and D. Malec. Sequential Auctions of Identical Items with Budget-Constrained Bidders. *ArXiv e-prints*, September 2012.

[38] Ramesh Johari and John N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):pp. 407–435, 2004.

[39] ToddR. Kaplan and Shmuel Zamir. Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economic Theory*, 50(2):269–302, 2012.

[40] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *The Journal of the Operational Research Society*, 49(3):pp. 237–252, 1998.

[41] Frank Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8(1):33–37, 1997.

[42] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th annual conference on Theoretical aspects of computer science*, STACS'99, pages 404–413, Berlin, Heidelberg, 1999. Springer-Verlag.

[43] Vijay Krishna. *Auction Theory*. Academic Press, March 2002.

[44] Jon Lee, Maxim Sviridenko, and Jan Vondrák. Submodular maximization over multiple matroids via generalized exchange properties. *Math. Oper. Res.*, 35(4):795–806, November 2010.

[45] 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.

[46] B. Lucier and A. Borodin. Price of anarchy for greedy auctions. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '10, pages 537–553, Philadelphia, PA, USA, 2010. Society for Industrial and Applied Mathematics.

[47] Brendan Lucier and Renato Paes Leme. Gsp auctions with correlated types. In *Proceedings of the 12th ACM conference on Electronic commerce*, EC '11, pages 71–80, New York, NY, USA, 2011. ACM.

[48] Brendan Lucier, Yaron Singer, Vasilis Syrgkanis, and Éva Tardos. Equilibrium in combinatorial public projects. In Yiling Chen and Nicole Immorlica, editors, *Web and Internet Economics*, volume 8289 of *Lecture Notes in Computer Science*, pages 347–360. Springer Berlin Heidelberg, 2013.

[49] Brendan Lucier and Vasilis Syrgkanis. Greedy auctions are smooth. Working Paper. 2014.

[50] Evangelos Markakis and Orestis Telelis. Uniform price auctions: Equilibria and efficiency. In Maria Serna, editor, *Algorithmic Game Theory*, Lecture Notes in Computer Science, pages 227–238. Springer Berlin Heidelberg, 2012.

[51] Eric Maskin and John Riley. Asymmetric auctions. *The Review of Economic Studies*, 67(3):pp. 413–438, 2000.

[52] Paul R Milgrom and Robert J Weber. A Theory of Auctions and Competitive Bidding. *Econometrica*, 50(5):1089–1122, September 1982.

[53] Paul R Milgrom and Robert J Weber. A Theory of Auctions and Competitive Bidding, II. *The Economic Theory of Auctions*, 1999.

[54] Uri Nadav and Tim Roughgarden. The limits of smoothness: a primal-dual framework for price of anarchy bounds. In *Proceedings of the 6th international conference on Internet and network economics*, WINE'10, pages 319–326, Berlin, Heidelberg, 2010. Springer-Verlag.

[55] S. S. Oren and M. H. Rothkopf. Optimal Bidding in Sequential Auctions. *Operations Research*, 23(6):1080–1090, November 1975.

[56] Renato Paes Leme, Vasilis Syrgkanis, and Éva Tardos. Sequential auctions and externalities. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '12, pages 869–886. SIAM, 2012.

[57] Renato Paes Leme and Eva Tardos. Pure and bayes-nash price of anarchy for generalized second price auction. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, FOCS '10, pages 735–744, Washington, DC, USA, 2010. IEEE Computer Society.

[58] Gustavo E. Rodriguez. Sequential auctions with multi-unit demands. *The B.E. Journal of Theoretical Economics*, 9(1):45, 2009.

[59] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, STOC '09, pages 513–522, New York, NY, USA, 2009. ACM.

[60] Tim Roughgarden. The price of anarchy in games of incomplete information. *SIGecom Exch.*, 11(1):18–20, June 2012.

[61] Tim Roughgarden and Florian Schoppmann. Local smoothness and the price of anarchy in atomic splittable congestion games. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 255–267. SIAM, 2011.

[62] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, March 2002.

[63] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency. Matroids,*

trees, stable sets. Number v. 2 in Algorithms and Combinatorics. Springer, 2003.

[64] Jeroen M. Swinkels. Efficiency of large private value auctions. *Econometrica*, 69(1):pp. 37–68, 2001.

[65] Vasilis Syrgkanis. Bayesian games and the smoothness framework. *CoRR*, abs/1203.5155, 2012.

[66] Vasilis Syrgkanis and Éva Tardos. Bayesian sequential auctions. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 929–944, New York, NY, USA, 2012. ACM.

[67] Vasilis Syrgkanis and Éva Tardos. Composable and efficient mechanisms. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 211–220, New York, NY, USA, 2013. ACM.

[68] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

[69] Robert J. Weber. Multiple-Object Auctions. Discussion Papers 496, Northwestern University, Center for Mathematical Studies in Economics and Management Science, August 1981.

APPENDIX A

COLLECTION OF LOWER BOUNDS

A.1 93% Lower Bound for Bayes-Nash Equilibria of the Asymmetric the First Price Auction

EXAMPLE A.1.1. (A 93% Lower Bound for the Asymmetric FPA) Player 1 has value v_1 with probability p_1 and 0 with probability $1-p_1$. Player 2 has value v_2 with probability p_2 and 0 with probability $1-p_2$. We assume that when some player has value 0 he doesn't even bid. Moreover, we assume that $v_2 \geq v_1$.

We need to determine the equilibrium strategies of player 1 and player 2 when they have positive value. The strategies are going to be mixed. Let $F_1(t), F_2(t)$ be the CDFs of the mixed strategies of the two bidders and U_1, U_2 be their supports. The supports are going to be identical $U_1 = U_2 = [0, \bar{b}]$ (the closedness of the upper and lower bounds might vary). The expected utility of player 1 and player 2 from bidding any $t \in [0, \bar{b}]$ is:

$$
\begin{aligned}
u_1(t) = &\ (1 - p_2 + p_2 F_2(t))(v_1 - t) \\
u_2(t) = &\ (1 - p_1 + p_1 F_1(t))(v_2 - t)
\end{aligned}
$$

We assume that the auction can differentiate between a player with positive value bidding zero and a player with zero value. Moreover, a player with positive value bidding zero always wins against a zero value player bidding zero. We will assume that player 1 bids with some positive probability 0, when he has a positive value, while player 2 doesn't have an atom at 0 and in fact 0 is not in the support of player 2's strategy but simply the lower bound. Since the expected utility must be constant across the support we have:

$$\begin{aligned}(1-p_2+p_2F_2(t))(v_1-t) &= \; v_1-\bar{b}=(1-p_2)v_1\\ (1-p_1+p_1F_1(t))(v_2-t) &= \; v_2-\bar{b}=v_2-v_1+(1-p_2)v_1=v_2-p_2v_1\end{aligned}$$

Thus $\bar{b}=v_1p_2$. Which then gives:

$$\begin{aligned}F_1(t) &= \frac{v_2-p_2v_1-(1-p_1)(v_2-t)}{p_1(v_2-t)}=\frac{p_1v_2-p_2v_1+(1-p_1)t}{p_1(v_2-t)}\\ F_2(t) &= \frac{1-p_2}{p_2}\frac{t}{v_1-t}\end{aligned}$$

with, $F_1(0)=\frac{p_1v_2-p_2v_1}{p_1v_2}$ and $F_2(0)=0$ and $F_1(\bar{b})=F_2(\bar{b})=1$. Optimizing over v_1, v_2, p_1, p_2, we get that the expected welfare is approximately $.93$ of the expected optimal welfare, when $v_2=1, v_1\approx 0.57$, $p_2=0.75$ and $p_1\to 0$. ■

A.2 Tight $\frac{e}{e-1}$ Lower Bound for Bayes-Nash of First Price Auction with Correlated Values

EXAMPLE A.2.1. ($\frac{e}{e-1}$ Lower Bound of FPA with correlated values) Consider the following setting: there are three players. Player one has value deterministically 1. The value of players two and three is perfectly correlated and identical. Moreover, their common value v is drawn from a probability distribution with density function $F(t)=\frac{1}{e}\frac{1}{1-t}$ and support $\left[0,\left(1-\frac{1}{e}\right)v\right]$. Thus the distribution has an atom at 0 with mass $\frac{1}{e}$ and has a well defined density $f(t)=\frac{1}{e}\frac{1}{(1-t)^2}$ for any $t\in\left(0,1-\frac{1}{e}\right]$.

The following is an equilibrium: Player one bids 0 (we assume that the auction favors player one in case of ties and set his bid to 0), while players two and three bid truthfully. This is obviously an equilibrium for the pair of correlated players.

Player one wins only when the pair of players has zero value, which happens with probability $\frac{1}{e}$. Moreover, any bid in the region $\left(0, 1-\frac{1}{e}\right]$, yields utility: $(1-t)F(t) = \frac{1}{e}$, while any bid above $1-\frac{1}{e}$, yields strictly lower utility. Thus it is an equilibrium for player 1 too.

The optimal social welfare is obviously 1, while the expected welfare at equilibrium is:

$$\frac{1}{e} \cdot 1 + \int_0^{1-1/e} \frac{1}{e} \frac{t}{(1-t)^2} dt = \frac{1}{e} + 1 - \frac{2}{e} = 1 - \frac{1}{e}$$

■

A.3 Tight $\frac{e}{e-1}$ Lower Bound for CCE of First Price Auction

We describe a continuous bidding example where we allow for the auctioneer to set any tie-breaking rule. We point out that if we discretize the bid space, then an example even under random tie-breaking that approaches $1 - 1/e$ can also be constructed that approximates the one that we describe here.

EXAMPLE A.3.1. (Tight Lower Bound for CCE of FPA in complete information) Consider a single-item auction among two bidders. Bidder 1 has value v and bidder 2 has value 0 for the item and this is common knowledge. We assume that the auctioneer breaks ties in favor of player 1 for any positive bid and in favor of bidder 2 when a tie occurs at a bid of 0.

We argue that the following is a coarse correlated equilibrium of the game:

a random number t is drawn from distribution with cummulative density function $F(t) = \frac{v}{e}\frac{1}{v-t}$ and support $\left[0, \left(1 - \frac{1}{e}\right) v\right]$. Thus the distribution has an atom at 0 with mass $\frac{1}{e}$ and has a well defined density $f(t) = \frac{v}{e}\frac{1}{(v-t)^2}$ for any $t \in \left(0, \left(1 - \frac{1}{e}\right) v\right]$. Then both players bid t.

Thus player 1 wins with probability equal to the probability that t is positive, which is $1-\frac{1}{e}$ and player 2 wins with probability equal to the point mass at 0, i.e. $1/e$. Thus the social welfare of this correlated distribution of bids is $\left(1 - \frac{1}{e}\right) v$, while obviously the optimal social welfare is v.

The expected utility of player 2 is 0 since whenever $t > 0$ he loses and whenever $t = 0$ he wins and pays 0. Thus it is trivially a coarse correlated equilibrium from his perspective.

The expected utility of player 1 can be computed as follows:

$$\begin{aligned} \mathbb{E}_t[u_1(t,t)] &= \int_0^{\left(1-\frac{1}{e}\right)v} (v-t)f(t)dt = \int_0^{\left(1-\frac{1}{e}\right)v} (v-t)\frac{v}{e}\frac{1}{(v-t)^2}dt \\ &= \frac{v}{e} \cdot [-\log(v-t)]_0^{\left(1-\frac{1}{e}\right)v} = \frac{v}{e} \end{aligned}$$

Last we need to check that the expected utility of player 1 from switching to any other fixed bid is at most $\frac{v}{e}$. We actually show that for any bid $b \in \left(0, \left(1 - \frac{1}{e}\right) v\right]$ the expected utility for switching to bidding b all the time is equal to $\frac{v}{e}$:

$$\mathbb{E}_t[u_1(b,t)] = (v-b) \cdot F(b) = (v-b)\frac{v}{e}\frac{1}{v-b} = \frac{v}{e} \tag{A.1}$$

In addition, it is easy to see that any bid above $\left(1 - \frac{1}{e}\right) v$ yields utility strictly smaller than $\frac{v}{e}$. ■

A.4 $4/3$ Lower bound for Bayes-Nash of All-Pay Auction with Correlated Values

EXAMPLE A.4.1. (All-pay auction lower bound) Consider a single-item all pay auction among three players. Player 1 has value 1. Players 2 and 3 have a common value which is 0 with probability p and $1-p$ with probability $1-p$.

At equilibrium, player 1 bids 0, while players 2 and 3 bid uniformly in $(0, 1-p]$, when they have value $1-p$ and bid 0 otherwise (we assume that the auction favors player 1 in case of ties).

The utility of player 1 is p, while his utility from any bid $t \in (0, 1-p]$ is: $p + (1-p)\frac{t}{1-p} - t = p$. The utility of players 2 and 3 is 0 for any bid in the support of their equilibrium strategy, conditional on having positive value.

The optimal welfare is obviously 1. The expected welfare at equilibrium is:

$$p \cdot 1 + (1-p) \cdot (1-p) \tag{A.2}$$

which is $3/4$, when $p = 1/2$. ■

A.5 $8/7$ Lower bound for mixed Nash of All-Pay Auction

EXAMPLE A.5.1. (Mixed Nash All-pay auction lower bound) Consider a single-item all pay auction among two players. Player 1 has value 1 while player 2 has value v.

At equilibrium, player 1 bids uniformly in $(0, v]$, while player 2 submits a bid with cummulative density function $G_2(t) = 1 - v + t$, i.e. player 2 bids 0

with probability $1 - v$ and otherwise bids uniformly between 0 and his value. The utility of player 1 is $1 - v$ for any bid in $(0, v]$. Player 2's utility is 0 for any bid in $(0, v]$.

The optimal welfare is obviously 1. The expected welfare at equilibrium is:

$$1 \cdot \int_0^v (1 - v + t) \cdot \frac{1}{v} dt + v \cdot \int_0^v \frac{t}{v} \cdot 1 dt = 1 - \frac{v}{2} + v \cdot \frac{v}{2} \tag{A.3}$$

which is $7/8$, when $p = 1/2$. ■

A.6 Inefficiency of GSP with per-impression values grows with slots

Theorem A.6.1. *The* POA *of GSP when bidders also have values per-impression is at least m (where m is the number of slots).*

Proof. Consider the following instance: we have m slots and m bidders that have a value of $v_p = 1$ per click and no value for the impression. Additionally there are m bidders that have a value of $v = H$ just for the impression and no value per click. The slots have click-through-rates: $a_1 = H$ and $a_i = 1$ for all $i > 1$. We consider the GSP auction where each bidder reports a per-click bid b_i and a bidder at slot s pays: $a_s \cdot b_{s+1}$.

We also assume that when bidders are tied, then the impression bidders are ranked to our favor, i.e. they are ranked first.

Now we argue that the following is an equillibrium: all per-click bidders bid 1 and all per-impression bidders bid 1^-. All players get utility 0. All the per-impression bidders cannot get better utility by deviating. All the per-click bidders if they bid 1 or higher they get the first slot and pay $a_1 \cdot v_p = H$. Thus their utility is also 0 from this deviation. Thus the latter is an equilibrium.

Under this equilibrium the welfare is $H + m - 1$. In the optimal all the impression bidders get all the slots and get welfare $H \cdot m$. For $H = m$, the price of anarchy is $O(m)$. ■