
Incentives in Dynamic Markets

A dissertation submitted towards the degree Doctor of Natural
Science of the Faculty of Mathematics and Computer Science of
Saarland University

by Bojana Kodric

Saarbrücken / 2017

Day of Colloquium: 23. July 2018
Dean of the Faculty: Prof. Dr. Sebastian Hack

Chair of the Committee: Prof. Dr. Markus Bläser
Reporters

First reviewer: Prof. Dr. Martin Hofer
Second reviewer: Prof. Dr. Dr. h.c. mult. Kurt Mehlhorn
Academic Assistant: Dr. Marvin Künnemann

Abstract

Abstract In this thesis, we consider a variety of combinatorial optimization problems within a common theme of *uncertainty and selfish behavior*.

In our first scenario, the input is collected from selfish players. Here, we study extensions of the so-called smoothness framework for mechanisms, a very useful technique for bounding the inefficiency of equilibria, to the cases of varying mechanism availability and participation of risk-averse players. In both of these cases, our main results are general theorems for the class of (λ, μ) -smooth mechanisms. We show that these mechanisms guarantee at most a (small) constant factor performance loss in the extended settings.

In our second scenario, we do not have access to the exact numerical input. Within this context, we explore combinatorial extensions of the well-known secretary problem under the assumption that the incoming elements only reveal their ordinal position within the set of previously arrived elements. We first observe that many existing algorithms for special matroid structures maintain their competitive ratio in the ordinal model. In contrast, we provide a lower bound for algorithms that are oblivious to the matroid structure. Finally, we design new algorithms that obtain constant competitive ratios for a variety of combinatorial problems.

Zusammenfassung In dieser Dissertation betrachten wir eine Auswahl kombinatorischer Optimierungsprobleme, denen allen das Thema *der Unsicherheit und des egoistischen Verhaltens* zugrunde liegt.

In unserem ersten Szenario wird die Eingabe von egoistischen Spielern empfangen. Hier studieren wir Erweiterungen des sogenannten „Smoothness-Frameworks“ für Mechanismen, einer sehr nützlichen Technik um die Ineffizienz von Gleichgewichten zu beschränken. Wir betrachten den Fall von variierender Mechanismen-Verfügbarkeit und die Teilnahme von risikoaversen Spielern. In beiden Fällen sind unsere Hauptresultate allgemeine Sätze für Mechanismen, die „ (λ, μ) -smooth“ sind. Wir zeigen, dass diese Mechanismen höchstens einen (kleinen) konstanten Faktor Effizienzverlust in den verallgemeinerten Situationen garantieren.

In unserem zweiten Szenario haben wir keinen Zugang zur exakten (numerischen) Eingabe. In diesem Kontext untersuchen wir kombinatorische Erweiterungen des wohlbekanntes Sekretärinnenproblems unter der Annahme, dass die eingehenden Elemente nur ihren Rang in der Ordnung der bisher empfangenen Elemente verraten. Wir beobachten zunächst, dass viele existierende Algorithmen für spezielle Matroidstrukturen ihre Kompetitivität in diesem Ordnungsmodell behalten. Im Gegensatz hierzu geben wir eine untere Schranke für Algorithmen an, die keine Matroidstrukturen erkennen können. Schließlich entwerfen wir neue Algorithmen mit konstanter Kompetitivität für eine Auswahl von kombinatorischen Problemen.

Acknowledgments

First of all, I would like to thank my advisor Martin Hofer for his support throughout my Ph.D., both on and off duty. I never expected the advisor-student relationship to be such an accurate reenactment of the customary growing up and I do hope I wasn't too naughty in my teenage years.

I am also very grateful to Thomas Kesselheim for his kind guidance. It was a pleasure experiencing an exceedingly clear way of thinking and posing questions, in a scientific world that is often cluttered with information and deprived of common sense.

Finally, I would like to thank Kurt Mehlhorn for accepting me to his group in the first place, for providing a lovely work environment, for always being there – not just for me but for all of us.

Contents

1	Overview	1
1.1	Part I: Extensions of the Smoothness Framework for Mechanisms	3
1.2	Part II: Combinatorial Secretary Problems with Ordinal Information . .	5
I	Extensions of the Smoothness Framework for Mechanisms	7
2	Introduction to Part I	9
2.1	Our Contribution	11
2.2	Related Work	13
3	Notation and Preliminaries	17
3.1	Setting	17
3.2	Solution Concepts and Benchmarks	19
3.3	Smoothness Framework	20
4	Risk-Averse Agents	21
4.1	Model	21
4.2	Single-Item Auctions in the Quasilinear Setting	22
4.3	Smoothness Beyond Quasilinear Utilities	23
4.4	Quasilinear Often Implies Risk-Averse Smoothness	28
4.5	Unbounded Price of Anarchy for All-Pay Auctions	32
4.6	Variance-Aversion Model	38
5	Simultaneous Composition of Mechanisms with Admission	41
5.1	Model	41
5.2	Composition with Independent Admission	43
5.3	Composition with Everybody-or-Nobody Admission	47
5.4	A Lower Bound for General XOS Functions	53
5.5	Applications beyond Auctions	55
5.6	Extension to Changing Unit-Demand Functions	57
II	Combinatorial Secretary Problems with Ordinal Information	59
6	Introduction to Part II	61
6.1	Our Contribution	62
6.2	Related Work	63
6.3	Notation and Preliminaries	64

7	Matroids	67
7.1	Submodular Matroids	67
7.2	A Lower Bound	68
8	Matching, Packing and Independent Set	73
8.1	Bipartite Matching	73
8.2	Packing	75
8.3	Matching in General Graphs	77
8.4	Independent Set and Local Independence	78
9	Conclusion and Open Problems	81
	List of Figures	83
	List of Tables	85
	List of Algorithms	87
	List of Acronyms	89
	Bibliography	91

CHAPTER 1

Overview

Algorithmic mechanism design strives towards designing games that have both good game-theoretical and algorithmic properties. It became a popular topic in the computer science community only in the last two decades, while traditionally mechanism design was extensively studied in economics. The algorithmic and computational aspects stem from a mechanism simply being an algorithm: upon receiving players' private information as input, it outputs an outcome for all of the participants. The players, on the other hand, are assumed to be selfish rational agents that are interested only in achieving the best possible outcome for themselves. To this end, they might misreport their private information. The goal of a mechanism, as opposed to a regular algorithm, is to align the incentives of the players such as to reach a desired objective in an equilibrium. An equilibrium is generally defined as a state in which no player can improve his outcome by changing his part of the input and players' preferences are modeled via valuation functions over the possible outcomes. It is commonly assumed that the utility a player experiences from an outcome is equal to his valuation for the outcome minus the payment imposed by the mechanism (if any). Payments allow the mechanism to encourage players to report truthfully, through maximizing their utilities for truthful actions.

A simple example of a truthful mechanism is the second price auction.¹ In a second-price auction, the highest bidder wins the item and pays the second highest bid. It is easy to check that submitting a bid equal to his valuation is each player's utility maximizing strategy. Second price auctions stay truthful also if the auctioneer in addition sets a reserve price such that the item is sold only if the bid of the highest bidder exceeds the reserve. The highest bidder has to, in this case, pay the maximum of the second highest bid and the reserve price.² An example of a non-truthful mechanism is the first price auction, where the highest bidder wins the item and pays his bid. If a player reports his valuation for obtaining the item correctly, his utility is zero, independently of whether he wins or loses.

The two commonly studied objectives in mechanism design are social welfare maximization and revenue maximization. Social welfare maximization is concerned with the sum of all the valuations in an equilibrium outcome, while revenue maximization aims at maximizing the sum of the payments. The second-price auction is maximizing social welfare, while a second-price auction with an accordingly set reserve maximizes revenue.

The domain of algorithmic mechanism design has a wide range of applications. Some of the most prominent real world examples are online auctions and exchanges, online advertising, and search engine's page ranking. These applications occur on a daily basis

¹In fact, the second price auction is a special case of the more general Vickrey-Clark-Groves mechanism (VCG), applied to the case of assigning one item.

²Second-price auction with reserve price is the application of the more general optimal mechanism by Myerson [74] to the case of a single item auction.

and constitute a multi-billion-dollar industry.

In traditional mechanism design, one of the most fundamental principles is the revelation principle. It states that if there exists a mechanism that implements a social choice function, then it is possible to implement the same function by an incentive-compatible mechanism (i.e., by a mechanism that allows every participant to achieve the best outcome for themselves by truthfully reporting their private information). The two mechanisms have the same equilibrium outcome and the same equilibrium payoffs. This postulate enabled researchers to restrict their attention solely to truthful mechanisms and made incentive compatibility into something of a standard requirement.

In the recent years, it has become apparent that truthful mechanism design is inherently complicated from an algorithmic point of view. The question of designing a mechanism can indeed be viewed as a standard optimization problem, where truthfulness is imposed as an additional constraint. The optimization itself, however, is very often computationally intractable and the resulting mechanism complex and impractical. Truthful mechanisms are, in fact, very rarely used in practice even in the tractable cases. For instance, it would be possible to use the celebrated VCG for the sponsored search application. Instead, simple and non-truthful procedures are used to allocate ads on search result pages.

A very important example of a non-truthful mechanism that is widely used is the generalized second-price auction (GSP). It is employed mainly in the context of keyword auctions, where sponsored search slots are sold on an auction basis. In GSP each bidder places a single bid. The highest bidder then wins the first sponsored search slot and pays the second highest bid, the second highest bidder wins the second slot and pays the third highest bid, and so on. This mechanism is used by Google's AdWords technology and has evolved into Google's main source of revenue, which hints to the fact that we are indeed dealing with a multi-billion auction market.

A recent trend in algorithmic mechanism design is, therefore, to study non-truthful and conceptually "simple" mechanisms for allocation in markets and their inherent loss in system performance. If one, in addition, wants to capture dynamic aspects, e.g., when agents arrive and depart over time, as is often the case with practically used mechanisms and specifically also with the multi-billion-dollar example of sponsored search, one needs to resort to online mechanism design. Once the online aspect is introduced, however, problems usually become much harder. For instance, if we consider the online version of a single-item auction, we already arrive at an algorithmic problem of selecting the maximum element in a worst case sequence that admits no non-trivial online approximation algorithm. The classical truthfulness objective is, not surprisingly, also harder to reach in the online setting. Both of these difficulties become less present when one assumes that the agents are arriving in a random order. In this case, the algorithmic problem corresponding to the single-item auction translates directly into the well-known secretary problem. Assuming random arrival, we are left with the problem of choosing the maximum element in a randomly ordered sequence. The (optimal) solution, rejecting all the elements in a sampling phase and then accepting the first one that has a higher value than any element seen so far [35], can be interpreted as a posted-price mechanism that both gives a constant approximation to the optimal social welfare and is truthful.

In fact, there is a rich interplay between secretary problems and online mechanism design. Algorithms for secretary problems can be directly transformed into truthful

online mechanisms which are constant-competitive for agents with random arrival order. At the same time, the goal of designing online mechanisms with various combinatorial constraints has led to the formulation and solution of new secretary problems that are interesting in their own right.

In this thesis, we consider a variety of combinatorial optimization problems within a common theme of *uncertainty and selfish behavior*. All of the considered problems can be seen from the perspective of studying incentives in dynamic markets, where the objective is to maximize social welfare either via designing truthful mechanisms or by pointing to non-truthful mechanisms with good performance guarantees. Our work can be grouped into two scenarios:

- (1) In the first scenario, the input is collected from selfish agents who might misreport their private information in order to achieve a better outcome for themselves. Alternatively, the format of the interaction with the agents does not even allow them to fully express their preferences. This scenario obviously falls within the domain of Mechanism Design. More precisely, we investigate extensions of the smoothness framework for mechanisms, a very useful technique for bounding the inefficiency of equilibria of the induced games, to the cases of varying mechanism availability and participation of risk-averse players.
- (2) In the second scenario, we assume that access to the exact numerical input is not possible. We nevertheless want to design algorithms with provable performance guarantees with respect to the optimal solution. This is motivated by problems where it is either generally difficult or even impossible to assign exact numerical values to the elements in the input or we believe there might be imprecisions in the values that we are provided with. Within this Online Mechanism Design scenario, we explore various extensions of the well-known secretary problem under the assumption that the incoming elements only reveal their ordinal position with respect to all elements seen so far, instead of their numerical value.

In what follows, we expand upon the specific problems considered in this thesis and give a high-level overview of the results. For more details, see Chapter 2 and Chapter 6.

1.1 Part I: Extensions of the Smoothness Framework for Mechanisms

The study of truthful mechanisms is a classic branch of microeconomics and has resulted in a variety of fundamental results, such as VCG for social welfare maximization or Myerson's revenue-optimal auctions. Strikingly, these techniques are only very rarely used in practice, as they often involve heavy algorithmic machinery, complicated allocation techniques, or other hurdles to easy and transparent implementation. This has led to the study of non-truthful and conceptually "simple" mechanisms, in which bidders might have the opportunity to gain from non-truthful bids.

The approach of formally analyzing "simple" mechanisms is to study the induced game among the bidders and bound the quality of (possibly manipulated) outcomes in equilibrium. In a seminal paper, Syrgkanis and Tardos [85] propose a general framework for bounding social welfare of these equilibria, based on a so-called "smoothness"

technique. In Part I of this thesis we generalize this framework into two directions in Chapters 4 and 5. We start by giving further introduction to non-truthful mechanism design in Chapter 2 and introducing the necessary notation and preliminaries in Chapter 3.

1.1.1 Risk Aversion

A standard assumption in Algorithmic Game Theory is that players are risk neutral, meaning that they do not distinguish between different strategies that give them the same expected utility. This is in turn modeled by defining utility to be the difference between the valuation and payment for any given outcome. So, an agent having a value of 1 for an item would be indifferent between getting this item with probability 0.1 for free and getting it all the time, paying 0.9. However, there are many reasons to believe that agents are not risk neutral. For instance, in the above example the agent might favor the certain outcome to the uncertain one. Therefore, in Chapter 4, we raise the following question: *What “simple” auction mechanisms preserve good performance guarantees in the presence of risk-averse agents?*

Risk averseness can be formalized in various ways. The two most common models are (1) defining utility as a concave function of the difference between the valuation and the payment, or (2) taking into account the standard deviation when computing the utility. These two are also the ones inspected in Chapter 4. We give bounds on the price of anarchy for Bayes-Nash and (coarse) correlated equilibria of mechanisms in the presence of risk-averse agents and expose how the two models lead to different results.

More specifically, our main positive result states that the loss of performance in model (1) compared to the risk neutral setting is bounded by a small constant if a slightly stronger smoothness condition is fulfilled. We also prove that this condition is necessary by showing that the second price auction has an unbounded price of anarchy in the presence of risk averse players. This is aligned with the intuition that players should be more unwilling to participate in an all-pay auction than in, say, a first price auction. The results in Chapter 4 give the first theoretical backup to this observation.

In model (2), we arrive at quite different results: first price and all-pay auctions do not significantly differ. Furthermore, (λ, μ) -smoothness of a mechanism does not bring any guarantees in the presence of risk-averse players. These results imply that the variance-aversion model is not necessarily the most natural model for risk aversion in the setting studied here.

1.1.2 Simultaneous Composition with Varying Availability

In Chapter 5, we study a variant of simultaneous composition of mechanisms. Our scenario is motivated by limited availability or admission: Suppose bidders try to acquire items in a repeated online market, in which m items are sold simultaneously via, say, first-price auctions. However, in each round only some of the items are actually available for purchase. More specifically, in each round each item is available for each bidder only with a certain probability. Our scenario is an elementary case of simple mechanism design with incomplete information, where availabilities are bidder types. It captures natural applications in online markets with limited supply and can be used to model access of unreliable channels in wireless networks. The main question that we pose in Chapter 5

is: *What bidding strategies give good performance guarantees in a market composed of multiple mechanisms and where bidders are facing limited availability?*”

To avoid the drawbacks of existing results in terms of plausibility and computational complexity, we assume that the players learn with no-regret strategies in a way that is oblivious to their own and all other bidders’ availabilities. Thereby, bidders arrive at what we term an *availability-oblivious* coarse-correlated equilibrium — a bid distribution not tailored to the specific availabilities of bidders, which can be computed (approximately) in polynomial time. Our main result is that for a large class of valuation functions, we can apply smoothness ideas in this framework and prove bounds that mirror the known guarantees for compositions of smooth mechanisms.

In more detail, we prove general composition theorems for smooth mechanisms when valuation functions of bidders are lattice-submodular. They rely on an interesting connection to the notion of correlation gap of submodular functions over product lattices. Our results hold for independent and fully correlated bidder availabilities. In addition, we give an almost logarithmic price of anarchy lower bound for general fractionally subadditive (XOS) valuation functions.

1.2 Part II: Combinatorial Secretary Problems with Ordinal Information

The secretary problem is a classic model for online decision making. Recently, combinatorial extensions such as matroid or matching secretary problems have become an important tool to study algorithmic problems in dynamic markets. Here the decision maker must know the numerical value of each arriving element, which can be a demanding informational assumption. In Part II of this thesis, we initiate the study of algorithms for combinatorial secretary problems that rely only on ordinal information. We assume that there is an unknown value for each element, but our algorithms only have access to the *total order* of the elements arrived so far, which is consistent with their values. We term this the *ordinal model*; as opposed to the *cardinal model*, in which the algorithm learns the exact values. We show bounds on the competitive ratio, i.e., we compare the quality of the computed solutions to the optima in terms of the exact underlying but unknown numerical values. Consequently, competitive ratios for our algorithms are robust guarantees against uncertainty in the input. The guiding question of Part II of this thesis is: *How can we design online algorithms with small competitive ratios for combinatorial secretary problems in the ordinal model?*

In Chapter 6, we give an introduction and motivation to the problems we consider, together with preliminaries and related work. For the matroid secretary problem, we observe that many existing algorithms for special matroid structures maintain their competitive ratios even in the ordinal model. In these cases, the restriction to ordinal information does not represent any additional obstacle.

In Chapter 7 we show that ordinal variants of the submodular matroid secretary problems can be solved using algorithms for the linear versions by extending the results from [39]. In contrast, we also provide a lower bound of $\Omega(\sqrt{n}/(\log n))$ for algorithms that are oblivious to the matroid structure, where n is the total number of elements. This contrasts an upper bound of $O(\log n)$ in the cardinal model, and it shows that the

technique of thresholding is not sufficient for good algorithms in the ordinal model.

In Chapter 8, we design new algorithms that obtain constant competitive ratios for a variety of combinatorial problems, such as bipartite matching, general packing LPs and independent set with bounded local independence.

PART I

Extensions of the Smoothness Framework for Mechanisms

This part is the result of close collaboration with Martin Hoefer and Thomas Kesselheim. It is based on an article that appeared in *Proceedings of the Web and Internet Economics - 12th International Conference (WINE) 2016*, pages 294–308, in December 2016 [50], and an article that appeared in *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP) 2018*, pages 155:1–155:14, in July 2018 [57]. Full versions are available at <https://arxiv.org/abs/1509.00337> and <https://arxiv.org/abs/1804.09468>, respectively.

CHAPTER 2

Introduction to Part I

A common way to understand the effects of strategic behavior is to analyze the induced game among the bidders and bound the quality of (possibly manipulated) outcomes in equilibrium. That is, one compares the social welfare that is achieved at the (worst) equilibrium of the induced game to the maximum possible welfare. Typical equilibrium concepts are Bayes-Nash equilibria and (coarse) correlated equilibria, which extend mixed Nash equilibria toward incomplete information or learning settings respectively.

In a seminal paper, Syrgkanis and Tardos [85] propose a general technique for bounding social welfare of equilibria of “simple” mechanisms, based on a so-called “smoothness” technique. These guarantees apply even to mixed Bayes-Nash equilibria in environments with composition of mechanisms. For example, in a combinatorial auction we might not sell all items via a complicated truthful mechanism, but instead sell each item simultaneously via simple individual single-item auctions. Such a mechanism is obviously not truthful, since bidders are not even able to express their valuations for all subsets of items. However, if bidders have complement-free fractionally subadditive (XOS) valuations, the (expected) social welfare of allocations in a mixed Bayes-Nash equilibrium turns out to be a constant-factor approximation of the optimal social welfare.

While this is a fundamental insight into non-truthful mechanisms, it is not well-understood how this result extends under more realistic conditions. In particular, there has been recent concern about the plausibility and computational complexity of exact and approximate Bayes-Nash equilibria [18]. For more general Bayesian concepts based on no-regret learning strategies in repeated games, there are two natural approaches – either bidder types are drawn newly with bids, or types are drawn only once initially. While the latter is not really in line with the idea of incomplete information (bidders could communicate their type in the course of learning [18]), the former is in general hard to obtain. Also, the composition theorem applies only if bidders’ types are drawn independently.

In Chapter 5 of this thesis, we study a variant of simultaneous composition of mechanisms and show how to avoid the drawbacks of the Bayesian approach. Our scenario is motivated by limited availability or admission: Suppose bidders try to acquire items in a repeated online market, in which m items are sold simultaneously via, say, first-price auctions. However, in each round only some of the items are actually available for purchase. This scenario can be phrased in the Bayesian framework when bidder i ’s type is given by the set of items available to him. To obtain an equilibrium in the Bayesian sense, each bidder would have to consider a complicated bid vector and satisfy an equilibrium condition for each of the possible 2^m subsets of items.

In contrast, we assume that bidders do not even get to know (or are not able to account for) *their own availabilities* before making bids in each round. We assume they learn with no-regret strategies in a way that is *oblivious* to their own and all other bidders’

availabilities. Thereby, bidders arrive at what one might term an *availability-oblivious* coarse-correlated equilibrium – a bid distribution not tailored to the specific availabilities of bidders, which can be computed (approximately) in polynomial time. Our main result in Chapter 5 is that for a large class of valuation functions, we can apply smoothness ideas in this framework and prove bounds that mirror the guarantees from [85]. The guarantees apply even if some bidders learn obliviously and others follow a Bayes-Nash bidding strategy. In particular, we cover a broad domain with simultaneous composition of weakly smooth mechanisms in the sense of [85] when bidders have lattice-submodular valuations. Our study covers cases where availabilities are correlated among bidders and provides lower bounds for combinatorial auctions with item-bidding and XOS valuations. As a part of our analysis, we use the concept of correlation gap from [2] for submodular functions over product lattices.

Another key assumption in analyses of “simple” mechanisms is that agents are *risk neutral*: Agents are assumed to maximize their expected quasilinear utility, which is defined to be the difference of the value associated to the outcome and payment imposed to the agent. As already discussed, this would imply that an agent having a value of 1 for an item would be indifferent between getting this item with probability 10% for free and getting it all the time, paying 0.9.

However, there are many reasons to believe that agents are not risk neutral. For instance, in the above example the agent might favor the certain outcome to the uncertain one. Therefore, in Chapter 4 of this thesis, we characterize “simple” auction mechanisms that preserve good performance guarantees in the presence of risk-averse agents.

The standard model of risk aversion in economics (see, e.g., [69]) is to apply a (weakly) concave function to the quasilinear term. That is, if agent i 's outcome is x_i and his payment is p_i , his utility is given as $u_i(x_i, p_i) = h_i(v_i(x_i) - p_i)$, where $h_i: \mathbb{R} \rightarrow \mathbb{R}$ is a weakly concave, monotone function. Agent i is risk neutral, if and only if h_i is a linear function. If the function is strictly concave, this has the effect that, by Jensen's inequality, the utility for fixed x_i and p_i is higher than for a randomized x_i and p_i with the same expected $v_i(x_i) - p_i$.

We compare outcomes based on their social welfare, which is defined to be the sum of utilities of all involved parties including the auctioneer. That is, it is the sum of agents' utilities and their payments $\text{SW}(\mathbf{x}, \mathbf{p}) = \sum_i u_i(x_i, p_i) + \sum_i p_i$. In the quasilinear setting this definition of social welfare coincides with the sum of values $\sum_i v_i(x_i)$. With risk-averse utilities these two quantities usually differ. However, all our results bound the sum of values and therefore also hold for this benchmark.

We assume that the mechanisms are oblivious to the h_i -functions and work like in the quasilinear model. Only the individual agent's perception changes. This makes it necessary to normalize the h_i -functions as they could be on different scales.¹ Therefore, we will assume that $u_i(\mathbf{x}, p_i) = v_i(\mathbf{x})$ if $p_i = 0$ and that $u_i(\mathbf{x}, p_i) = 0$ if $p_i = v_i$. That is, for the two cases that p_i is either 0 or the full value, the utility matches exactly the quasilinear one. However, due to risk aversion, the agents might be less sensitive to payments. We note here that this will not in turn allow the mechanism to arrive at huge utility gains, as compared to the quasilinear model, by just cleverly splitting the quasilinear utility among the players. Indeed, Lemma 4.1 in Section 4.1 will show that

¹E.g., the functions could be such that $h_1(y) = y$ and $h_2(y) = 1000 \cdot y$, which would be impossible for the mechanism to cope with without additional information.

the difference between the two optima is bounded by at most a factor of 2.

For sake of simplicity, after providing the notation and preliminaries in Chapter 3, we start first with risk aversion in Chapter 4. Then, we present the technically more demanding material dealing with mechanism availability in Chapter 5.

2.1 Our Contribution

2.1.1 Risk Aversion

We give bounds on the price of anarchy for Bayes-Nash and (coarse) correlated equilibria of mechanisms in the presence of risk-averse agents. Our positive results are stated within the smoothness framework, which was introduced by [81]. We use the version that is tailored to quasilinear utilities by [85], which we extend to mechanism settings with general utilities (for a formal definition see Section 4.3). Our main positive result in Chapter 4 states that the loss of performance compared to the quasilinear setting is bounded by a constant if a slightly stronger smoothness condition is fulfilled.

Main Result 1. *Given a mechanism with price of anarchy α in the quasilinear model provable via smoothness such that the deviation guarantees non-negative utility, then this mechanism has price of anarchy at most 2α in the risk-averse model.*

This result relies on the fact that the deviation action to establish smoothness guarantees agents to have non-negative utility. A sufficient condition is that all undominated strategies never have negative utility. First-price and second-price auctions satisfy this condition, we thus get constant price-of-anarchy bounds for both of these auction formats.

In an all-pay auction every positive bid can lead to negative utility. Therefore, the positive result does not apply. As a matter of fact, this is not a coincidence because, as we show, equilibria can be arbitrarily bad.

Main Result 2. *The single-item all-pay auction has unbounded price of anarchy for Bayes-Nash equilibria, even with only three agents.*

This means that although equilibria of first-price and all-pay auctions have very similar properties with quasilinear utilities, in the risk-averse setting they differ by a lot. We feel that this to some extent matches the intuition that agents should be more reluctant to participate in an all-pay auction compared to a first-price auction.

In our construction, we give a symmetric Bayes-Nash equilibrium for two agents. The equilibrium is designed in such a way that a third agent of much higher value would lose with some probability with every possible bid. Losing in an all-pay auction means that the agent has to pay without getting anything, resulting in negative utility. In the quasilinear setting, this negative contribution to the utility would be compensated by respective positive amounts when winning. For the risk-averse agent in our example, this is not true. Because of the risk of negative utility, he prefers to opt out of the auction entirely.

We also consider a different model of aversion to uncertainty, in which solution concepts are modified. Instead of evaluating a distribution over utilities in terms of their expectation, agents evaluate them based on the expectation minus a second-order term. We find that this model has entirely different consequences on the price of anarchy. For

example, the all-pay auction has a constant price of anarchy in correlated and Bayes-Nash equilibria, whereas the second-price auction can have an unbounded price of anarchy in correlated equilibria.

2.1.2 Simultaneous Composition with Varying Availability

In the simultaneous composition setting, we assume that every mechanism satisfies a weak smoothness bound (for more details see Section 4.1) with parameters $\lambda, \mu_1, \mu_2 \geq 0$. It is known that for each individual mechanism, this implies an upper bound of $(\max(1, \mu_1) + \mu_2)/\lambda$ on the price of anarchy for no-regret learning outcomes and Bayes-Nash equilibria. Furthermore, the same bound also applies for outcomes of multiple simultaneous mechanisms that are tailored to availabilities, i.e., not oblivious.

In Section 5.2 we consider smoothness for oblivious learning and composition with independent availabilities, where in each round t , each mechanism j is available to each bidder i independently with probability $q_{i,j}$. Our smoothness bound involves the above parameters and the correlation gap of the class of valuation functions. In particular, if valuations v_i come from a class \mathcal{V} with a correlation gap of $\gamma(\mathcal{V})$, the price of anarchy becomes $\gamma(\mathcal{V}) \cdot (\max(1, \mu_1) + \mu_2)/\lambda$.

Main Result 3. *The price of anarchy for oblivious learning with monotone valuations that come from a class \mathcal{V} with a correlation gap of $\gamma(\mathcal{V})$ and fully independent admission is at most $\gamma(\mathcal{V}) \cdot (\mu_2 + \max(1, \mu_1))/\lambda$.*

Our construction uses smoothness of simultaneous composition from [85]. However, since learning is oblivious, the deviations establishing smoothness must be independent of availability. Here we use correlation gap to relate the value for independent deviations to that of type-dependent Bayesian deviations. Correlation gap is a notion originally defined for submodular set functions in [3]. It captures the worst-case ratio between the expected value of independent and correlated distributions over elements with the same marginals. We use an extension of this notion from [2] to Cartesian products of outcome spaces such as product lattices. For the class \mathcal{V} of monotone lattice-submodular valuations, we prove a correlation gap of $\gamma(\mathcal{V}) = e/(e - 1)$, which simplifies and slightly extends previous results.

In Section 5.3, we analyze oblivious learning for composition with correlated availabilities in the form of “everybody-or-nobody” – each mechanism is either available to all bidders or to no bidder. The probability for availability of mechanism j is q_j , and availabilities are independent among mechanisms. In this case, we simulate independence by assuming that each bidder draws random types and outcomes for himself. We also consider distributions where outcomes are drawn independently according to the marginals from the optimal correlated distribution over outcomes. While these two distributions are directly related via correlation gap, the technical challenge is to show that there is a connection to the value obtained by the bidder. For lattice-submodular functions, we show a smoothness bound that implies a price of anarchy of $4e/(e - 1) \cdot (\max(1, \mu_1) + \mu_2)/\lambda^2$.

Main Result 4. *The price of anarchy for oblivious learning with monotone lattice-submodular valuations and everybody-or-nobody admission is at most $4e/(e - 1) \cdot (\mu_2 + \max(1, \mu_1))/\lambda^2$.*

For neither of the results is it necessary that all bidders follow our oblivious-learning approach. We only require that bidders have no regret compared to this strategy. This is also fulfilled if some or all bidders determine their bids based on the actually available items rather than in the oblivious way.

Finally, in Section 5.4 we show a lower bound for simultaneous composition of single-item first-price auctions with general XOS valuation functions. The correlation gap for such functions is known to be large [3], but this does not directly imply a lower bound on the price of anarchy for oblivious learning. We provide a class of instances where the price of anarchy for oblivious learning becomes $\Omega((\log m)/(\log m \log m))$. This shows that for XOS functions it is impossible to generalize the constant price of anarchy for single-item first-price auctions.

Main Result 5. *The price of anarchy for pure Nash equilibria with oblivious bidding can be as large as $\Omega((\log m)/(\log \log m))$ if we allow XOS valuation functions.*

Our results have additional implications beyond auctions for the analysis of regret learning in wireless networks. We discuss these in Section 5.5.

2.2 Related Work

The smoothness framework was introduced by [81, 79] to analyze correlated and Bayes-Nash equilibria of general games. In [85] it was adjusted to the quasilinear case of mechanisms, and it was shown that simultaneous or sequential composition of smooth mechanisms is again smooth. Combinatorial auctions with item bidding are an example of a simultaneous composition. To show smoothness of the combined mechanism, it is thus enough to show smoothness of each single auction. Other examples of smooth mechanisms are position auctions with generalized second price [19, 77] and greedy auctions [65]. The smoothness approach for mixed Bayes-Nash equilibria shown in [85] is, in fact, slightly more general and continues to hold for variants of Bayesian correlated equilibrium [48].

Closely related to our work on mechanism availability are combinatorial auctions with item bidding, where multiple items are being sold in separate auctions. Bidders are generally interested in multiple items. However, depending on the bidder, some items may be substitutes for others. As the auctions work independently, bidders have to strategize in order to buy not too many items simultaneously. In a number of papers [23, 16, 49, 36] the efficiency of Nash and Bayes-Nash equilibria has been studied. It has been shown that, if the single items are sold in first or second price auctions and if the valuation functions are XOS or subadditive, the price of anarchy is constant. Limitations of this approach are shown in [24, 80].

The complexity of finding Bayes-Nash equilibria and Bayesian correlated equilibria has been studied only very recently. It has been shown in [18, 31] that equilibria are hard to find in some settings. In contrast, in [28] a different auction format is studied that yields good bounds on social welfare for equilibria that can be found more easily. Although similar in spirit, our approach is different – it shows that in some scenarios agents can reduce the computational effort and still obtain reasonably good states with existing mechanisms.

As such, our approach is closer to recent work [27] that shows hardness results for learning full-information coarse-correlated equilibria in simultaneous single-item second-price auctions with unit-demand bidders. As a remedy, a form of so-called no-envy learning is proposed, in which bidders use a different form of bidding that enables convergence in polynomial time. While achieving a general no-regret guarantee against all possible bid vectors is hard, we note here that our approach based on smoothness requires only a guarantee with respect to bids that are derived directly from the XOS representation of the bidder valuation. As such, bidders can obtain the guarantees required for our results in polynomial time. Conceptually, we here treat a different problem – the impact of availabilities, and more generally, different bidder types on learning outcomes in repeated mechanism design.

A model with dynamic populations in games has recently been considered in [66]. Each round a small portion of players are replaced by others with different utility functions. When players use algorithms that minimize a notion called adaptive regret, smoothness conditions and the resulting bounds on the price of anarchy continue to hold if there are solutions which remain near-optimal over time with a small number of structural changes. Using tools from differential privacy, these conditions are shown for some special classes of games, including first-price auctions with unit-demand or gross-substitutes valuations. In contrast, our scenario is orthogonal, since we consider much more general classes of mechanisms and allow changes in each round for possibly all players. However, our model of change captures the notion of availability and therefore is much more specific than the adversarial approach of [66].

The notion of correlation gap was defined and analyzed for stochastic optimization in [3, 2]. It was used in [87] for analyzing revenue maximization with sequential auctions, which is very different from our approach.

Studying the impact of risk-averseness is a regularly reoccurring theme in the literature. A proposal to distinguish between money and the utility of money, and to model risk aversion by a utility function that is concave first appeared in [14]. The expected utility theory, which basically states that the agent’s behavior can be captured by a utility function and the agent behaves as a maximizer of the expectation of his utility, was postulated in [86]. This theory does not capture models that are standardly used in portfolio theory, “expectation minus variance” or “expectation minus standard deviation” [67], the latter of which we also consider in Section 4.6.

In the context of mechanisms, one usually models risk aversion by concave utility functions. One research direction in this area is to understand the effects of risk aversion on a given mechanism. For example, in [42] the authors study symmetric equilibria in all-pay auctions with a homogeneous population of risk-averse players. They compare the bidding behavior to the risk-neutral case. In [71] a similar analysis for auctions with a buyout option is performed; in [54] customers with heterogeneous risk attitudes in mechanisms for cloud resources are considered. In [34] it is shown that for certain classes of mechanisms the correlated equilibrium is unique and has a certain structure. One consequence of this result is that risk aversion does not influence the form of the equilibria or the revenue.

Another direction is to design mechanisms for the risk-averse setting. For example, the optimal revenue is higher because buyers are less sensitive to payments. In a number of papers, mechanisms for revenue maximization are proposed [72, 70, 84, 55, 15, 44].

Furthermore, randomized mechanisms that are *truthful in expectation* lose their incentive properties if agents are not risk neutral. Black-box transformations from truthful-in-expectation mechanisms into ones that fulfill stronger properties are given in [33] and [51].

Studying the effects of risk aversion also has a long history in game theory, where different models of agents' attitudes towards risk are analyzed. One major question is, for example, if equilibria still exist and if they can be computed [41, 53]. Price of anarchy analyses have so far only been carried out for congestion games. Tight bounds on the price of anarchy for atomic congestion games with affine cost functions under a range of risk-averse decision models are given in [78].

It is important to remark here that our approach to risk is different from the one taken by [73]. They use the smoothness framework to prove generalized price of anarchy bounds for games in which players have biased utility functions. They assume that players are playing the “wrong game” and their point of comparison is the “true” optimal social welfare, meaning that the biases only determine the equilibria but do not affect the social welfare. We take the utility functions as they are, including the risk aversion, to evaluate social welfare in equilibria and also to determine the optimum, which makes our models incomparable.

For precise relation of von Neumann-Morgenstern preferences to mean-variance preferences, see for instance [68]. Mean-variance preferences were explored for congestion games in [75, 76], while the authors in [62] study the bidding behavior in an all-pay auction depending on the level of variance-averseness.

CHAPTER 3

Notation and Preliminaries

3.1 Setting

We consider the following setting: There is a set N of n players and \mathcal{X} is the set of possible outcomes. Each player i has a utility function $u_i^{\theta_i}$, which is parameterized by his type $\theta_i \in \Theta_i$. Given a type θ_i , an outcome $\mathbf{x} \in \mathcal{X}$, and a payment $p_i \geq 0$, his utility is $u_i^{\theta_i}(\mathbf{x}, p_i)$. The traditionally most studied case are quasilinear utilities, in which types are valuation functions $v_i \in \mathcal{V}_i$, $v_i: \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ and $u_i^{\theta_i}(\mathbf{x}, p_i) = v_i(\mathbf{x}) - p_i$.

For fixed utility functions and types, the social welfare of an outcome $\mathbf{x} \in \mathcal{X}$ and payments $(p_i)_{i \in N}$ is defined as

$$\text{SW}^\theta(\mathbf{x}, \mathbf{p}) := \sum_{i \in N} u_i^{\theta_i}(\mathbf{x}, p_i) + \sum_{i \in N} p_i . \quad (3.1)$$

In the quasilinear case, this simplifies to

$$\text{SW}^\theta(\mathbf{x}, \mathbf{p}) = \sum_{i \in N} v_i(\mathbf{x}) . \quad (3.2)$$

Unless noted otherwise, by $\text{OPT}(\theta)$, we will refer to the optimal social welfare under type profile θ , i.e.,

$$\text{OPT}(\theta) = \max_{\mathbf{x}, \mathbf{p}} \text{SW}^\theta(\mathbf{x}, \mathbf{p}) . \quad (3.3)$$

A *mechanism* \mathcal{M} is a triple (\mathcal{A}, f, p) , where $\mathcal{A} = \times_i \mathcal{A}_i$ is the set of actions and \mathcal{A}_i is the set of actions for each player i , $f: \mathcal{A} \rightarrow \mathcal{X}$ is an allocation function that maps actions to outcomes and $p: \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}^n$ is a payment function that maps actions to payments p_i for each player i . Given an action profile $\mathbf{a} \in \mathcal{A}$, we will use the short-hand notation $u_i^{\theta_i}(\mathbf{a})$, or sometimes even $u_i(\mathbf{a})$, to denote $u_i^{\theta_i}(f(\mathbf{a}), p_i)$.

We assume that players always have the possibility of not participating, hence any rational outcome has non-negative utility in expectation over the non-available information and the randomness of other players and the mechanism.

3.1.1 Simultaneous Composition

In Chapter 5 we will focus on the following generalized setting: There are, as before, n players but they are participating in m mechanisms, where $m \geq 1$. The mechanisms are not running in isolation, but rather take place simultaneously. Each mechanism \mathcal{M}_j has its own outcome space \mathcal{X}_j and consists of a triple $(\mathcal{A}_j, f_j, p_j)$ as described previously, i.e., $\mathcal{A}_j = \times_i \mathcal{A}_{i,j}$ is the action space, $f_j: \mathcal{A}_j \rightarrow \mathcal{X}_j$ is the allocation function and $p_j: \mathcal{A}_j \rightarrow \mathbb{R}_{\geq 0}^n$ the payment function.

In this case, we assume that a player has a valuation over vectors of outcomes from the different mechanisms: $v_i : \times_j \mathcal{X}_j \rightarrow \mathbb{R}_{\geq 0}$. A player's utility will be quasi-linear in this extended setting in the sense that his utility from an allocation vector $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_m) \in \times_j \mathcal{X}_j$ and payment vector $\mathbf{p}_i = (p_{i,1}, \dots, p_{i,m})$ is given by:

$$u_i^{v_i}(\mathbf{x}, \mathbf{p}_i) = v_i(\mathbf{x}_1, \dots, \mathbf{x}_m) - \sum_{j=1}^m p_{i,j} . \quad (3.4)$$

Players can have valuations that are complex functions of the outcomes of different mechanisms. The main results for simultaneous composition of mechanisms in [85], which we will use in Chapter 5, hold for the class of valuation functions known as XOS.

3.1.2 Valuation Function Classes

Definition 3.1 (XOS). *Valuation v_i of player i is XOS if there exists a set \mathcal{L} of additive valuations $v_{i,j}^\ell : \mathcal{X}_j \rightarrow \mathbb{R}_{\geq 0}, \forall \ell \in L, i \in N, j \in [m]$, such that: $v_i(\mathbf{x}) = \max_{\ell \in \mathcal{L}} \sum_j v_{i,j}^\ell(\mathbf{x}_j)$.*

Our main results in Chapter 5 will hold for the class of lattice-submodular valuations. We first give the definition of a submodular set function.

Definition 3.2 (Submodular set function). *Let Ω be a finite set. A function $f : 2^\Omega \rightarrow \mathbb{R}$ is set submodular, if $\forall S, T \subseteq \Omega$ such that $S \subseteq T$ and $\forall x \in \Omega \setminus T$*

$$f(S \cup \{x\}) - f(S) \geq f(T \cup \{x\}) - f(T) .$$

Equivalently, $\forall S, T \subseteq \Omega$

$$f(S \cup T) + f(S \cap T) \leq f(S) + f(T) .$$

Before delivering the definition of a lattice-submodular function, we state the definition of a lattice.

Definition 3.3 (Lattice). *A partially ordered set (L, \succeq) is called a lattice if each two-element subset $\{a, b\} \subseteq L$ has a join (i.e., least upper bound) and a meet (i.e., greatest lower bound), denoted by $a \vee b$ and $a \wedge b$, respectively.*

Definition 3.4 (Lattice-submodular valuation). *Suppose for every mechanism j the set $\mathcal{X}_{i,j}$ of possible outcomes for bidder i forms a lattice $(\mathcal{X}_{i,j}, \succeq_{i,j})$ with a partial order $\succeq_{i,j}$. Bidder i has a lattice-submodular valuation v_i if and only if it is submodular on the product lattice $(\mathcal{X}_i, \succeq_i)$ of outcomes for bidder i :*

$$\forall \mathbf{x}_i, \tilde{\mathbf{x}}_i \in \mathcal{X}_i : v_i(\mathbf{x}_i \vee \tilde{\mathbf{x}}_i) + v_i(\mathbf{x}_i \wedge \tilde{\mathbf{x}}_i) \leq v_i(\mathbf{x}_i) + v_i(\tilde{\mathbf{x}}_i) .$$

Lattice-submodular functions generalize submodular set functions but are a strict subclass of XOS functions.

A further prominent valuation class is the class of unit-demand valuations.

Definition 3.5 (Unit-demand). *Valuation v_i of player i is unit-demand, if there exist valuations $v_{i,j}$ such that*

$$v_i(S) = \max_{j \in S} v_{i,j} .$$

3.2 Solution Concepts and Benchmarks

In the setting of *complete information*, the type profile θ is fixed. We consider (coarse) correlated equilibria, which generalize Nash equilibria and are the outcome of (no-regret) learning dynamics.

Definition 3.6. A correlated equilibrium (CE) is a distribution \mathbf{a} over action profiles from \mathcal{A} such that for every player i and every strategy a_i in the support of \mathbf{a} and every action $a'_i \in \mathcal{A}_i$, player i does not benefit from switching to a'_i whenever he was playing a_i . Formally,

$$\mathbf{E}_{\mathbf{a}_{-i}|a_i}[u_i(\mathbf{a})] \geq \mathbf{E}_{\mathbf{a}_{-i}|a_i}[u_i(a'_i, \mathbf{a}_{-i})], \forall a'_i \in \mathcal{A}_i, \forall i .$$

Definition 3.7. A coarse correlated equilibrium (CCE) is a distribution \mathbf{a} over action profiles from \mathcal{A} such that for every player i and every action $a'_i \in \mathcal{A}_i$, player i does not benefit from switching to a'_i . Formally,

$$\mathbf{E}_{\mathbf{a}}[u_i(\mathbf{a})] \geq \mathbf{E}_{\mathbf{a}}[u_i(a'_i, \mathbf{a}_{-i})], \forall a'_i \in \mathcal{A}_i, \forall i .$$

In *incomplete information*, the type of each player is drawn from a distribution F_i over his type space Θ_i . The distributions are common knowledge and the draws are independent among players. The solution concept we consider in this setting is the Bayes-Nash equilibrium. Here, the strategy of each player is now a (possibly randomized) function $s_i: \Theta_i \rightarrow \mathcal{A}_i$.

Definition 3.8. A Bayes-Nash equilibrium (BNE) is a distribution \mathbf{s} over functions s_i such that for every player i , every type θ_i and every strategy $a_i \in \mathcal{A}_i$, player i does not benefit from switching to a_i whenever he was playing $s_i(\theta_i)$. Formally,

$$\mathbf{E}_{\theta_{-i}|\theta_i}[u_i^{\theta_i}(\mathbf{s}(\theta))] \geq \mathbf{E}_{\theta_{-i}|\theta_i}[u_i^{\theta_i}(a_i, \mathbf{s}_{-i}(\theta_{-i}))], \forall a_i \in \mathcal{A}_i, \forall \theta_i \in \Theta_i, \forall i .$$

The measure of efficiency is the expected social welfare over the types of the players: Given a strategy profile $\mathbf{s}: \times_i \Theta_i \rightarrow \times_i \mathcal{A}_i$, we consider $\mathbf{E}_{\theta}[\text{SW}^{\theta}(\mathbf{s}(\theta))]$. We compare the efficiency of our solution concept with respect to the expected optimal social welfare $\mathbf{E}_{\theta}[\text{OPT}(\theta)]$.

Definition 3.9. The Price of Anarchy (PoA) with respect to an equilibrium concept is the worst possible ratio between the optimal expected welfare and the expected welfare at equilibrium. Formally,

$$\text{PoA} = \max_F \max_{D \in \text{EQ}(F)} \frac{\mathbf{E}_{\theta \sim F}[\text{OPT}(\theta)]}{\mathbf{E}_{\theta \sim F, \mathbf{a} \sim D}[\text{SW}^{\theta}(\mathbf{a})]} ,$$

where by $F = F_1 \times \dots \times F_n$ we denote the product distribution of the players' type distributions and by $\text{EQ}(F)$ the set of all equilibria, which are probability distributions over action profiles.

Note that PoA generally depends on the set of considered equilibria and can therefore differ for different equilibrium concepts.

3.3 Smoothness Framework

The results in Part I of this thesis will mostly be based on the smoothness framework for mechanisms, as given in [85]. Here we only introduce the basic definitions and main theorems. All definitions and theorems in this section assume quasilinear utilities. For the sake of simplicity, we defer the definition and discussion of further concepts (as, for instance, “weak smoothness”) to Chapter 4 and Chapter 5.

Definition 3.10 (Smooth Mechanism). *A mechanism is (λ, μ) -smooth if for any valuation profile $\mathbf{v} \in \times_i \mathcal{V}_i$ and for any action profile \mathbf{a} there exists a randomized action $a_i^*(\mathbf{v}, a_i)$ for each player i , s.t.:*

$$\sum_i u_i^{v_i}(a_i^*(\mathbf{v}, a_i), \mathbf{a}_{-i}) \geq \lambda \text{OPT}(\mathbf{v}) - \mu \sum_i P_i(\mathbf{a}) \quad (3.5)$$

for some $\lambda, \mu \geq 0$. If \mathbf{a} is a vector of randomized strategies, $u_i^{v_i}(\mathbf{a})$ denotes the expected utility of a player .

The following theorems from [85] reveal the appeal of checking the very technical condition that a mechanism needs to satisfy in order to be (λ, μ) -smooth.

Theorem 3.11. *If a mechanism is (λ, μ) -smooth, then any correlated equilibrium in the full information setting and any Bayes-Nash equilibrium in the Bayesian setting achieves efficiency of at least a fraction of $\frac{\lambda}{\max\{1, \mu\}}$ of $\text{OPT}(\mathbf{v})$ or of $\mathbf{E}_{\mathbf{v}}[\text{OPT}(\mathbf{v})]$, respectively.*

Lastly, the next theorem bounds the price of anarchy of a simultaneous composition of (λ, μ) -smooth mechanisms.

Theorem 3.12 (Simultaneous Composition). *Consider the mechanism defined by the simultaneous composition of m mechanisms. Suppose that each mechanism \mathcal{M}_j is (λ, μ) -smooth when the mechanism restricted valuations of the players come from a class of valuations $(\mathcal{V}_{i,j})_{i \in N}$. If the valuation $v_i : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ of each player across mechanisms is XOS, meaning it can be expressed by component valuations $v_{i,j}^\ell \in \mathcal{V}_{i,j}$, then the global mechanism is also (λ, μ) -smooth.*

CHAPTER 4

Risk-Averse Agents

4.1 Model

When modeling risk aversion, one wants to capture the fact that a random payoff (lottery) X is less preferred than a deterministic one of value $\mathbf{E}[X]$. The standard approach is, therefore, to apply a concave non-decreasing function $h: \mathbb{R} \rightarrow \mathbb{R}$ to X and consider $h(X)$ instead. By Jensen's inequality, we now know $\mathbf{E}[h(X)] \leq h(\mathbf{E}[X])$.

In the case of mechanism design, the utility of a risk-neutral agent is defined as the quasilinear utility $v_i(\mathbf{x}) - p_i$. That is, if an agent has a value of 1 for an item and has to pay 0.9 for it, then the resulting utility is 0.1. The expected utility is identical if the agent only gets the item with probability $\frac{1}{10}$ for free. To capture the effect that the agent prefers the certain outcome to the uncertain one, we again apply a concave function $h_i: \mathbb{R} \rightarrow \mathbb{R}$ to the quasilinear term $v_i(\mathbf{x}) - p_i$. We therefore consider utility functions of the form $u_i(\mathbf{x}, p_i) = h_i(v_i(\mathbf{x}) - p_i)$ in the general setting described in Chapter 3. Note that the mechanisms we consider do not know the h_i -functions. They work as if all utility functions were quasilinear. Throughout this chapter, we will refer to quasilinear utilities by \hat{u}_i .

We want to compare outcomes based on their social welfare. We use the definition of social welfare being the sum of utilities of all involved parties including the auctioneer. That is,

$$\text{SW}(\mathbf{x}, \mathbf{p}) = \sum_{i \in N} u_i(\mathbf{x}, p_i) + \sum_{i \in N} p_i . \quad (4.1)$$

It is impossible for any mechanism to choose good outcomes for this benchmark if the h_i -function are arbitrary and unknown. Therefore, we assume that utility functions are normalized so that the utility matches the quasilinear one for $p_i = 0$ and $p_i = v_i(\mathbf{x})$ (see Figure 4.1). In more detail, we assume the following *normalized risk-averse utilities*:

- (1) $u_i^{v_i}(\mathbf{x}, p_i) \geq u_i^{v_i}(\mathbf{x}, p'_i)$, if $p_i \leq p'_i$ (monotonicity)
- (2) $u_i^{v_i}(\mathbf{x}, p_i) = 0$, if $p_i = v_i(\mathbf{x})$ (normalization at $p_i = v_i(\mathbf{x})$)
- (3) $u_i^{v_i}(\mathbf{x}, p_i) = v_i(\mathbf{x})$, if $p_i = 0$ (normalization at $p_i = 0$)
- (4) $u_i^{v_i}(\mathbf{x}, p_i) \geq v_i(\mathbf{x}) - p_i$, if $0 \leq p_i \leq v_i(\mathbf{x})$;
 $u_i^{v_i}(\mathbf{x}, p_i) \leq v_i(\mathbf{x}) - p_i$, otherwise (relaxed concavity)

Assumption (4) is a relaxed version of concavity that suffices our needs for the positive results. Our negative results always fulfill concavity.

As an effect of normalization, the optimal social welfare of the risk-averse setting can be bounded in terms of the optimal sum of values, which coincides with the social welfare for quasilinear utilities.




Figure 4.1: Normalized risk-averse utility function (bold) and quasilinear utility function for a fixed allocation \mathbf{x} and varying payment p_i .

Lemma 4.1. *Given valuation functions $(v_i)_{i \in N}$ and normalized risk-averse utilities $(u_i^{v_i})_{i \in N}$, let OPT denote the optimal social welfare with respect to utilities $(u_i^{v_i})_{i \in N}$ and $\widehat{\text{OPT}}$ denote the optimal social welfare with respect to quasilinear utilities $(\hat{u}_i^{v_i})_{i \in N}$. Then, $\text{OPT} \leq 2\widehat{\text{OPT}}$.*

Proof. Let (\mathbf{x}, \mathbf{p}) denote the outcome and payment profile that maximizes the social welfare $\sum_{i \in N} u_i^{v_i}(\mathbf{x}, p_i) + \sum_{i \in N} p_i$ respectively. We observe we can safely assume that $0 \leq p_i \leq v_i(\mathbf{x})$. Otherwise, we know from property (4) that $u_i^{v_i}(\mathbf{x}, p_i) + p_i \leq v_i(\mathbf{x})$, so changing p_i to 0 could only increase social welfare.

By monotonicity of $u_i^{v_i}(\mathbf{x}, \cdot)$ and Assumption (3), $\text{OPT} = \sum_{i \in N} u_i^{v_i}(\mathbf{x}, p_i) + \sum_{i \in N} p_i \leq \sum_{i \in N} u_i^{v_i}(\mathbf{x}, 0) + \sum_{i \in N} p_i \leq \sum_{i \in N} v_i(\mathbf{x}) + \sum_{i \in N} v_i(\mathbf{x}) = 2\widehat{\text{OPT}}$. \square

As a consequence, the optimal social welfare changes only within a factor of 2 by risk aversion, and we may as well take $\widehat{\text{OPT}}$ as our point of comparison. A Vickrey-Clark-Groves mechanism (VCG) mechanism, for example, is still incentive compatible under risk-averse utilities but optimizes the wrong objective function. Lemma 4.1 shows that it is still a constant-factor approximation to the optimal social welfare. However, in simple mechanisms, the agents' behavior and the resulting equilibria may or may not change drastically under risk aversion, depending on the mechanism.

4.2 Single-Item Auctions in the Quasilinear Setting

In the standard single-item setting, one item is auctioned among n players, with their valuations and actions (bids) both being real numbers. In the common auction formats, the item is given to the bidder with the highest bid.

First Price Auction

In a first price auction, the highest bidder wins the item and pays his bid. In order to show that this is indeed a (λ, μ) -smooth mechanism, we need to find deviations that satisfy equation (5.1). If we let the highest value player, with a valuation of v_h , deviate to half of his value and everybody else to 0:

$$\sum_i \hat{u}_i(a_i^*, \mathbf{a}_{-i}) \geq \hat{u}_h\left(\frac{1}{2}v_h, \mathbf{a}_{-h}\right) \geq \frac{1}{2}v_h - \max_{i \neq h} a_i \geq \frac{1}{2}\widehat{\text{OPT}} - \max_i a_i \geq \frac{1}{2}\widehat{\text{OPT}} - \sum_i p_i(\mathbf{a}),$$

where for deriving the second inequality we verify that it holds in both the case of the highest player winning and also losing. We conclude that first price auction is a $(\frac{1}{2}, 1)$ -smooth mechanism and therefore the price of anarchy of its correlated and Bayes-Nash equilibria is at most 2.

All-Pay Auction

In an all-pay auctions, the highest bidder wins the item but all the players pay their bid. If we again let the highest value player deviate to half of his value and everyone else to 0:

$$\sum_i \hat{u}_i(a_i^*, \mathbf{a}_{-i}) \geq \hat{u}_h\left(\frac{1}{2}v_h, \mathbf{a}_{-h}\right) \geq \frac{1}{2}v_h - 2 \cdot \max_{i \neq h} a_i \geq \frac{1}{2}v_h - 2 \sum_i a_i = \frac{1}{2}\widehat{\text{OPT}} - 2 \sum_i p_i(\mathbf{a}) ,$$

therefore all-pay auction is a $(\frac{1}{2}, 2)$ -smooth mechanism and the price of anarchy of its correlated and Bayes-Nash equilibria is bounded by 4.

Second Price Auction

In a second price auction, the highest bidder wins the item and pays the second highest bid. Here, we need an no-overbidding assumption in order to guarantee good performance in equilibria. In other words, we want to show that second price auction is a (λ, μ_1, μ_2) -smooth mechanism. If we let the highest value player deviate to his value and everyone else to 0:

$$\sum_i \hat{u}_i(a_i^*, \mathbf{a}_{-i}) \geq \hat{u}_h(v_h, \mathbf{a}_{-h}) \geq v_h - \max_{i \neq h} a_i \geq v_h - \sum_i a_i = \widehat{\text{OPT}} - \sum_i W_i(a_i, f(\mathbf{a})) ,$$

so second price auction is a $(1, 0, 1)$ -weakly smooth mechanism, meaning that the price of anarchy of its correlated and Bayes-Nash equilibria in which the no-overbidding assumption is satisfied is upper bounded by 2.

These smoothness results were given by [85]. As we saw in Theorem 3.12 of Section 3.3, a simultaneous composition of (λ, μ_1, μ_2) -smooth mechanisms is again (λ, μ_1, μ_2) -smooth, under a condition that the players' valuations across mechanism come from a certain valuation function class, namely XOS. Going back to the presented examples, this in particular also means that we can state price of anarchy bounds for a simultaneous composition of an arbitrary number of first price, second price or all-pay auctions.

What is remarkable here is that first-price and all-pay auctions achieve nearly the same welfare guarantees. We will show that in the risk-averse setting this is not true. While the first-price auction almost preserves its constant price of anarchy, the all-pay auction has an unbounded price of anarchy, even with only three players.

4.3 Smoothness Beyond Quasilinear Utilities

Most of our positive results in this chapter rely on the *smoothness* framework for quasilinear mechanism design by [85]. Since our utility functions will not be quasilinear, in this section we first observe that the framework can be extended to general utility functions. Note that throughout this section, the exact definition of $\text{OPT}(\boldsymbol{\theta})$ is irrelevant. Therefore, it can be set to the optimal social welfare but also to weaker benchmarks depending on the setting.

Definition 4.2 (Smooth Mechanism). *A mechanism \mathcal{M} is (λ, μ) -smooth with respect to utility functions $(u_i^{\theta_i})_{\theta_i \in \Theta_i, i \in N}$ for $\lambda, \mu \geq 0$, if for any type profile $\boldsymbol{\theta} \in \times_i \Theta_i$ and for any*

action profile \mathbf{a} there exists a randomized action $a_i^*(\boldsymbol{\theta}, a_i)$ for each player i , such that $\sum_i u_i^{\theta_i}(a_i^*(\boldsymbol{\theta}, a_i), \mathbf{a}_{-i}) \geq \lambda \text{OPT}(\boldsymbol{\theta}) - \mu \sum_i p_i(\mathbf{a})$. We denote by $u_i^{\theta_i}(\mathbf{a})$ the expected utility of a player if \mathbf{a} is a vector of randomized strategies.

Mechanism smoothness implies bounds on the price of anarchy. The following theorem and its proof are analogous to the theorems in [85]. In cases where the deviation required by smoothness does not depend on a_i , the results extend to coarse correlated equilibria. The important point is that the respective bounds mostly do not depend on the assumption of quasilinearity.

Theorem 4.3. *If a mechanism \mathcal{M} is (λ, μ) -smooth with respect to utility functions $(u_i^{\theta_i})_{\theta_i \in \Theta_i, i \in N}$, then any correlated equilibrium in the full information setting and any Bayes-Nash equilibrium in the Bayesian setting achieves efficiency of at least a fraction of $\frac{\lambda}{\max\{1, \mu\}}$ of $\text{OPT}(\boldsymbol{\theta})$ or of $\mathbf{E}_{\boldsymbol{\theta}}[\text{OPT}(\boldsymbol{\theta})]$, respectively.*

Proof. **Full Information Setting**

Let \mathbf{a} be a correlated equilibrium. This means that for every a_i in the support of \mathbf{a}

$$\mathbf{E}_{\mathbf{a}_{-i}|a_i} \left[u_i^{\theta_i}(a_i, \mathbf{a}_{-i}) \right] \geq \mathbf{E}_{\mathbf{a}_{-i}|a_i} \left[u_i^{\theta_i}(a'_i, \mathbf{a}_{-i}) \right], \forall a'_i \in \mathcal{A}_i, \forall i .$$

Applying the equilibrium property to $a'_i = a_i^*(\boldsymbol{\theta}, a_i)$, we know that for every a_i in the support of \mathbf{a} :

$$\mathbf{E}_{\mathbf{a}_{-i}|a_i} \left[u_i^{\theta_i}(a_i, \mathbf{a}_{-i}) \right] \geq \mathbf{E}_{\mathbf{a}_{-i}|a_i} \left[u_i^{\theta_i}(a_i^*(\boldsymbol{\theta}, a_i), \mathbf{a}_{-i}) \right], \forall i .$$

If we now take the expectation over a_i and add over all players:

$$\mathbf{E}_{\mathbf{a}} \left[\sum_i u_i^{\theta_i}(\mathbf{a}) \right] \geq \mathbf{E}_{\mathbf{a}} \left[\sum_i u_i^{\theta_i}(a_i^*(\boldsymbol{\theta}, a_i), \mathbf{a}_{-i}) \right] \geq \lambda \text{OPT}(\boldsymbol{\theta}) - \mu \mathbf{E}_{\mathbf{a}} \left[\sum_i p_i(\mathbf{a}) \right] ,$$

and further by adding $\mathbf{E}_{\mathbf{a}}[\sum_i p_i(\mathbf{a})]$ to both sides

$$\begin{aligned} \mathbf{E}_{\mathbf{a}} \left[\sum_i u_i^{\theta_i}(\mathbf{a}) + \sum_i p_i(\mathbf{a}) \right] &\geq \mathbf{E}_{\mathbf{a}} \left[\sum_i u_i^{\theta_i}(a_i^*(\boldsymbol{\theta}, a_i), \mathbf{a}_{-i}) \right] \\ &\geq \lambda \text{OPT}(\boldsymbol{\theta}) + (1 - \mu) \mathbf{E}_{\mathbf{a}} \left[\sum_i p_i(\mathbf{a}) \right] . \end{aligned}$$

The result follows by doing a case distinction over $\mu \leq 1$ and $\mu > 1$. In the first case, we immediately get

$$\mathbf{E}_{\mathbf{a}} \left[\sum_i u_i^{\theta_i}(\mathbf{a}) + \sum_i p_i(\mathbf{a}) \right] \geq \lambda \text{OPT}(\boldsymbol{\theta}) + (1 - \mu) \mathbf{E}_{\mathbf{a}} \left[\sum_i p_i(\mathbf{a}) \right] \geq \lambda \text{OPT}(\boldsymbol{\theta}) ,$$

and in the latter case we use the fact that $\mathbf{E}_{\mathbf{a}} \left[\sum_i u_i^{\theta_i}(\mathbf{a}) \right] \geq 0$. Then also

$$\mathbf{E}_{\mathbf{a}} \left[\sum_i u_i^{\theta_i}(\mathbf{a}) \right] + \mathbf{E}_{\mathbf{a}} \left[\sum_i p_i(\mathbf{a}) \right] \geq \mathbf{E}_{\mathbf{a}} \left[\sum_i p_i(\mathbf{a}) \right] ,$$

which results in

$$\mathbf{E}_{\mathbf{a}} \left[\sum_i u_i^{\theta_i}(\mathbf{a}) + \sum_i p_i(\mathbf{a}) \right] \geq \lambda \text{OPT}(\boldsymbol{\theta}) + (1 - \mu) \mathbf{E}_{\mathbf{a}} \left[\sum_i u_i^{\theta_i}(\mathbf{a}) + \sum_i p_i(\mathbf{a}) \right]$$

and finally

$$\mathbf{E}_{\mathbf{a}} \left[\sum_i u_i^{\theta_i}(\mathbf{a}) + \sum_i p_i(\mathbf{a}) \right] \geq \frac{\lambda}{\mu} \cdot \text{OPT}(\boldsymbol{\theta}) .$$

Bayesian Setting

For reasons of clarity, we prove the claim for the simpler case of pure BNE. First, we let each player i sample a type profile $\boldsymbol{\zeta} \sim \times_i F_i$ and play $a_i^*((\theta_i, \boldsymbol{\zeta}_{-i}), s_i(\zeta_i))$.

$$\begin{aligned} \mathbf{E}_{\boldsymbol{\theta}} \left[u_i^{\theta_i}(s(\boldsymbol{\theta})) \right] &\geq \mathbf{E}_{\boldsymbol{\theta}, \boldsymbol{\zeta}} \left[u_i^{\theta_i}(a_i^*((\theta_i, \boldsymbol{\zeta}_{-i}), s_i(\zeta_i)), s_{-i}(\boldsymbol{\theta}_{-i})) \right] \\ &= \mathbf{E}_{\boldsymbol{\theta}, \boldsymbol{\zeta}} \left[u_i^{\zeta_i}(a_i^*((\zeta_i, \boldsymbol{\zeta}_{-i}), s_i(\theta_i)), s_{-i}(\boldsymbol{\theta}_{-i})) \right] \\ &= \mathbf{E}_{\boldsymbol{\theta}, \boldsymbol{\zeta}} \left[u_i^{\zeta_i}(a_i^*(\boldsymbol{\zeta}, s_i(\theta_i)), s_{-i}(\boldsymbol{\theta}_{-i})) \right] \end{aligned}$$

Summing over the players and using the smoothness property, we get

$$\begin{aligned} \mathbf{E}_{\boldsymbol{\theta}} \left[\sum_i u_i^{\theta_i}(s(\boldsymbol{\theta})) \right] &\geq \mathbf{E}_{\boldsymbol{\theta}, \boldsymbol{\zeta}} \left[\sum_i u_i^{\zeta_i}(a_i^*(\boldsymbol{\zeta}, s_i(\theta_i)), s_{-i}(\boldsymbol{\theta}_{-i})) \right] \\ &\geq \mathbf{E}_{\boldsymbol{\theta}, \boldsymbol{\zeta}} \left[\lambda \text{OPT}(\boldsymbol{\zeta}) - \mu \sum_i p_i(s(\boldsymbol{\theta})) \right] \\ &= \lambda \mathbf{E}_{\boldsymbol{\theta}} [\text{OPT}(\boldsymbol{\theta})] - \mu \mathbf{E}_{\boldsymbol{\theta}} \left[\sum_i p_i(s(\boldsymbol{\theta})) \right] , \end{aligned}$$

and therefore

$$\mathbf{E}_{\boldsymbol{\theta}} \left[\sum_i u_i^{\theta_i}(s(\boldsymbol{\theta})) + \sum_i p_i(s(\boldsymbol{\theta})) \right] \geq \lambda \mathbf{E}_{\boldsymbol{\theta}} [\text{OPT}(\boldsymbol{\theta})] + (1 - \mu) \mathbf{E}_{\boldsymbol{\theta}} \left[\sum_i p_i(s(\boldsymbol{\theta})) \right] ,$$

from where the result follows by case distinction over μ , as in the proof for the full information setting.

The generalization to a mixed Bayes-Nash equilibrium is now straightforward. \square

4.3.1 Weak Smoothness

For second-price auctions and their generalizations, for example, the already stated theorems do not suffice to prove guarantees on the quality of equilibria. For such mechanisms we additionally need an no-overbidding assumption. To state this assumption, we first discuss the notion of willingness-to-pay that was originally defined in [85].

Definition 4.4 (Willingness-to-pay). *Given a mechanism (\mathcal{A}, f, p) a player's maximum willingness-to-pay for an allocation \mathbf{x} when using strategy a_i is defined as the maximum he could ever pay conditional on allocation \mathbf{x} :*

$$W_i(a_i, \mathbf{x}) = \max_{\mathbf{a}_{-i}: f(\mathbf{a})=\mathbf{x}} p_i(\mathbf{a}) . \quad (4.2)$$

Now, we can state weak smoothness.

Definition 4.5 (Weakly Smooth Mechanism). *A mechanism \mathcal{M} is weakly (λ, μ_1, μ_2) -smooth with respect to utility functions $(u_i^{\theta_i})_{\theta_i \in \Theta_i, i \in N}$ for $\lambda, \mu_1, \mu_2 \geq 0$, if for any type profile $\boldsymbol{\theta} \in \times_i \Theta_i$ and for any action profile \mathbf{a} there exists a randomized action $a_i^*(\boldsymbol{\theta}, a_i)$ for each player i , s.t.:*

$$\sum_i u_i^{\theta_i}(a_i^*(\boldsymbol{\theta}, a_i), \mathbf{a}_{-i}) \geq \lambda \text{OPT}(\boldsymbol{\theta}) - \mu_1 \sum_i p_i(\mathbf{a}) - \mu_2 \sum_i W_i(a_i, f(\mathbf{a})) .$$

We denote by $u_i^{\theta_i}(\mathbf{a})$ the expected utility of a player if \mathbf{a} is a vector of randomized strategies.

Note that (λ, μ) -smoothness implies weak $(\lambda, \mu, 0)$ -smoothness. We get the following generalization of the price-of-anarchy guarantees for equilibria that fulfill the aforementioned no-overbidding assumption on the players' willingness-to-pay:

Theorem 4.6. *If a mechanism is weakly (λ, μ_1, μ_2) -smooth with respect to utility functions $(u_i^{\theta_i})_{\theta_i \in \Theta_i, i \in N}$, then any correlated equilibrium in the full information setting and any Bayes-Nash equilibrium in the Bayesian setting that satisfies*

$$\mathbf{E}_{\mathbf{a}}[W_i(a_i, f(\mathbf{a}))] \leq \mathbf{E}_{\mathbf{a}}[u_i^{\theta_i}(\mathbf{a}) + p_i(\mathbf{a})] \quad (4.3)$$

achieves efficiency of at least a fraction of $\frac{\lambda}{(\mu_2 + \max\{\mu_1, 1\})}$ of $\text{OPT}(\boldsymbol{\theta})$ or of $\mathbf{E}_{\boldsymbol{\theta}}[\text{OPT}(\boldsymbol{\theta})]$, respectively.

In the quasilinear setting, (4.3) simplifies to the no-overbidding assumption

$$\mathbf{E}_{\mathbf{a}}[W_i(a_i, f(\mathbf{a}))] \leq \mathbf{E}_{\mathbf{a}}[v_i(f(\mathbf{a}))] \quad (4.4)$$

that was introduced in [85], and that is a generalization of the no-overbidding assumptions previously used in the literature [23, 16, 19]. That is, players cannot pay more than their respective value, regardless of the other players' actions.

Proof of Theorem 4.6. For the case of a correlated equilibrium in complete information setting, we start from observing that for every a_i in the support of \mathbf{a}

$$\mathbf{E}_{\mathbf{a}_{-i}|a_i} \left[u_i^{\theta_i}(a_i, \mathbf{a}_{-i}) \right] \geq \mathbf{E}_{\mathbf{a}_{-i}|a_i} \left[u_i^{\theta_i}(a'_i, \mathbf{a}_{-i}) \right], \forall a'_i \in \mathcal{A}_i, \forall i .$$

Again, applying the equilibrium property to $a'_i = a_i^*(\boldsymbol{\theta}, a_i)$, we know that for every a_i in the support of \mathbf{a} :

$$\mathbf{E}_{\mathbf{a}_{-i}|a_i} \left[u_i^{\theta_i}(a_i, \mathbf{a}_{-i}) \right] \geq \mathbf{E}_{\mathbf{a}_{-i}|a_i} \left[u_i^{\theta_i}(a_i^*(\boldsymbol{\theta}, a_i), \mathbf{a}_{-i}) \right], \forall i .$$

Taking the expectation over a_i and adding over all players gives

$$\begin{aligned} \mathbf{E}_{\mathbf{a}} \left[\sum_i u_i^{\theta_i}(\mathbf{a}) \right] &\geq \mathbf{E}_{\mathbf{a}} \left[\sum_i u_i^{\theta_i}(a_i^*(\boldsymbol{\theta}, a_i), \mathbf{a}_{-i}) \right] \\ &\geq \lambda \text{OPT}(\boldsymbol{\theta}) - \mu_1 \sum_i p_i(\mathbf{a}) - \mu_2 \sum_i W_i(a_i, f(\mathbf{a})) . \end{aligned}$$

Using (4.3),

$$(1 + \mu_2) \left[\sum_i u_i^{\theta_i}(\mathbf{a}) + \sum_i p_i(\mathbf{a}) \right] \geq \lambda \text{OPT}(\boldsymbol{\theta}) - (\mu_1 - 1) \sum_i p_i(\mathbf{a}) .$$

By doing a case distinction over $\mu_1 \leq 1$ and $\mu_1 > 1$, we get the claimed result. When $\mu_1 \leq 1$, then immediately

$$(1 + \mu_2) \left[\sum_i u_i^{\theta_i}(\mathbf{a}) + \sum_i p_i(\mathbf{a}) \right] \geq \lambda \text{OPT}(\boldsymbol{\theta}) - (\mu_1 - 1) \sum_i p_i(\mathbf{a}) \geq \lambda \text{OPT}(\boldsymbol{\theta}) ,$$

and when $\mu_1 > 1$, we additionally use that $\sum_i u_i^{\theta_i}(\mathbf{a}) \geq 0$. So,

$$\sum_i u_i^{\theta_i}(\mathbf{a}) + \sum_i p_i(\mathbf{a}) \geq \sum_i p_i(\mathbf{a}) ,$$

from where it follows that

$$(1 + \mu_2) \left[\sum_i u_i^{\theta_i}(\mathbf{a}) + \sum_i p_i(\mathbf{a}) \right] \geq \lambda \text{OPT}(\boldsymbol{\theta}) - (\mu_1 - 1) \left[\sum_i u_i^{\theta_i}(\mathbf{a}) + \sum_i p_i(\mathbf{a}) \right] ,$$

and finally

$$\left[\sum_i u_i^{\theta_i}(\mathbf{a}) + \sum_i p_i(\mathbf{a}) \right] \geq \frac{\lambda}{\mu_2 + \mu_1} \cdot \text{OPT}(\boldsymbol{\theta}) .$$

For the incomplete information setting, by using the same arguments as in Theorem 4.3 we arrive at

$$\begin{aligned} \mathbf{E}_{\boldsymbol{\theta}} \left[\sum_i u_i^{\theta_i}(s(\boldsymbol{\theta})) \right] &\geq \lambda \mathbf{E}_{\boldsymbol{\theta}} [\text{OPT}(\boldsymbol{\theta})] - \mu_1 \mathbf{E}_{\boldsymbol{\theta}} \left[\sum_i p_i(s(\boldsymbol{\theta})) \right] \\ &\quad - \mu_2 \mathbf{E}_{\boldsymbol{\theta}} \left[\sum_i W_i(s_i(\boldsymbol{\theta}_i), f(s(\boldsymbol{\theta}))) \right] . \end{aligned}$$

The result now follows by using assumption (4.3) and case distinction, in the same way as in the full information case. \square

4.4 Quasilinear Often Implies Risk-Averse Smoothness

Our main positive result in this chapter is that many price-of-anarchy guarantees that are proved via smoothness in the quasilinear setting transfer to the risk-averse one. First, we consider mechanisms that are (λ, μ) -smooth with respect to quasilinear utility functions. We show that if the deviation strategy \mathbf{a}^* that is used to establish smoothness ensures non-negative utility, then the price-of-anarchy bound extends to risk-averse settings at a constant loss.

Theorem 4.7. *If a mechanism is (λ, μ) -smooth with respect to quasilinear utility functions $(\hat{u}_i^{v_i})_{i \in N, v_i \in \mathcal{V}_i}$ and the actions in the support of the smoothness deviations satisfy*

$$\hat{u}_i(a_i^*, \mathbf{a}_{-i}) \geq 0, \forall \mathbf{a}_{-i}, \forall i,$$

then any correlated equilibrium in the full information setting and any Bayes-Nash equilibrium in the Bayesian setting achieves efficiency of at least a factor of $\frac{\lambda}{2 \cdot \max\{1, \mu\}}$ of the expected optimal social welfare even in the presence of risk averse bidders.

Using Theorem 4.3, it suffices to prove the following lemma.

Lemma 4.8. *If a mechanism is (λ, μ) -smooth with respect to quasilinear utility functions $(\hat{u}_i^{v_i})_{i \in N, v_i \in \mathcal{V}_i}$ and the actions in the support of the smoothness deviations satisfy*

$$\hat{u}_i(a_i^*, \mathbf{a}_{-i}) \geq 0, \forall \mathbf{a}_{-i}, \forall i, \quad (4.5)$$

then the mechanism is $(\lambda/2, \mu)$ -smooth with respect to any normalized risk-averse utility functions $(u_i^{v_i})_{i \in N, v_i \in \mathcal{V}_i}$.

Proof. We start from an arbitrary action profile \mathbf{a} and want to satisfy Definition 4.2. Since there exist smoothness deviations s.t. $\hat{u}_i(a_i^*, \mathbf{a}_{-i}) = v_i(f(a_i^*, \mathbf{a}_{-i})) - p_i \geq 0, \forall \mathbf{a}_{-i}, \forall i$, we know from property (4) of the risk aversion definition that $u_i^{v_i}(a_i^*, \mathbf{a}_{-i}) \geq \hat{u}_i^{v_i}(a_i^*, \mathbf{a}_{-i})$. Therefore,

$$\begin{aligned} \sum_i u_i^{v_i}(a_i^*, \mathbf{a}_{-i}) &\geq \sum_i \hat{u}_i^{v_i}(a_i^*, \mathbf{a}_{-i}) \\ &\geq \lambda \widehat{\text{OPT}} - \mu \sum_i p_i(\mathbf{a}) \\ &\geq \frac{\lambda}{2} \text{OPT} - \mu \sum_i p_i(\mathbf{a}), \end{aligned}$$

where the last inequality follows from Lemma 4.1. \square

Note that in order for (4.5) to hold, it is sufficient if all undominated strategies guarantee non-negative quasilinear utility. For example, in a first-price auction, the only undominated bids are the ones from 0 to v_i . Regardless of the other players' bids, these can never result in negative utility.

Corollary 4.9. *Under normalized risk-averse utilities, the first-price auction has a constant price of anarchy for correlated and Bayes-Nash equilibria.*

In addition, we note here that the first part of Property (4) of the normalization assumption, $u_i^{v_i}(\mathbf{x}, p_i) \geq v_i(\mathbf{x}) - p_i, 0 \leq p_i \leq v_i(\mathbf{x})$, is not crucial for obtaining a result similar to Theorem 4.7. Indeed, a relaxation of the form $u_i^{v_i}(\mathbf{x}, p_i) \geq C \cdot (v_i(\mathbf{x}) - p_i), 0 \leq p_i \leq v_i(\mathbf{x})$ for $0 < C < 1$, C constant, would incur a loss of at most a factor of C in the efficiency bound of Theorem 4.7.

Formally, let us assume the following *relaxed normalized risk-averse utilities*:

- (1) $u_i^{v_i}(\mathbf{x}, p_i) \geq u_i^{v_i}(\mathbf{x}, p'_i)$, if $p_i \leq p'_i$ (monotonicity)
- (2) $u_i^{v_i}(\mathbf{x}, p_i) = 0$, if $p_i = v_i(\mathbf{x})$ (normalization at $p_i = v_i(\mathbf{x})$)
- (3) $u_i^{v_i}(\mathbf{x}, p_i) = v_i(\mathbf{x})$, if $p_i = 0$ (normalization at $p_i = 0$)
- (4) $u_i^{v_i}(\mathbf{x}, p_i) \geq C \cdot (v_i(\mathbf{x}) - p_i)$, if $0 \leq p_i \leq v_i(\mathbf{x})$, $0 < C < 1$, C constant ;
 $u_i^{v_i}(\mathbf{x}, p_i) \leq v_i(\mathbf{x}) - p_i$, otherwise (extra relaxed concavity)

Lemma 4.10. *If a mechanism is (λ, μ) -smooth with respect to quasilinear utility functions $(\hat{u}_i^{v_i})_{i \in N, v_i \in \mathcal{V}_i}$ and the actions in the support of the smoothness deviations satisfy*

$$\hat{u}_i(a_i^*, \mathbf{a}_{-i}) \geq 0, \forall \mathbf{a}_{-i}, \forall i, \quad (4.6)$$

then the mechanism is $(C\lambda/2, C\mu)$ -smooth with respect to any relaxed normalized risk-averse utility functions $(u_i^{v_i})_{i \in N, v_i \in \mathcal{V}_i}$.

Proof. We start from an arbitrary action profile \mathbf{a} and want to satisfy Definition 4.2. Since there exist smoothness deviations s.t. $\hat{u}_i(a_i^*, \mathbf{a}_{-i}) = v_i(f(a_i^*, \mathbf{a}_{-i})) - p_i \geq 0, \forall \mathbf{a}_{-i}, \forall i$, we know from property (4) of the relaxed risk aversion definition that $u_i^{v_i}(a_i^*, \mathbf{a}_{-i}) \geq C \cdot \hat{u}_i^{v_i}(a_i^*, \mathbf{a}_{-i})$. Therefore,

$$\begin{aligned} \sum_i u_i^{v_i}(a_i^*, \mathbf{a}_{-i}) &\geq \sum_i C \cdot \hat{u}_i^{v_i}(a_i^*, \mathbf{a}_{-i}) \\ &\geq C\lambda \cdot \widehat{\text{OPT}} - C\mu \sum_i p_i(\mathbf{a}) \\ &\geq \frac{C\lambda}{2} \cdot \text{OPT} - C\mu \sum_i p_i(\mathbf{a}), \end{aligned}$$

where the last inequality follows from Lemma 4.1. □

Using Theorem 4.3, we obtain the following theorem.

Theorem 4.11. *If a mechanism is (λ, μ) -smooth with respect to quasilinear utility functions $(\hat{u}_i^{v_i})_{i \in N, v_i \in \mathcal{V}_i}$ and the actions in the support of the smoothness deviations satisfy $\hat{u}_i(a_i^*, \mathbf{a}_{-i}) \geq 0, \forall \mathbf{a}_{-i}, \forall i$, then any Correlated Equilibrium in the full information setting and any Bayes-Nash Equilibrium in the Bayesian setting achieves efficiency of at least a fraction of $\frac{C\lambda}{2 \cdot \max\{1, C\mu\}}$ of the expected optimal social welfare even in the presence of risk averse bidders.*

4.4.1 Weak Smoothness

We will assume the following pointwise condition:

Definition 4.12 (Pointwise No-Overbidding). *A randomized strategy profile \mathbf{a} satisfies the pointwise no-overbidding assumption if for every player i and every action in the support of \mathbf{a} the following holds:*

$$W_i(a_i, \mathbf{x}) := \max_{\mathbf{a}_{-i}: f(\mathbf{a})=\mathbf{x}} p_i(\mathbf{a}) \leq v_i(\mathbf{x}) , \quad (4.7)$$

i.e., no player is pointwise bidding in a way that he could potentially pay more than his value, subject to his allocation remaining the same.

Theorem 4.13. *If a mechanism is weakly (λ, μ_1, μ_2) -smooth with respect to quasilinear utility functions $(\hat{u}_i^{v_i})_{i \in N, v_i \in \mathcal{V}_i}$, the actions in the support of the smoothness deviations satisfy $\hat{u}_i(a_i^*, \mathbf{a}_{-i}) \geq 0, \forall \mathbf{a}_{-i}, \forall i$, then any correlated equilibrium in the full information setting and any Bayes-Nash equilibrium in the Bayesian setting that satisfies the pointwise no-overbidding assumption achieves efficiency of at least a fraction of $\frac{\lambda}{2 \cdot (\mu_2 + \max\{\mu_1, 1\})}$ of the expected optimal social welfare even in the presence of risk-averse bidders.*

Proof. First, we show that weak smoothness with respect to quasilinear utility functions with the additional constraint that players have non-negative utility from the smoothness deviation implies weak smoothness with respect to risk-averse players.

Lemma 4.14. *If a mechanism is weakly (λ, μ_1, μ_2) -smooth with respect to quasilinear utility functions $(\hat{u}_i^{v_i})_{i \in N, v_i \in \mathcal{V}_i}$ and the actions in the support of the smoothness deviations satisfy $\hat{u}_i(a_i^*, \mathbf{a}_{-i}) \geq 0, \forall \mathbf{a}_{-i}, \forall i$, then the mechanism is weakly $(\lambda/2, \mu_1, \mu_2)$ -smooth with respect to risk-averse utility functions $(u_i^{v_i})_{i \in N, v_i \in \mathcal{V}_i}$.*

Proof. We start from an arbitrary action profile \mathbf{a} and want to satisfy Definition 4.2. Since there exist smoothness deviations s.t. $\hat{u}_i(a_i^*, \mathbf{a}_{-i}) = v_i(f(a_i^*, \mathbf{a}_{-i})) - p_i \geq 0, \forall \mathbf{a}_{-i}, \forall i$, we know from property (4) of the risk aversion definition that $u_i^{v_i}(a_i^*, \mathbf{a}_{-i}) \geq \hat{u}_i^{v_i}(f(a_i^*, \mathbf{a}_{-i}))$. Therefore,

$$\begin{aligned} \sum_i u_i^{v_i}(a_i^*, \mathbf{a}_{-i}) &\geq \sum_i \hat{u}_i^{v_i}(a_i^*, \mathbf{a}_{-i}) \\ &\geq \lambda \widehat{\text{OPT}} - \mu_1 \sum_i p_i(\mathbf{a}) - \mu_2 \sum_i W_i(a_i, f(\mathbf{a})) \\ &\geq \frac{\lambda}{2} \text{OPT} - \mu_1 \sum_i p_i(\mathbf{a}) - \mu_2 \sum_i W_i(a_i, f(\mathbf{a})) , \end{aligned}$$

where the last inequality follows from Lemma 4.1. □

Next, we will show that pointwise no-overbidding indeed implies the no-overbidding assumption (4.3):

Using the pointwise no-overbidding assumption $v_i(\mathbf{x}) \geq p_i$, we know that $u_i^{v_i}(\mathbf{x}, p_i) \geq v_i(\mathbf{x}) - p_i$. From here, $W_i(a_i, \mathbf{x}) \leq v_i(\mathbf{x}) \leq u_i^{v_i}(\mathbf{a}) + p_i(\mathbf{a})$, so we can conclude that

$$\mathbf{E}_{\mathbf{a}}[W_i(a_i, f(\mathbf{a}))] \leq \mathbf{E}_{\mathbf{a}}[u_i^{v_i}(\mathbf{a}) + p_i(\mathbf{a})] .$$

Theorem 4.6 now completes the proof. □

Using that the second-price auction is weakly $(1, 0, 1)$ -smooth with respect to quasilinear utilities, we immediately get that its price of anarchy is also constant in the risk-averse setting.

Corollary 4.15. *Under normalized risk-averse utilities, the second-price auction has a constant price of anarchy for correlated and Bayes-Nash equilibria with pointwise no-overbidding.*

4.4.2 Budget Constraints

The techniques and results so far have striking similarities to settings with budget constraints, where players do not have quasilinear preferences already in the risk neutral case. As it turns out, under very mild additional assumptions, we can also add (a generalized form of) hard budget constraints to our consideration.

We now assume that types are pairs $\theta_i = (v_i, B_i)$, where $B_i: \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ is an outcome-dependent budget function. Depending on which outcome is achieved, the agent may have different amounts of liquidity. We assume that there is a normalized risk-averse utility function $u_i^{v_i}$ such that for a player of type $\theta_i = (v_i, B_i)$

$$u_i^{\theta_i}(\mathbf{x}, p_i) = \begin{cases} u_i^{v_i}(\mathbf{x}, p_i), & \text{if } p_i \leq B_i(\mathbf{x}), \\ -\infty, & \text{otherwise .} \end{cases} \quad (4.8)$$

In the budgeted setting, one cannot hope to achieve full welfare. This is due to low budget participants not being able to maximize their contribution. Therefore, we will replace $\text{OPT}(\boldsymbol{\theta})$ in the price-of-anarchy and smoothness definition by the optimal *effective* or *liquid* welfare, given as $\max_{\mathbf{x}, \mathbf{p}} \sum_i \min\{u_i^{\theta_i}(\mathbf{x}, p_i) + p_i, B_i(\mathbf{x})\}$. This benchmark, introduced in [32], reflects that players with low budgets cannot be expected to be effective at maximizing their own value.

The effect of budgets on efficiency in the risk neutral case was already studied in [85], where the authors, in order to be able to prove efficiency bounds, introduced the notion of a *conservatively smooth mechanism* that has the following additional assumption on the smoothness deviations:

$$\max_{\mathbf{a}_{-i}} p_i(a_i^*(\mathbf{v}, a_i), \mathbf{a}_{-i}) \leq \max_{\mathbf{x}} v_i(\mathbf{x}) . \quad (4.9)$$

Conservatively smooth mechanisms are then shown to allow the budgeted scenario without any further loss of efficiency. Note that (4.9) is a weaker assumption than the Condition (4.5) we ask for. Therefore, we can easily extend our results for risk-averse bidders to the budgeted setting.

Our main result is that if the type space is chosen in a way that taking the pointwise minimum of a valuation function and a budget function yields again a feasible valuation function, meaning that we stay within the “permitted” valuation space when applying the budget constraints, then the price-of-anarchy guarantee is again preserved. The valuation space being closed under capping is a crucial requirement both for our result and the result in [85].

Theorem 4.16. *If a mechanism is (λ, μ) -smooth with respect to quasilinear utility functions $(\hat{u}_i^{v_i})_{i \in N, v_i \in \mathcal{V}_i}$, its valuation space is closed under capping with budget functions,*

the actions in the support of the smoothness deviations satisfy $\hat{u}_i(a_i^*, \mathbf{a}_{-i}) \geq 0, \forall \mathbf{a}_{-i}, \forall i$, then the social welfare at any correlated equilibrium and at any Bayes-Nash equilibrium is at least a fraction of $\frac{\lambda}{2 \cdot \max\{1, \mu\}}$ of the expected maximum effective welfare even in the presence of risk-averse bidders.

For the proof, as before, we show a lemma connecting smoothness with respect to quasilinear utilities to smoothness with respect to risk-averse ones.

Lemma 4.17. *If a mechanism is (λ, μ) -smooth with respect to quasilinear utility functions $(\hat{u}_i^{v_i})_{i \in N, v_i \in \mathcal{V}_i}$, its valuation space is closed under capping with the budget functions, and the actions in the support of the smoothness deviations satisfy $\hat{u}_i(a_i^*, \mathbf{a}_{-i}) \geq 0, \forall \mathbf{a}_{-i}, \forall i$, then the mechanism is $(\lambda/2, \mu)$ -smooth with respect to risk-averse budgeted utility functions $(u_i^{\theta_i})_{\theta_i \in \Theta_i, i \in N}$.*

Proof. We start from an arbitrary action profile \mathbf{a} and keep in mind that the risk-averse budgeted utility function $u_i^{\theta_i}$ has type $\theta_i = (v_i, B_i)$. By $\hat{u}^{\bar{v}_i}$, we denote the quasilinear utility of player i with the capped valuation function \bar{v}_i . Formally,

$$\hat{u}_i^{\bar{v}_i}(\mathbf{x}, p_i) = \bar{v}_i(\mathbf{x}) - p_i = \min\{v_i(\mathbf{x}), B_i(\mathbf{x})\} - p_i .$$

Since the valuation space is closed under capping with the budget function, we can find smoothness deviations $a_i^*(\bar{\mathbf{v}}, a_i)$ s.t. $\hat{u}_i^{\bar{v}_i}(f(a_i^*, \mathbf{a}_{-i}), p_i(a_i^*, \mathbf{a}_{-i})) = \bar{v}_i(f(a_i^*, \mathbf{a}_{-i})) - p_i(a_i^*, \mathbf{a}_{-i}) \geq 0, \forall \mathbf{a}_{-i}$ and therefore $u_i^{\bar{v}_i}(a_i^*, \mathbf{a}_{-i}) \geq \hat{u}_i^{\bar{v}_i}(a_i^*, \mathbf{a}_{-i})$. It follows that

$$\begin{aligned} \sum_i u_i^{\theta_i}(a_i^*(\bar{\mathbf{v}}, a_i), \mathbf{a}_{-i}) &= \sum_i u_i^{v_i}(a_i^*(\bar{\mathbf{v}}, a_i), \mathbf{a}_{-i}) \geq \sum_i u_i^{\bar{v}_i}(a_i^*, \mathbf{a}_{-i}) \geq \sum_i \hat{u}_i^{\bar{v}_i}(a_i^*, \mathbf{a}_{-i}) \\ &\geq \lambda \cdot \widehat{\text{OPT}}_{\bar{\mathbf{v}}} - \mu \cdot \sum_i p_i(\mathbf{a}) \geq \frac{\lambda}{2} \cdot \text{OPT}_{\bar{\mathbf{v}}} - \mu \cdot \sum_i p_i(\mathbf{a}) , \end{aligned}$$

where the first equality holds because the deviations are such that the payments are below the budgets, the first inequality because

$$u_i^{v_i}(\mathbf{a}) = h(v_i(f(\mathbf{a})) - p_i(\mathbf{a})) \geq h(\min\{v_i(f(\mathbf{a})), B_i\} - p_i(\mathbf{a})) = u_i^{\bar{v}_i}(\mathbf{a}), \forall \mathbf{a} ,$$

and the third because the valuation space is closed under capping. \square

The generality of Theorem 4.3 allows us to now obtain Theorem 4.16. Note that $\text{OPT}_{\bar{\mathbf{v}}}$, where $\bar{\mathbf{v}}$ is the vector of capped valuation functions, indeed aligns correctly with the effective welfare benchmark.

4.5 Unbounded Price of Anarchy for All-Pay Auctions

From the previous section, we infer that the constant price-of-anarchy bounds for first-price and second-price auctions immediately extend to the risk-averse setting. This is not true for all-pay auctions; by definition there is no non-trivial bid that always ensures non-negative utility. Indeed, as we show in this section, the price of anarchy is unbounded in the presence of risk-averse players.

Theorem 4.18. *In an all-pay auction with risk-averse players, the price of anarchy is unbounded.*

The general idea is to construct a Bayes-Nash equilibrium with two players that very rarely have high values and only then bid high values. We then add a third player who always has a high value. However, as the first two players bid high values occasionally, there is no possible bid that ensures he will surely win. This means, any bid has a small probability of not getting the item but having to pay. Risk-averse players are more inclined to avoid this kind of lotteries. In particular, making our third player risk-averse enough, he prefers the sure zero utility of not participating to any way of bidding that always comes with a small probability of negative utility.

Proof of Theorem 4.18. We consider two (mildly) risk-averse players who both have the same valuation distributions and a third (very) risk-averse player with a constant value. For a large number $M > 5$, the first two players have values v_1 and v_2 drawn independently from distributions with density functions of value $2 \cdot (1 - (M - 1) \cdot \varepsilon)$ on the interval $[1/2, 1)$ and value ε on the interval $[1, M]$, where $\varepsilon = 1/M^2$. The third player always has value $1/3 \cdot \ln(M/2)$ for winning.

We will construct a symmetric pure Bayes-Nash equilibrium involving only the first two players. It will be designed such that for the third player it is a best response to always bid 0, i.e., to opt out of the mechanism and never win the item. So, the combination of these strategies will be a pure Bayes-Nash equilibrium for all three players.

Note that the social welfare of any equilibrium of this form is upper-bounded by the optimal social welfare that can be achieved by the first two bidders. By Lemma 4.1, it is bounded by

$$\mathbf{E}[\text{SW}] \leq 2 \cdot \mathbf{E}[\max\{v_1, v_2\}] \leq 2 \cdot \mathbf{E}[v_1 + v_2] = 2 \cdot (\mathbf{E}[v_1] + \mathbf{E}[v_2]) = 4 \cdot \mathbf{E}[v_1] \leq 4 .$$

Furthermore, the third player's value $v_3 = 1/3 \cdot \ln(M/2)$ is a lower bound to the optimal social welfare in the construction containing all three players. So, as pointwise $\text{OPT}(v) \geq 1/3 \cdot \ln(M/2)$, where $v = (v_1, v_2, v_3) \in \mathcal{V}$ denotes the valuation profile, this implies that the price of anarchy can be arbitrarily high.

We define the utility functions by setting

$$u_i^{v_i}(b_i) = \begin{cases} \frac{h_i(v_i - b_i)}{h(v_i)} \cdot v_i, & \text{if } b_i \text{ is the winning bid,} \\ \frac{h_i(-b_i)}{h(v_i)} \cdot v_i, & \text{otherwise .} \end{cases} \quad (4.10)$$

For the first two players, we use $h_i(x) := 1 - e^{-x}$, $i \in \{1, 2\}$, which is increasing and concave. For the third player, we set $h_i(x) = x$ for $x \geq 0$ and $h_i(x) = C \cdot x$ for $x < 0$, where $C = (16 \cdot \frac{1}{3} \cdot \ln M/2) \cdot M^2 \geq 1$. Again this function is increasing and concave. We further note that its slope is not an absolute constant. This is indeed necessary because the price of anarchy can be bounded in terms of the slopes of the h_i -functions as we show in the following Lemma.

Lemma 4.19. *In an all-pay auction with risk-averse players whose utilities are of the form $h(v_i(\mathbf{x}) - p_i)$, where h is a concave function s.t. $h(x) = C \cdot x$ for $x < 0$, $C \geq 1$ constant, the Price of Anarchy is at most $4(C + 1)$.*

Proof. We use the following smoothness deviation: The highest value player with value v_h deviates to $\frac{1}{2}v_h$ and everybody else to 0. Now, it is easy to see that the following inequality holds independent of whether the highest value player obtains the item or not

$$u_h^{v_h}\left(\frac{v_h}{2}, \mathbf{a}_{-i}\right) \geq \frac{1}{2}v_h - (C+1) \max_{i \neq h} a_i \geq \frac{1}{2}\widehat{\text{OPT}} - (C+1) \sum_i a_i ,$$

so then

$$\sum_i u_i^{v_i}(a_i^*, \mathbf{a}_{-i}) \geq \frac{1}{2}\widehat{\text{OPT}} - (C+1) \sum_i p_i(\mathbf{a}) \geq \frac{1}{4}\text{OPT} - (C+1) \sum_i p_i(\mathbf{a}) .$$

The claim follows by applying Theorem 4.3. \square

Note that the utility functions also satisfy normalizations at $p_i = v_i(\mathbf{x})$ and at $p_i = 0$.

In conclusion, we see that in our example risk aversion has the effect of heavily penalizing payments without winning the auction.

Claim 1. *With the third player not participating, it is a symmetric pure Bayes-Nash equilibrium for the first two players to play according to bidding function $\beta: \mathcal{V}_i \rightarrow \mathbb{R}_{\geq 0}, i \in \{1, 2\}$, such that*

$$\beta(x) = \int_{\frac{1}{2}}^x \frac{f(t)(e^t - 1)}{F(t) + (1 - F(t))e^t} dt , \quad (4.11)$$

where F denotes the cumulative distribution function of the value and f denotes its density.

Proof. We will argue that playing according to β is always the unique best response if the other player is playing according to β , too. Due to symmetry reasons, it is enough to argue about the first player.

Let us fix player 1's value $v_1 = x$ and consider the function $g: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ that is defined by $g(y) = \mathbf{E}[u_1^x(b_1 = y, b_2 = \beta(v_2), b_3 = 0)]$. We claim that g is indeed maximized at $y = \beta(x)$. We have¹

$$\begin{aligned} g(y) &= \mathbf{Pr}[\beta(v_2) \leq y] \cdot \frac{h_1(x-y)}{h_1(x)} \cdot x + (1 - \mathbf{Pr}[\beta(v_2) \leq y]) \cdot \frac{h_1(-y)}{h_1(x)} \cdot x \\ &= \frac{x}{h_1(x)} \left[F(\beta^{-1}(y)) (h_1(x-y) - h_1(-y)) + h_1(-y) \right] \\ &= xe^y F(\beta^{-1}(y)) + \frac{x(1 - e^y)}{1 - e^{-x}} . \end{aligned}$$

The first derivative of this function is given by

$$g'(y) = xe^y F(\beta^{-1}(y)) + xe^y \frac{d}{dy} F(\beta^{-1}(y)) - \frac{x}{1 - e^{-x}} e^y .$$

The inverse function theorem implies

$$\frac{d}{dy} F(\beta^{-1}(y)) = \frac{f(\beta^{-1}(y))}{\beta'(\beta^{-1}(y))} .$$

¹Note that the first step assumes tie breaking in favor of player 1. This is irrelevant for the future steps as the involved probability distributions are continuous.

Furthermore, as $\beta'(t) = \frac{f(t)(e^t-1)}{F(t)+(1-F(t))e^t}$, we get for all t that

$$\frac{f(t)}{\beta'(t)} = \frac{F(t) + (1 - F(t))e^t}{e^t - 1} = (1 - F(t)) + \frac{1}{e^t - 1} .$$

This simplifies $g'(y)$ to

$$g'(y) = xe^y + \frac{xe^y}{e^{\beta^{-1}(y)} - 1} - \frac{xe^y}{1 - e^{-x}} = \frac{xe^y}{(1 - e^{-x})(e^{\beta^{-1}(y)} - 1)} \left(1 - e^{-x+\beta^{-1}(y)}\right) .$$

Notice that the factor $\frac{xe^y}{(1-e^{-x})(e^{\beta^{-1}(y)}-1)}$ is always positive. Therefore, we observe that $g'(y) = 0$ if and only if $e^{-x+\beta^{-1}(y)} = 1$, which is equivalent to $y = \beta(x)$. Furthermore, $g'(y) > 0$ for $y < \beta(x)$ and $g'(y) < 0$ for $y > \beta(x)$. This means that $y = \beta(x)$ has to be the (unique) global maximum of $g(y)$. \square

Claim 2. *If the first two players are bidding according to (4.11), then it is a best response for the third player to always bid 0.*

Proof. We now show that the very risk-averse third player with valuation $1/3 \cdot \ln(M/2)$ will indeed bid 0 because every bid $b'_3 > 0$ will cause negative expected utility.

We distinguish two cases. For values of $b'_3 > \frac{1}{16}$, we use that with a small probability one of the two remaining players has a valuation of at least $M - 1$, which leads to negative utility. For $b'_3 \leq \frac{1}{16}$ on the other hand, he loses so often that his expected utility is again negative.

Let us first assume that the third player bids b'_3 with $\frac{1}{16} < b'_3 \leq v_3$. In this case, with probability more than ε one of the first two players has value of at least $M - 1$. The bid of this player with $v_i \geq M - 1$ can be estimated as follows

$$\begin{aligned} \beta(v_i) &\geq \beta(M - 1) \\ &\geq \int_{M/2}^{M-1} \frac{f(t)(e^t - 1)}{1 + (1 - F(t))e^t} dt \\ &= \int_{M/2}^{M-1} \frac{\varepsilon(e^t - 1)}{1 + \varepsilon(M - t)e^t} dt \\ &\geq \frac{1}{2}(1 - e^{-\frac{M}{2}}) \int_{M/2}^{M-1} \frac{\varepsilon e^t}{\varepsilon(M - t)e^t} dt \\ &= \frac{1}{2}(1 - e^{-\frac{M}{2}}) \ln(M/2) \\ &> \frac{1}{3} \ln(M/2) , \end{aligned}$$

which means that by bidding b'_3 the third player loses with probability of at least $\varepsilon = 1/M^2$.

For the expected utility, this implies

$$\begin{aligned}
 \mathbf{E} [u_3(b'_3, \mathbf{b}_{-3})] &\leq (1 - \varepsilon) \cdot \left(\frac{1}{3} \cdot \ln M/2 - b'_3 \right) + \varepsilon \cdot \left(- \left(16 \cdot \frac{1}{3} \cdot \ln M/2 - 1 \right) \cdot M^2 \cdot b'_3 \right) \\
 &< \frac{1}{3} \cdot \ln M/2 - b'_3 - \frac{1}{M^2} \cdot \left(16 \cdot \frac{1}{3} \cdot \ln M/2 - 1 \right) \cdot M^2 \cdot b'_3 \\
 &= \frac{1}{3} \cdot \ln M/2 - b'_3 \left(1 + 16 \cdot \frac{1}{3} \cdot \ln M/2 - 1 \right) \\
 &< \frac{1}{3} \ln M/2 - \frac{1}{16} \cdot 16 \cdot \frac{1}{3} \ln M/2 \\
 &= 0 .
 \end{aligned}$$

In the case where the third player bids b'_3 , $0 < b'_3 \leq \frac{1}{16}$, we need to be a bit more careful with estimating the winning probability. We first give a lower bound on the bidding function of the first player for $v_1 < 1$

$$\begin{aligned}
 \beta(v_1) &\geq \int_{1/2}^{v_1} \frac{2(1 - \frac{M-1}{M^2})(e^t - 1)}{2(t - \frac{1}{2})(1 - \frac{M-1}{M^2}) + (1 - 2(t - \frac{1}{2})(1 - \frac{M-1}{M^2})) \cdot e^t} dt \\
 &> \int_{1/2}^{v_1} \frac{\frac{3}{2}(e^t - 1)}{2t - 1 + 2 \cdot e^t} dt \\
 &= \frac{3}{4} \int_{1/2}^{v_1} \frac{e^t - 1}{e^t + t - \frac{1}{2}} dt \\
 &\geq \frac{3}{4} \int_{1/2}^{v_1} \left(1 - \frac{1}{\sqrt{e}} \right) dt \\
 &= \frac{3}{4} \left(1 - \frac{1}{\sqrt{e}} \right) \left(v_1 - \frac{1}{2} \right) \\
 &> \frac{1}{4} \left(v_1 - \frac{1}{2} \right) .
 \end{aligned}$$

For $v_1 \geq 1$, $\beta(v_1) > \frac{1}{16}$ with probability 1. This implies that with b'_3 , the third player has a winning probability of at most

$$\mathbf{Pr} [\beta(v_1) \leq b'_3] \leq \mathbf{Pr} \left[\frac{1}{4} \left(v_1 - \frac{1}{2} \right) \leq b'_3 \right] = \mathbf{Pr} \left[v_1 \leq 4b'_3 + \frac{1}{2} \right] < 2 \cdot 4b'_3 .$$

Now, having in mind that $C = (16v_3) \cdot M^2 \geq 32 \cdot v_3$, the utility can be estimated as

follows

$$\begin{aligned}
 \mathbf{E} [u_3(b'_3, \mathbf{b}_{-3})] &\leq \mathbf{Pr} [\beta(v_1) \leq b'_3] \cdot (v_3 - b'_3) - \mathbf{Pr} [\beta(v_1) > b'_3] \cdot 32 \cdot v_3 \cdot b'_3 \\
 &< 2 \cdot 4b'_3(v_3 - b'_3) - (1 - 2 \cdot 4b'_3) \cdot 32 \cdot v_3 \cdot b'_3 \\
 &= 8b'_3v_3 - 8(b'_3)^2 - 32b'_3v_3 + 8 \cdot 32v_3(b'_3)^2 \\
 &= 8b'_3(-3v_3 + b'_3(32v_3 - 1)) \\
 &\leq 8b'_3 \left(-3v_3 + 2v_3 - \frac{1}{16} \right) \\
 &= 8b'_3 \left(-v_3 - \frac{1}{16} \right) \\
 &< 0 .
 \end{aligned}$$

So also in this case, the expected utility is negative. \square

Combining the two claims, we have constructed a class of distributions and Bayes-Nash equilibria with unbounded price of anarchy. \square

As a final remark, we note that the first two bidders occasionally bid high only due to risk aversion. In a symmetric Bayes-Nash equilibrium of the all-pay auction in the quasilinear setting, all bids are always bounded by the expected value of a player. Therefore, such an equilibrium would not work as a point of departure.

Claim 3. *In a symmetric BNE of the all-pay auction in the quasilinear setting, all bids are bounded by the expected value of a player.*

Proof. Due to symmetry, it is enough to argue about the first player. Let β denote the equilibrium bidding function. We fix player 1's value $v_1 = x$ and consider his expected utility for bidding y :

$$\begin{aligned}
 \mathbf{E} [u_1^x(b_1 = y, b_2 = \beta(v_2))] &= \mathbf{Pr} [\beta(v_2) < y] \cdot (x - y) + \mathbf{Pr} [\beta(v_2) \geq y] \cdot (-y) \\
 &= \mathbf{Pr} [\beta(v_2) < y] \cdot x - y \\
 &= F(\beta^{-1}(y)) \cdot x - y .
 \end{aligned}$$

By taking the derivative and setting it to zero, we arrive at

$$\frac{f(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))} \cdot x - 1 = 0 ,$$

so

$$\beta'(x) = x \cdot f(x) .$$

Now it is obvious that

$$\beta(x) = \int_0^x t \cdot f(t) \leq \mathbf{E} [v_1] . \quad \square$$

4.6 Variance-Aversion Model

In this section, we consider a different model that tries to capture the effect that agents prefer certain outcomes to uncertain ones. It is inspired by similar models in game theory and penalizes variance of random variables. Rather than reflecting the aversion in the utility functions, it is modeled by adapting the solution concept.

In the usual definition of equilibria involving randomization, the utility of a randomized strategy profile is set to be the expectation over the pure strategies. The definition we consider here is modified by subtracting the respective standard deviation. For a player i , the utility of a randomized strategy profile \mathbf{a} is given as

$$u_i^{v_i}(\mathbf{a}) = \mathbf{E}_{\mathbf{b} \sim \mathbf{a}}[\hat{u}_i^{v_i}(\mathbf{b})] - \gamma \sqrt{\text{Var}[\hat{u}_i^{v_i}(\mathbf{b})]} ,$$

so a player's utility for an action profile is his expected quasilinear utility for this profile minus the standard deviation multiplied by a parameter γ that determines the degree of variance-averseness, $0 \leq \gamma \leq 1$. As already mentioned, $\hat{u}_i(\mathbf{a})$ denotes the quasilinear utility of player i for the action profile \mathbf{a} .

Bayes-Nash equilibria and correlated equilibria can be defined the same way as before, always replacing expectations by the difference of expectation and standard deviation. The formal definition for $s(\mathbf{v})$ being a Bayes-Nash equilibrium in this setting is that $\forall i \in N, \forall v_i \in \Theta_i, a_i \in \mathcal{A}_i$,

$$\begin{aligned} \mathbf{E}_{\mathbf{v}_{-i}}[\hat{u}_i^{v_i}(s(\mathbf{v})) \mid v_i] - \gamma \sqrt{\text{Var}[\hat{u}_i^{v_i}(s(\mathbf{v})) \mid v_i]} \\ \geq \mathbf{E}_{\mathbf{v}_{-i}}[\hat{u}_i^{v_i}(a_i, s_{-i}(\mathbf{v}_{-i})) \mid v_i] - \gamma \sqrt{\text{Var}[\hat{u}_i^{v_i}(a_i, s_{-i}(\mathbf{v}_{-i})) \mid v_i]} . \end{aligned}$$

Note that we again evaluate social welfare as agents perceive it. That is, for a randomized strategy profile \mathbf{a} , we set

$$\text{SW}^{\mathbf{v}}(\mathbf{a}) = \sum_i u_i^{v_i}(\mathbf{a}) + \sum_i p_i(\mathbf{a}) .$$

Our first result shows that first-price and notably also all-pay auctions have a constant price of anarchy in this setting. Note that, even though the proof looks a lot like smoothness proofs, it is not possible to phrase it within the smoothness framework, since here we are dealing with a different solution concept.

Theorem 4.20. *Bayes-Nash equilibria and correlated equilibria of the first-price and all-pay auction have a constant price of anarchy in the variance-aversion model.*

Proof. For simplicity, we will show the claim only for Bayes-Nash equilibria. The proof for correlated equilibria works the same way with minor modifications to the notation.

Assume \mathbf{b} is a Bayes-Nash equilibrium. We claim that $\mathbf{E}_{\mathbf{v}}[\text{SW}^{\mathbf{v}}(\mathbf{b})] \geq \frac{1}{16} \cdot \mathbf{E}_{\mathbf{v}}[\text{OPT}]$, where OPT denotes the value of social welfare in the allocation that maximizes it, i.e., maximized sum of utility and payments of the agents.

Consider a fixed player j and a fixed valuation v_j . Let $q = \Pr[\max_{i \neq j} b_i \leq \frac{1}{4} \cdot v_j]$ denote the probability that no other player's bid exceeds $\frac{1}{4} \cdot v_j$.

Assume first that $q \leq \frac{3}{4}$. Then, because the total social welfare is lower bounded by the payments, we have

$$\begin{aligned} \mathbb{E}_{\mathbf{v}_{-j}|v_j} [\text{SW}^{\mathbf{v}}(\mathbf{b})] &\geq \mathbb{E}_{v_{-i}|v_i} \left[\hat{u}_i^{v_i}(0, b_{-i}) - \gamma \sqrt{\text{Var}[\hat{u}_i^{v_i}(0, b_{-i})]} \right] + \mathbb{E}_{v_{-i}|v_i} \left[\sum_i p_i(b) \right] \\ &\geq \mathbb{E}_{v_{-j}|v_j} \left[\sum_i p_i(b) \right] \geq (1-q) \frac{1}{4} v_j \geq \frac{1}{16} v_j . \end{aligned}$$

On the other hand, if $q \geq \frac{3}{4}$, we use that $\mathbb{E}_{\mathbf{v}_{-j}|v_j} [\text{SW}^{\mathbf{v}}(\mathbf{b})] \geq \mathbb{E}_{\mathbf{v}_{-j}|v_j} \left[u_j^{v_j} \left(\frac{v_j}{4}, \mathbf{b}_{-j} \right) \right]$. It follows

$$\mathbb{E}_{\mathbf{v}_{-j}|v_j} [\text{SW}^{\mathbf{v}}(\mathbf{b})] \geq v_j q - \frac{1}{4} v_j - \gamma v_j \sqrt{q(1-q)} \geq \left(\frac{2 - \gamma\sqrt{3}}{4} \right) v_j \geq \frac{1}{16} v_j ,$$

where the first inequality is in fact an equality for the all-pay auction.

From here, by taking the expectation over v_j and by weighing the right hand side by the probability that OPT takes a particular agent, the theorem follows. \square

This is contrasted by a correlated equilibrium with 0 social welfare in a setting with positive values. Indeed, for the special case of $\lambda = 1$, we see that the variance-averse model further differs from the risk-averse model described in previous sections.

Observation 4.21. *The price of anarchy for correlated equilibria of second price auctions is unbounded if $\gamma = 1$.*

Proof. Consider two bidders that both have a valuation of 1. They will be in an equilibrium if they both bid 1 with probability $\frac{1}{2}$ and 0 with the remaining probability, but in a correlated manner, such that always just one player submits a non-zero bid. Let us now calculate the utilities:

$$u_i(\mathbf{b}) = \mathbb{E}_{\mathbf{a} \sim \mathbf{b}} [\hat{u}_i(\mathbf{a})] - \sqrt{\mathbb{E}[\hat{u}_i^2(\mathbf{a})] - (\mathbb{E}[\hat{u}_i(\mathbf{a})])^2} = \frac{1}{2} - \sqrt{\frac{1}{2} - \left(\frac{1}{2}\right)^2} = \frac{1}{2} - \frac{1}{2} = 0 .$$

Since the payments are also 0, the social welfare in this equilibrium is 0, meaning that the price of anarchy is unbounded. \square

This is not only a difference between smoothness and weak smoothness. Our final result is a mechanism that is (λ, μ) -smooth for constant λ and μ but has unbounded price of anarchy.

Theorem 4.22. *For any constant $\gamma > 0$ there is a mechanism that is (λ, μ) -smooth with respect to quasilinear utility functions for constant λ and μ but has unbounded price of anarchy in the variance-aversion model.*

Proof. Consider a setting with two items and two players, who have unit-demand valuation functions (see Definition 3.5) such that $\frac{1}{c} v_{i,1} \leq v_{i,2} \leq c v_{i,1}$ for constant $c \geq 1$. The players' possible actions are to either report one of the two items as preferred or to opt out entirely. Our mechanism first assigns player 1 his (claimed) favorite item, then

assigns player 2 the remaining one unless he opts out. There are no payments. Obviously, this mechanism is $(\frac{1}{c}, 0)$ -smooth because the allocation is within a $\frac{1}{c}$ -factor of the optimal allocation by construction of the valuation functions.

We will now construct a mixed Nash equilibrium of bad welfare. To this end, let $v_{1,1} = v_{1,2} = \varepsilon$ for some small $\varepsilon > 0$. This makes player 1 indifferent between items 1 and 2. In particular, it is a best response to ask for item 1 with probability $\frac{q-1}{q}$ and for item 2 with probability $\frac{1}{q}$. We note at this point that in a Bayes-Nash equilibrium we could make this respective action the unique best response by having random types.

For player 2, we set $v_{2,1} = c$, $v_{2,2} = 1$. She has the choice of participating or opting out. Opting out implies utility 0, whereas participating implies

$$u_2(\mathbf{a}) = \frac{c+q-1}{q} - \gamma \sqrt{\frac{(c-1)^2(q-1)}{q^2}} = \frac{(c-1)(1-\gamma\sqrt{q-1})}{q} + 1$$

Now, if we set $q = c - 1$, then $u_2(\mathbf{a}) = 2 - \gamma\sqrt{c-2}$ which is negative for $c > \frac{4}{\gamma^2} + 2$. We further set $c = \frac{4}{\gamma^2} + 3$. That is, player 2 prefers to opt out. This outcome has social welfare ε whereas the optimal social welfare is c . \square

Note that this last example shows that variance-averseness yields very strange preferences for lotteries. In our example, the variance-averse player prefers not to participate although any outcome in the (free) lottery has positive value.

CHAPTER 5

Simultaneous Composition of Mechanisms with Admission

5.1 Model

There are n bidders that participate in m simultaneous mechanisms. The action space of each player will be limited to bids. Each mechanism $j \in [m]$ is therefore a triple $\mathcal{M}_j = (B_j, f_j, p_j)$, consisting of the bid space B_j , an outcome function and payment functions. More formally, function $f_j: B_j \rightarrow \mathcal{X}_j$ maps every bid vector $b_{\cdot,j}$ on mechanism j into an outcome space \mathcal{X}_j . The function $p_j = (p_{1,j}, \dots, p_{n,j})$ defines a payment for each bidder. That is, depending on the bid vector, $p_{i,j}: B_j \rightarrow \mathbb{R}_{\geq 0}$ defines the non-negative payment for bidder i in mechanism j . For convenience, we denote by $f = (f_j)_{j \in [m]}$ the composed mechanism and by $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_m$ its outcome space. We will not notationally distinguish between scalars and vectors in this chapter, as we will mostly be working with vectors anyway.

We consider a repeated framework with oblivious learning in a simultaneous composition of mechanisms with availabilities. There are T rounds and in each round the bidders participate in m simultaneous mechanisms. In round $t = 1, \dots, T$, each bidder places a bid $b_{i,j}^t$ for each mechanism, the mechanism determines the outcome and the payments, and bidder i has a quasilinear utility function $u_i(b^t) = v_i(f(b^t)) - p_i(b^t)$, where v_i is a valuation function over vectors of outcomes and $p_i = \sum_j p_{i,j}(b^t)$. In addition, in each round we assume that each mechanism is available to each bidder with a certain probability. We let the Bernoulli random variable $A_{i,j} = 1$ if mechanism j is available to bidder i . Due to availability, the mechanisms must also be applicable when only subsets of bidders are placing bids. For this reason, it will be convenient to assume that the outcome space for mechanism $j \in [m]$ is $\mathcal{X}_j = \mathcal{X}_{1,j} \times \dots \times \mathcal{X}_{n,j}$ and $x_j \in \mathcal{X}_j$ is $x_j = (x_{i,j})_{i \in [n]}$. We assume that each bidder, for whom the mechanism is not available, must place a bid of “0”. If bidder i bids 0 for mechanism j , we assume $f_j(0, b_{-i,j}) = \perp_{i,j}$, where $\perp_{i,j}$ is a “losing” outcome, and payment $p_{i,j}(0, b_{-i,j}) = 0$.

Oblivious Learning We assume *oblivious learning* – each bidder runs a single no-regret learning algorithm and uses the utility of every round as feedback, no matter how the availability in each round turns out. In hindsight, the average history of play for oblivious learning becomes an availability-oblivious variant of coarse-correlated equilibrium [17]. Hence, the outcomes of oblivious learning are captured by the coarse-correlated equilibria in the following one-shot game: First, all bidders simultaneously place a bid for every mechanism. They know only the probability distribution of the availabilities. Only after they placed their bids, the availability of each mechanism for each bidder is determined at random.

Definition 5.1. An availability-oblivious coarse-correlated equilibrium is a distribution over bid vectors b (independent of A) such that, in expectation over all availabilities, it is not beneficial for any bidder i to switch to another bid b'_i . For each i and each b'_i , we have $\mathbf{E}[u_i(b'_i, b_{-i})] \leq \mathbf{E}[u_i(b)]$.

Indeed, our results also hold for a larger class of equilibria, in which a subset of bidders might not be oblivious to availabilities. For our guarantees, it is enough to consider distributions over bidding strategies b which might be dependent on A such that, in expectation over all availabilities, it is not beneficial for any bidder i to switch to another bid b'_i . For each i and each b'_i , we have $\mathbf{E}[u_i(b'_i, b_{-i})] \leq \mathbf{E}[u_i(b)]$. Note that both ordinary coarse-correlated equilibria and availability-oblivious ones fulfill this property.

We bound the performance of these equilibria by deriving suitable smoothness bounds.

Smoothness We assume that each mechanism j satisfies *weak smoothness* as defined in Subsection 4.3.1 of Chapter 4. For any valuations $v_{i,j}: \mathcal{X}_j \rightarrow \mathbb{R}^{\geq 0}$ there are (possibly randomized) deviations¹ $b'_{i,j}$ for each $i \in [n]$ such that for all bid vectors $b_{\cdot,j}$

$$\begin{aligned} & \mathbf{E} \left[\sum_{i \in [n]} v_{i,j}(f_j(b'_{i,j}, b_{-i,j})) - p_{i,j}(b'_{i,j}, b_{-i,j}) \right] \\ & \geq \lambda \cdot \max_{x_j \in \mathcal{X}_j} \sum_{i \in [n]} v_{i,j}(x_j) - \mu_1 \cdot \sum_{i \in [n]} p_{i,j}(b_{\cdot,j}) - \mu_2 \sum_{i \in [n]} W_{i,j}(b_{i,j}, f_j(b_{\cdot,j})) \quad , \quad (5.1) \end{aligned}$$

where $W_{i,j}(b_{i,j}, x_j) = \max_{b_{-i,j}: f_j(b_{\cdot,j})=x_j} p_{i,j}(b_{\cdot,j})$. For intuition, assume that (5.1) holds with $\mu_2 = 0$. Consider a learning outcome with a no-regret guarantee where every bidder i can gain at most ε in any fixed deviation, i.e., $\mathbf{E}[v_{i,j}(f_j(b_{\cdot,j})) - p_{i,j}(b_{\cdot,j})] \geq \mathbf{E}[v_{i,j}(f_j(b'_{i,j}, b_{-i,j})) - p_{i,j}(b'_{i,j}, b_{-i,j})] - \varepsilon$. Applying (5.1) pointwise

$$\sum_{i \in [n]} \mathbf{E}[v_{i,j}(f_j(b_{\cdot,j})) - p_{i,j}(b_{\cdot,j})] \geq \lambda \cdot \max_{x_j \in \mathcal{X}_j} \sum_{i \in [n]} v_{i,j}(x_j) - \mu_1 \cdot \sum_{i \in [n]} \mathbf{E}[p_{i,j}(b_{\cdot,j})] - n\varepsilon,$$

which implies for social welfare

$$\sum_{i \in [n]} \mathbf{E}[v_{i,j}(f_j(b_{\cdot,j}))] \geq \lambda \cdot \max_{x_j \in \mathcal{X}_j} \sum_{i \in [n]} v_{i,j}(x_j) + (1 - \mu_1) \cdot \sum_{i \in [n]} \mathbf{E}[p_{i,j}(b_{\cdot,j})] - n\varepsilon .$$

Every bidder i can stay away from the market and payments are non-negative, so $0 \leq \mathbf{E}[p_{i,j}(b_{\cdot,j})] \leq \mathbf{E}[v_{i,j}(f_j(b_{\cdot,j}))] + \varepsilon$ and

$$\max(1, \mu_1) \sum_{i \in [n]} \mathbf{E}[v_{i,j}(f_j(b_{\cdot,j}))] \geq \lambda \cdot \max_{x_j \in \mathcal{X}_j} \sum_{i \in [n]} v_{i,j}(x_j) - (n + \mu_1)\varepsilon .$$

Thus, for $\varepsilon \rightarrow 0$, the price of anarchy tends to $\max(1, \mu_1)/\lambda$. More generally, (5.1) implies a bound on the price of anarchy of $(\mu_2 + \max(1, \mu_1))/\lambda$ for many equilibrium concepts. If $\mu_2 > 0$, then the bound relies on an additional no-overbidding assumption, which directly transfers to our results. For details see Subsection 4.3.1 in Chapter 4.

¹In slight contrast to [85], we here assume that the smoothness deviations of a bidder do not depend on his own current bid. This serves to simplify our exposition and can be incorporated into our analysis.

Valuation Functions Our main results in this chapter apply for the class of monotone lattice-submodular valuations (see Definition 3.4). In this thesis, we concentrate on distributive lattices, for which this definition is equivalent to the diminishing marginal returns property:

$$\forall z_i \succeq_i y_i \in \mathcal{X}_i \implies \forall t \in \mathcal{X}_i : v_i(t \vee y_i) - v(y_i) \geq v_i(t \vee z_i) - v(z_i).$$

As already mentioned in Subsection 3.1.2, lattice-submodular functions generalize submodular set functions but are a strict subclass of XOS functions. For the definition of the XOS valuation class, see Definition 3.1.

5.2 Composition with Independent Admission

We first consider simultaneous composition of smooth mechanisms with independent availabilities. Here, all random variables $A_{i,j}$ are independent, and we let $q_{i,j} = \Pr[A_{i,j} = 1]$.

Definition 5.2. *Let v be a valuation function on a product lattice, coming from a class of valuation functions \mathcal{V} . Given vectors x^1, \dots, x^k and numbers $\alpha_1, \dots, \alpha_k \in [0, 1]$ such that $\sum_{j=1}^k \alpha_j = 1$, determine another vector y at random by setting component y_i to x_i^j independently with probability α_j . Then, the smallest γ s.t. $\sum_{j=1}^k \alpha_j v(x^j) \leq \gamma \cdot \mathbf{E}[v(y)]$ is the correlation gap of class \mathcal{V} .*

Theorem 5.3. *Suppose bidder valuations are monotone and come from a class \mathcal{V} with a correlation gap of $\gamma(\mathcal{V})$. The price of anarchy for oblivious learning for simultaneous composition of weakly (λ, μ_1, μ_2) -smooth mechanisms with valuations from \mathcal{V} and fully independent availability is at most $\gamma(\mathcal{V}) \cdot (\mu_2 + \max(1, \mu_1)) / \lambda$.*

Before the proof of the main theorem of this section, we prove an upper bound on $e/(e-1)$ on the correlation gap of lattice-submodular valuations with diminishing marginal returns. This result slightly generalizes the result of [2] from composition of totally ordered sets to arbitrary product lattices.

Lemma 5.4 (Correlation Gap on a Product Lattice). *Let v be a function with diminishing marginal returns on a product lattice. Given vectors x^1, \dots, x^k and numbers $\alpha_1, \dots, \alpha_k \in [0, 1]$ such that $\sum_{j=1}^k \alpha_j = 1$, determine another vector y at random by setting component y_i to x_i^j independently with probability α_j . Then $\mathbf{E}[v(y)] \geq (1 - \frac{1}{e}) \sum_{j=1}^k \alpha_j v(x^j)$.*

Proof. Without loss of generality, let $v(x^1) \geq v(x^2) \geq \dots \geq v(x^k)$. For each component $i \in [m]$, let $J_i \in [k]$ be the random variable of the index of the vector from which y_i was taken.

Let z be defined by

$$z_i = \begin{cases} \perp_i, & \text{if } J_i = 1, \\ y_i, & \text{otherwise.} \end{cases}$$

If $J_i \neq 1$, we have

$$v(y_1, \dots, y_i, z_{i+1}, \dots, z_m) - v(y_1, \dots, y_{i-1}, z_i, \dots, z_m) = 0.$$

Otherwise, if $J_i = 1$, we have

$$\begin{aligned}
 & v(y_1, \dots, y_i, z_{i+1}, \dots, z_m) - v(y_1, \dots, y_{i-1}, z_i, \dots, z_m) \\
 & \geq v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, y_i, z_{i+1}, \dots, z_m) - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, \perp_i, z_{i+1}, \dots, z_m) \\
 & = v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, x_i^1, z_{i+1}, \dots, z_m) - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, \perp_i, z_{i+1}, \dots, z_m).
 \end{aligned}$$

That is, in combination, we get

$$\begin{aligned}
 & \mathbf{E} [v(y_1, \dots, y_i, z_{i+1}, \dots, z_m) - v(y_1, \dots, y_{i-1}, z_i, \dots, z_m)] \\
 & \geq \alpha_1 \mathbf{E} [v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, x_i^1, z_{i+1}, \dots, z_m) \\
 & \quad - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, \perp_i, z_{i+1}, \dots, z_m) \mid J_i = 1] \\
 & = \alpha_1 \mathbf{E} [v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, x_i^1, z_{i+1}, \dots, z_m) \\
 & \quad - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, \perp_i, z_{i+1}, \dots, z_m)] \quad ,
 \end{aligned}$$

where the last step uses the independence of the components.

Note that, by diminishing marginal returns, we have

$$\begin{aligned}
 & v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, x_i^1, z_{i+1}, \dots, z_m) - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, \perp_i, z_{i+1}, \dots, z_m) \\
 & \geq v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, x_i^1 \vee z_i, z_{i+1}, \dots, z_m) \\
 & \quad - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, z_i, z_{i+1}, \dots, z_m).
 \end{aligned}$$

Applying furthermore $x_i^1 \vee z_i = x_i^1 \vee y_i$ and taking the expectation, we get

$$\begin{aligned}
 & \mathbf{E} [v(y_1, \dots, y_i, z_{i+1}, \dots, z_m) - v(y_1, \dots, y_{i-1}, z_i, \dots, z_m)] \\
 & \geq \alpha_1 \mathbf{E} [v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, x_i^1 \vee y_i, z_{i+1}, \dots, z_m) \\
 & \quad - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, z_i, z_{i+1}, \dots, z_m)].
 \end{aligned}$$

Overall, we get

$$\begin{aligned}
 \mathbf{E} [v(y)] &= \mathbf{E} \left[v(z) + \sum_{i=1}^m v(y_1, \dots, y_i, z_{i+1}, \dots, z_m) - v(y_1, \dots, y_{i-1}, z_i, \dots, z_m) \right] \\
 &\geq \mathbf{E} [v(z)] + \sum_{i=1}^m \alpha_1 \mathbf{E} [v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, x_i^1 \vee y_i, z_{i+1}, \dots, z_m) \\
 & \quad - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, z_i, \dots, z_m)] \\
 &= \mathbf{E} [v(z)] + \alpha_1 (\mathbf{E} [v(x^1 \vee y)] - \mathbf{E} [v(z)]) \\
 &\geq (1 - \alpha_1) \mathbf{E} [v(z)] + \alpha_1 v(x^1) \quad .
 \end{aligned}$$

By applying this argument inductively, we also get

$$\mathbf{E} [v(y)] \geq \sum_{i=1}^k \alpha_i v(x^1) \prod_{i'=1}^{i-1} (1 - \alpha_{i'}) \quad .$$

To proceed with the proof, we use the following inequality.

Fact 5.5 (Generalized Chebyshev's Sum Inequality (see for instance [47])). *Let $a_1 \geq a_2 \geq \dots \geq a_k$ and $b_1 \geq b_2 \geq \dots \geq b_k$ be any real numbers, and m_1, \dots, m_k non-negative real numbers whose sum is 1. Then*

$$\sum_{i=1}^k a_i b_i m_i \geq \left(\sum_{i=1}^k a_i m_i \right) \left(\sum_{i=1}^k b_i m_i \right). \quad (5.2)$$

For $a_i = v(x^1)$, $b_i = \prod_{i'=1}^{i-1} (1 - \alpha_{i'})$, $m_i = \alpha_i$, $i \in \{1, \dots, k\}$, (5.2) gives us

$$\sum_{i=1}^k \alpha_i v(x^1) \prod_{i'=1}^{i-1} (1 - \alpha_{i'}) \geq \left(\sum_{i=1}^k \alpha_i \prod_{i'=1}^{i-1} (1 - \alpha_{i'}) \right) \left(\sum_{i=1}^k \alpha_i v(x^1) \right),$$

This further implies

$$\mathbf{E}[v(y)] \geq \left(\frac{1}{k} \sum_{i=1}^k \left(1 - \frac{1}{k}\right)^{i-1} \right) \left(\sum_{i=1}^k \alpha_i v(x^1) \right) \geq \left(1 - \frac{1}{e}\right) \sum_{i=1}^k \alpha_i v(x^1),$$

as the expression $\left(\sum_{i=1}^k \alpha_i \prod_{i'=1}^{i-1} (1 - \alpha_{i'}) \right)$ is minimized for $\alpha_i = \frac{1}{k}$, $i \in \{1, \dots, k\}$. Since $v(x^1) \geq v(x^i)$, $i \in \{1, \dots, k\}$, the claim is proved. \square

From here, we arrive at the following corollary of the main theorem.

Corollary 5.6. *The price of anarchy for oblivious learning for simultaneous composition of weakly (λ, μ_1, μ_2) -smooth mechanisms with monotone lattice-submodular valuations and fully independent availability is at most $e/(e-1) \cdot (\mu_2 + \max(1, \mu_1))/\lambda$.*

We proceed to prove Theorem 5.3.

Proof Theorem 5.3. We will prove the theorem by defining an availability-oblivious (randomized) deviation b'_i for each player i such that the following inequality will hold for any (not necessarily availability-oblivious) bidding strategy b :

$$\begin{aligned} & \sum_i \mathbf{E}[u_i(b'_i, b_{-i})] \\ & \geq \frac{1}{\gamma(\mathcal{V})} \cdot \lambda \cdot \sum_i \mathbf{E}[v_i(x^*)] - \mu_1 \sum_i \mathbf{E}[p_i(b)] - \mu_2 \sum_i \mathbf{E}[W_i(b_i, f(b))] \quad , \end{aligned} \quad (5.3)$$

where x^* denotes the (random) optimal outcome. From this inequality, whose form is in fact exactly that of the smoothness condition (5.1), the claim of the theorem follows as described in Section 4.1.

In more detail, to attain the aforementioned inequality, we will relate each player's utility for deviating to b'_i to the utility he could achieve if he was allowed to see and react upon the availabilities. In that case, he could simply use the smoothness deviation tailored to the specific availability profile $A_i = (A_{i,1}, \dots, A_{i,m})$ that he is encountering. We denote this non-oblivious smoothness deviation by $b_i^{A_i}$. Because the global mechanism is a simultaneous composition of (λ, μ_1, μ_2) -smooth mechanisms, it is again (λ, μ_1, μ_2) -smooth. Therefore we know that the non-oblivious deviations $b_i^{A_i}$ do exist, and they satisfy the smoothness inequality (5.1) by definition.

We proceed to define, for each player i , the availability-oblivious deviation b'_i . First, bidder i assumes for himself a reduced valuation function $\bar{v}_i = \alpha \cdot v_i$, for some appropriate α to be chosen later. The deviation b'_i is a composition of component-wise independent deviations $b'_{i,j}$, i.e., $b'_i = (b'_{i,1}, \dots, b'_{i,m})$ where each $b'_{i,j}$ is chosen independently. To arrive at $b'_{i,j}$, bidder i assumes that mechanism j is available to him and draws all other availabilities independently according to probabilities $q_{i',j'}$. This means that he draws availabilities for all other players on all mechanisms and also his own availabilities on all mechanisms other than j . Now he has a full availability profile, and therefore he can consider the non-oblivious smoothness deviation. He observes the j -th component of this smoothness deviation and sets $b'_{i,j}$ to be equal to it. Note that $b'_{i,j}$ will be applied only with the probability that mechanism j is in fact available to bidder i , i.e., with probability $q_{i,j}$.

Next, we want to compare $u_i(b'_i, b_{-i})$ and $u_i(b_i^{A_i}, b_{-i})$. Let us focus on the valuation $v_i(f(b'_i, b_{-i}))$ first. The non-oblivious smoothness deviation $b_i^{A_i}$ is a vector whose components are correlated. More precisely, to form this bid we observe A_i , sample the availabilities A_{-i} and bids b_{-i} of other players, and take the optimal allocation x^* for the resulting availability profile A . Then, we determine the ℓ for which $\bar{v}_i(x_i^*) = \sum_j \bar{v}_{i,j}^\ell(x_{i,j}^*)$ and use $\bar{v}_{i,j}^\ell$ for determining $b_{i,j}^{A_i}$ (note that A_i can be regarded as bidder i 's type in a Bayesian sense, for more details see [85]). Therefore, the components of $b_i^{A_i}$ are correlated through the common choice of ℓ . Our deviation b'_i is assembled by setting $b'_{i,j} = (b_{i,j}^{A_i})_{k_j}$ independently for each j .

Formally, let $r_{i,j}^\ell$ denote the conditional probability that the optimum yields an outcome vector x^* that attains its maximum value for bidder i in \bar{v}_i^ℓ , given that $A_{i,j} = 1$. Then, the marginal probability of observing $b_{i,j}^{A_i} = (b_{i,j}^{A_i})_\ell$ is $r_{i,j}^\ell q_{i,j}$. In b'_i we pick ℓ independently for each mechanism with probability $r_{i,j}^\ell$, which yields a combined probability of $r_{i,j}^\ell q_{i,j}$ for availability and deviation. Thus, b'_i simulates the marginal probabilities of outcomes in $b_i^{A_i}$, i.e., $\Pr[f_j(b'_i, b_{-i}) = y_{i,j} \mid A_{-i}, b_{-i}] = \Pr[f_j(b_i^{A_i}, b_{-i}) = y_{i,j} \mid A_{-i}, b_{-i}]$ for all $y_{i,j} \in \mathcal{X}_{i,j}$, for each $j \in [m]$. Hence, for fixed A_{-i}, b_{-i} , the two expected valuations $\mathbf{E}[v_i(f(b'_i, b_{-i})) \mid A_{-i}, b_{-i}]$ and $\mathbf{E}[v_i(f(b_i^{A_i}, b_{-i})) \mid A_{-i}, b_{-i}]$ are related via correlation gap.

Thus, setting $\alpha = 1/\gamma(\mathcal{V})$ and $\bar{v}_i(x) = 1/\gamma(\mathcal{V}) \cdot v_i(x)$ we get

$$\begin{aligned} \mathbf{E}[v_i(f(b'_i, b_{-i})) \mid A_{-i}, b_{-i}] &= \sum_{y \in \mathcal{X}} v_i(y) \cdot \Pr[f(b'_i, b_{-i}) = y \mid A_{-i}, b_{-i}] \\ &= \sum_{y \in \mathcal{X}} v_i(y) \cdot \prod_j \Pr[f_j(b'_i, b_{-i}) = y_{i,j} \mid A_{-i}, b_{-i}] \\ &\geq \frac{1}{\gamma(\mathcal{V})} \cdot \sum_{y \in \mathcal{X}} v_i(y) \cdot \Pr[f(b_i^{A_i}, b_{-i}) = y \mid A_{-i}, b_{-i}] \\ &= \frac{1}{\gamma(\mathcal{V})} \cdot \mathbf{E}[v_i(f(b_i^{A_i}, b_{-i})) \mid A_{-i}, b_{-i}] \\ &= \mathbf{E}[\bar{v}_i(f(b_i^{A_i}, b_{-i})) \mid A_{-i}, b_{-i}] . \end{aligned}$$

In addition, because payments are simply additive across mechanisms, it is straightfor-

ward to see that for every bidder i

$$\mathbf{E} [p_i(b'_i, b_{-i}) \mid A_{-i}, b_{-i}] = \mathbf{E} [p_i(b_i^{A_i}, b_{-i}) \mid A_{-i}, b_{-i}].$$

This allows to apply the smoothness bound for Bayesian mechanisms with independent types from Subsection 4.3.1 in Chapter 4 to derive

$$\begin{aligned} & \sum_i \mathbf{E} [u_i(b'_i, b_{-i})] \\ &= \sum_i \mathbf{E} [v_i(f(b'_i, b_{-i}))] - \mathbf{E} [p_i(b'_i, b_{-i})] \\ &\geq \sum_i \mathbf{E} [\bar{v}_i(f(b_i^{A_i}, b_{-i}))] - \mathbf{E} [p_i(b_i^{A_i}, b_{-i})] \\ &\geq \lambda \cdot \sum_i \mathbf{E} [\bar{v}_i(x^*)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [W_i(b_i, f(b))] \\ &= \frac{\lambda}{\gamma(\mathcal{V})} \cdot \sum_i \mathbf{E} [v_i(x^*)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [W_i(b_i, f(b))] \end{aligned}$$

This proves the desired smoothness guarantee and implies the theorem. \square

5.3 Composition with Everybody-or-Nobody Admission

We consider the case in which at each point in time each mechanism is either available to all bidders or to none. We let $A_j = A_{i,j}$ for all $i \in [n]$ and $q_j = \Pr [A_j = 1]$. Note that all A_j are assumed to be independent.

Let the social optimum be denoted by x^* . We assume that $x_j^* = \perp_j$ if $A_j = 0$. Otherwise, x^* might have different values, depending on the availabilities of other mechanisms. Let us denote the possible outcomes by x_j^1, x_j^2, \dots and let $r_j^\ell := \Pr [x_j^* = x_j^\ell \mid A_j = 1]$. That is, r_j^ℓ is the marginal probability of x_j^ℓ conditioned on j being available. Theorem 5.7 formulates our main result in this section.

Theorem 5.7. *The price of anarchy for oblivious learning for simultaneous composition of weakly (λ, μ_1, μ_2) -smooth mechanisms with monotone lattice-submodular valuations and everybody-or-nobody admission is at most $4e/(e-1) \cdot (\mu_2 + \max(1, \mu_1))/\lambda^2$.*

Proof. We will prove that, for each bidder i and each mechanism j there are randomized deviation strategies $b'_{i,j}$ that are independent of the availabilities such that the following smoothness guarantee holds against any (potentially non-oblivious) bidding strategy b :

$$\begin{aligned} & \sum_i \mathbf{E} [u_i(b'_i, b_{-i})] \\ &\geq \left(1 - \frac{1}{e}\right) \frac{\lambda^2}{4} \sum_i \mathbf{E} [v_i(x^*)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [W_i(b_i, f(b))] . \end{aligned}$$

From this guarantee the claim of the theorem again follows as described in Section 4.1.

To define $b'_{i,j}$, every bidder i draws two vectors z^i and \tilde{t}^i at random as follows. He sets z_j^i to x_j^ℓ with probability r_j^ℓ/α , where $\alpha = 2/\lambda$, and to \perp_j with the remaining

probability. Furthermore, he sets \tilde{t}_j^i to x_j^ℓ with probability $q_j r_j^\ell$ and to \perp_j with the remaining probability. These draws are performed independent of any availabilities. Observe that for each i , we have $\mathbf{E} [\sum_{i'} v_{i'}(\tilde{t}^i)] \geq (1 - \frac{1}{e}) \mathbf{E} [\sum_{i'} v_{i'}(x^*)]$ by Lemma 5.4.

Due to the random draws, each bidder i' defines functions $w_{i',j}^{i'}: \Omega_j \rightarrow \mathbb{R}$ for each bidder i and each mechanism j . Function $w_{i',j}^{i'}$ maps an outcome of mechanism j , denoted by y_j , to a real number as follows

$$w_{i',j}^{i'}(y_j) = v_i(\tilde{t}_1^{i'}, \dots, \tilde{t}_{j-1}^{i'}, y_j \wedge z_j^{i'}, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^{i'}, \dots, \tilde{t}_{j-1}^{i'}, \perp_j, \dots, \perp_m) .$$

Note that these functions do not necessarily reflect the actual value any outcome might have. They are only used to define the deviation strategy: bidder i' pretends all bidders i , including himself, would have valuations $w_{i',j}^{i'}$ for the outcome of mechanism j . This gives him a deviation strategy $b'_{i',j}$ by setting $b'_{i',j} = b_{i',j}^*(w_{1,j}^{i'}, \dots, w_{n,j}^{i'})$ as defined by the smoothness of mechanism j .

The proofs for the following three lemmas are, for ease of exposition, presented after the proof of the theorem.

Lemma 5.8. *For every bidder i and deviating bids $b'_{i,j} = b_{i,j}^*(w_{1,j}^i, \dots, w_{n,j}^i)$,*

$$\mathbf{E} [v_i(f(b'_i, b_{-i}))] \geq \sum_j \mathbf{E} [w_{i,j}^i(f_j(b'_{i,j}, b_{-i}))] - \frac{1}{\alpha(\alpha + 1)} \mathbf{E} [v_i(\tilde{t}^i)] .$$

Lemma 5.9. *For the adjusted functions w we can apply smoothness to obtain*

$$\begin{aligned} & \sum_i \sum_j \mathbf{E} [w_{i,j}^i(f_j(b'_{i,j}, b_{-i})) - p_{i,j}(b'_{i,j}, b_{-i})] \\ & \geq \lambda \sum_i \sum_j q_j \mathbf{E} [w_{i,j}^1(z_j^1)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [W_i(b_i, f(b))] . \end{aligned}$$

Lemma 5.10. *For function w^1 , random vectors z_j^1 and \tilde{t}^1 , and every mechanism j*

$$\sum_j q_j \mathbf{E} [w_{i,j}^1(z_j^1)] = \frac{1}{\alpha} \mathbf{E} [v_i(\tilde{t}^1)] .$$

The bound from Lemma 5.9 has striking similarities to the smoothness bound (5.1). However, it is expressed in terms of the functions $w_{i',j}^{i'}$ rather than the actual valuation functions v_i . The other two Lemmas show that, in expectation, these functions are close enough to the functions v_i so that this bound actually suffices to prove the main result:

$$\begin{aligned}
 & \sum_i \mathbf{E} [u_i(b'_i, b_{-i})] \\
 &= \sum_i \mathbf{E} \left[v_i(f(b'_i, b_{-i})) - \sum_j p_{i,j}(b'_{i,j}, b_{-i}) \right] \\
 &\geq \sum_i \sum_j \mathbf{E} [w_{i,j}^i(f_j(b'_{i,j}, b_{-i})) - p_{i,j}(b'_{i,j}, b_{-i})] - \frac{1}{\alpha(\alpha+1)} \sum_i \mathbf{E} [v_i(\tilde{t}^i)] \\
 &\hspace{20em} \text{(by Lemma 5.8)} \\
 &\geq \lambda \sum_i \sum_j q_j \mathbf{E} [w_{i,j}^1(z_j^1)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [W_i(b_i, f(b))] \\
 &\hspace{15em} - \frac{1}{\alpha(\alpha+1)} \sum_i \mathbf{E} [v_i(\tilde{t}^1)] \quad \text{(by Lemma 5.9)} \\
 &= \sum_i \left(\frac{\lambda}{\alpha} - \frac{1}{\alpha(\alpha+1)} \right) \mathbf{E} [v_i(\tilde{t}^i)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [W_i(b_i, f(b))] \ . \\
 &\hspace{20em} \text{(by Lemma 5.10)}
 \end{aligned}$$

By setting $\alpha = \frac{2}{\lambda}$

$$\begin{aligned}
 \sum_i \mathbf{E} [u_i(b'_i, b_{-i})] &\geq \frac{\lambda^2}{4} \sum_i \mathbf{E} [v_i(\tilde{t}^1)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [W_i(b_i, f(b))] \\
 &\geq \left(1 - \frac{1}{e}\right) \frac{\lambda^2}{4} \sum_i \mathbf{E} [v_i(x^*)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [W_i(b_i, f(b))] \ .
 \end{aligned}$$

The last step follows from Lemma 5.4. This proves Theorem 5.7. \square

Note that technically the mechanism could be randomized itself. Our results extend to this case in a straightforward way.

5.3.1 Proof of Lemma 5.8

Let $y_j^i = f_j(b'_{i,j}, b_{-i}) \wedge z_j^i$ and $\hat{z}_j^i = z_j^i$ if j is available and $\hat{z}_j^i = \perp_j$ otherwise. Notice that $f_j(b'_{i,j}, b_{-i}) \wedge \hat{z}_j^i = f_j(b'_{i,j}, b_{-i}) \wedge z_j^i$. This is because $\hat{z}_j^i = z_j^i$ when j is available and $f_j(b'_{i,j}, b_{-i}) = \perp_j$ when j is not available.

By monotonicity, we have $v_i(f(b'_i, b_{-i})) \geq v_i(f(b'_i, b_{-i}) \wedge z^i) = v_i(y^i)$. Furthermore, we can decompose $v_i(y^i)$ into a telescoping sum by

$$v_i(y^i) = \sum_j v_i(y_1^i, \dots, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(y_1^i, \dots, y_{j-1}^i, \perp_j, \dots, \perp_m) \ .$$

Next, we bound each of these terms independently using diminishing marginal returns

multiple times

$$\begin{aligned}
 & v_i(y_1^i, \dots, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(y_1^i, \dots, y_{j-1}^i, \perp_j, \dots, \perp_m) \\
 & \geq v_i(\hat{z}_1^i, \dots, \hat{z}_{j-1}^i, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i, \dots, \hat{z}_{j-1}^i, \perp_j, \dots, \perp_m) \\
 & \geq v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m) \\
 & = v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m) \\
 & \quad - (v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) \\
 & \quad \quad \quad - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, y_j^i, \perp_{j+1}, \dots, \perp_m)) \\
 & \geq v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m) \\
 & \quad - (v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, y_j^i, \perp_{j+1}, \dots, \perp_m)) \\
 & = v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m) \\
 & \quad - (v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)) \\
 & \quad + v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m) .
 \end{aligned}$$

That is, by linearity of expectation, we have

$$\begin{aligned}
 & \mathbf{E} [v_i(y_1^i, \dots, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(y_1^i, \dots, y_{j-1}^i, \perp_j, \dots, \perp_m)] \geq \\
 & \underbrace{\mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)]}_{\text{part 1}} \\
 & - \underbrace{\mathbf{E} [v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)]}_{\text{part 2}} \\
 & + \underbrace{\mathbf{E} [v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)]}_{\text{part 3}} .
 \end{aligned}$$

To simplify part 2, we use the fact that $\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i$ and \hat{z}_j^i are independent. Therefore, we have

$$\begin{aligned}
 & \mathbf{E} [v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
 & = \sum_{\ell} \frac{q_j r_j^\ell}{\alpha} \mathbf{E} [v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, x_j^\ell, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
 & = \frac{1}{\alpha} \mathbf{E} [v_i(\tilde{t}_1^i, \dots, \tilde{t}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] .
 \end{aligned}$$

For the same reason, we can also bound part 1 by using

$$\begin{aligned}
 & \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_j^i \vee \tilde{t}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
 &= \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i \vee \tilde{t}_j^i, \perp_{j+1}, \dots, \perp_m) \\
 &\quad - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m)] \\
 &\quad + \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) \\
 &\quad - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
 &\leq \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \tilde{t}_j^i, \perp_{j+1}, \dots, \perp_m) \\
 &\quad - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
 &\quad + \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) \\
 &\quad - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
 &= (\alpha + 1) \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) \\
 &\quad - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)],
 \end{aligned}$$

which implies

$$\begin{aligned}
 & \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
 &\geq \frac{1}{\alpha + 1} \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_j^i \vee \tilde{t}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)]
 \end{aligned}$$

Finally, part 3 is precisely the definition of $\mathbf{E} [w_{i,j}^i(y_j^i)]$. Therefore, in combination, we get

$$\begin{aligned}
 & \mathbf{E} [v_i(y_1^i, \dots, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(y_1^i, \dots, y_{j-1}^i, \perp_j, \dots, \perp_m)] \\
 &\geq \frac{1}{\alpha + 1} \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_j^i \vee \tilde{t}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
 &\quad - \frac{1}{\alpha} \mathbf{E} [v_i(\tilde{t}_1^i, \dots, \tilde{t}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
 &\quad + \mathbf{E} [w_{i,j}^i(y_j^i)]
 \end{aligned}$$

Taking the sum over all j , we get two telescoping sums, which simplify to $\mathbf{E} [v_i(\hat{z}^i \vee \tilde{t}^i)]$ (part 1) and $\mathbf{E} [v_i(\tilde{t}^i)]$ (part 2). This gives us

$$\begin{aligned}
 \mathbf{E} [v_i(f(b'_i, b_{-i}))] &\geq \mathbf{E} [v_i(y^i)] \\
 &\geq \frac{1}{\alpha + 1} \mathbf{E} [v_i(\hat{z}^i \vee \tilde{t}^i)] - \frac{1}{\alpha} \mathbf{E} [v_i(\tilde{t}^i)] + \sum_j \mathbf{E} [w_{i,j}^i(y_j^i)] \\
 &\geq \sum_j \mathbf{E} [w_{i,j}^i(y_j^i)] - \frac{1}{\alpha(\alpha + 1)} \mathbf{E} [v_i(\tilde{t}^i)] \\
 &= \sum_j \mathbf{E} [w_{i,j}^i(f(b'_i, b_{-i}))] - \frac{1}{\alpha(\alpha + 1)} \mathbf{E} [v_i(\tilde{t}^i)] \quad . \quad \square
 \end{aligned}$$

5.3.2 Proof of Lemma 5.9

Note that functions $w_{i,j}^{i'}$ are identically distributed for different i' and independent of any availabilities. Therefore, we have

$$\mathbf{E} \left[w_{i,j}^{i'}(f_j(b'_{i,j}, b_{-i}) \mid A_j = 1) \right] = \mathbf{E} \left[w_{i,j}^1(f_j(b_{i,j}^*(w_{1,j}^1, \dots, w_{n,j}^1), b_{-i})) \mid A_j = 1 \right] \quad (5.4)$$

and

$$\mathbf{E} [p_{i,j}(b'_{i,j}, b_{-i}) \mid A_j = 1] = \mathbf{E} [p_{i,j}(b_{i,j}^*(w_{1,j}^1, \dots, w_{n,j}^1), b_{-i}) \mid A_j = 1] . \quad (5.5)$$

Next, we apply the smoothness of each separate mechanism. Let us first assume mechanism j is available and z^1 and \tilde{t}^1 are fixed arbitrarily. This also fixes the functions $w_{1,j}^1, \dots, w_{n,j}^1$. We pretend these are the actual valuation functions. Then smoothness gives us

$$\begin{aligned} & \sum_i w_{i,j}^1(f_j(b_{i,j}^*(w_{1,j}^1, \dots, w_{n,j}^1), b_{-i})) - p_{i,j}(b_{i,j}^*(w_{1,j}^1, \dots, w_{n,j}^1), b_{-i}) \\ & \geq \lambda \left(\max_{y \in \Omega_j} \sum_i w_{i,j}^1(y) \right) - \mu_1 \sum_i p_{i,j}(b) - \mu_2 \sum_i h_{i,j}(b_i, f(b)) \\ & \geq \lambda \left(\sum_i w_{i,j}^1(z_j^1) \right) - \mu_1 \sum_i p_{i,j}(b) - \mu_2 \sum_i h_{i,j}(b_i, f(b)) . \end{aligned}$$

Taking the expectation over z^1 and \tilde{t}^1 , we can combine this bound with (5.4) and (5.5) to get

$$\begin{aligned} & \sum_i \mathbf{E} [w_{i,j}^i(f_j(b'_{i,j}, b_{-i})) - p_{i,j}(b'_{i,j}, b_{-i}) \mid A_j = 1] \\ & = \sum_i \mathbf{E} [w_{i,j}^1(b_{i,j}^*(w_{1,j}^1, \dots, w_{n,j}^1), b_{-i})) - p_{i,j}(b_{i,j}^*(w_{1,j}^1, \dots, w_{n,j}^1), b_{-i}) \mid A_j = 1] \\ & \geq \mathbf{E} \left[\lambda \sum_i w_{i,j}^1(z_j^1) - \mu_1 \sum_i p_{i,j}(b) - \mu_2 \sum_i h_{i,j}(b_i, f(b)) \mid A_j = 1 \right] . \end{aligned}$$

If j is not available, then $f_j(b'_{i,j}, b_{-i}) = \perp_j$ and $p_{i,j}(b'_{i,j}, b_{-i}) = p_{i,j}(b) = 0$. By definition, however, $w_{i,j}^1(z_j^1)$ is independent of the fact whether j is available or not. Therefore, we have

$$\begin{aligned} & \sum_i \mathbf{E} [w_{i,j}^i(f_j(b'_{i,j}, b_{-i})) - p_{i,j}(b'_{i,j}, b_{-i})] \\ & \geq q_j \mathbf{E} \left[\lambda \sum_i w_{i,j}^1(z_j^1) \mid A_j = 1 \right] - q_j \mathbf{E} \left[\mu_1 \sum_i p_{i,j}(b) + \mu_2 \sum_i h_{i,j}(b_i, f(b)) \mid A_j = 1 \right] \\ & = q_j \mathbf{E} \left[\lambda \sum_i w_{i,j}^1(z_j^1) \right] - \mathbf{E} \left[\mu_1 \sum_i p_{i,j}(b) \right] - \mathbf{E} \left[\mu_2 \sum_i h_{i,j}(b_i, f(b)) \right] . \end{aligned}$$

We can take the sum over all j to get

$$\begin{aligned} & \sum_i \sum_j \mathbf{E} [w_{i,j}^i(f_j(b'_{i,j}, b_{-i})) - p_{i,j}(b'_{i,j}, b_{-i})] \\ & \geq \lambda \sum_i \sum_j q_j \mathbf{E} [w_{i,j}^1(z_j^1)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [W_i(b_i, f(b))] . \quad \square \end{aligned} \quad (5.6)$$

5.3.3 Proof of Lemma 5.10

By simply plugging in the definitions of $w_{i,j}^1$, z_j^1 , and \tilde{t}^1 , we get

$$\begin{aligned} & \sum_j q_j \mathbf{E} [w_{i,j}^1(z_j^1)] \\ & = \sum_j q_j \mathbf{E} [v_i(\tilde{t}_1^1, \dots, \tilde{t}_{j-1}^1, z_j^1, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^1, \dots, \tilde{t}_{j-1}^1, \perp_j, \dots, \perp_m)] \\ & = \sum_j q_j \sum_{\ell} \frac{r_{\ell}}{\alpha} \mathbf{E} [v_i(\tilde{t}_1^1, \dots, \tilde{t}_{j-1}^1, x_j^{\ell}, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^1, \dots, \tilde{t}_{j-1}^1, \perp_j, \dots, \perp_m)] \\ & = \sum_j \frac{1}{\alpha} \sum_{\ell} q_j r_{\ell} \mathbf{E} [v_i(\tilde{t}_1^1, \dots, \tilde{t}_{j-1}^1, x_j^{\ell}, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^1, \dots, \tilde{t}_{j-1}^1, \perp_j, \dots, \perp_m)] \\ & = \sum_j \frac{1}{\alpha} \mathbf{E} [v_i(\tilde{t}_1^1, \dots, \tilde{t}_j^1, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^1, \dots, \tilde{t}_{j-1}^1, \perp_j, \dots, \perp_m)] \\ & = \frac{1}{\alpha} \mathbf{E} [v_i(\tilde{t}^1)] . \quad \square \end{aligned}$$

5.4 A Lower Bound for General XOS Functions

In this section we consider combinatorial auctions with item bidding and first-price auctions. We can apply the previous analysis, since for each bidder the outcomes form a trivial 2-element lattice – winning an item is the supremum outcome, not winning is the infimum outcome. In the analysis, observe that each bidder determines a random allocation of items according to the probabilities in the optimum. Based on these allocations, bidders determine the valuations $w_{i,j}^{i'}$, which in turn form the basis for the deviation. The first-price auction with general bidding space is $(1 - 1/e, 1, 0)$ -smooth [85]. If valuation functions are submodular, the composition theorems can be applied to yield the following corollary.

Corollary 5.11. *The price of anarchy for oblivious learning for simultaneous composition of single-item first-price auctions with monotone submodular valuations and fully independent availability is at most $1/(1 - 1/e)^2$; for everybody-or-nobody admission it is at most $4/(1 - 1/e)^3$.*

For more general XOS valuations, we prove a lower bound that with oblivious bidding we will not be able to show a guarantee based on the smoothness parameters – even for a single bidder, so the bound applies without assumptions on correlation among bidders.

Theorem 5.12. *In a simultaneous composition of discrete first-price single-item auctions with m items and XOS valuations, the price of anarchy for pure Nash equilibria with oblivious bidding can be as large as $\Omega((\log m)/(\log \log m))$, while each single mechanism is weakly $(1/2, 1, 0)$ -smooth.*

Proof. Consider the following single-item first-price auction with a discrete bidding space. Each bidder has valuation 0 or 2 for the item, and the set of possible bids is $\{0, 1, 2\}$. The item is sold to one of the bidders with the maximal bid (arbitrary but fixed deterministic tie-breaking), and only this bidder pays his bid. If all bidders bid 0, the item is not given away. It is easy to see that if each bidder i in this auction deviates to half of his valuation, the auction becomes smooth with $\lambda = 1/2$, $\mu_1 = 1$ and $\mu_2 = 0$. Hence, the auction has a price of anarchy of at most 2.

We compose this auction for a set $[m]$ of $m = k^2$ items, for some integer $k > 0$, and every item is sold simultaneously via the first-price auction above. There is a single bidder, and he has an XOS valuation function v as follows. The items are grouped into k groups M_1, \dots, M_k of k items each. For a set of items $S \subseteq [m]$ we have $v(S) = \max_{\ell=1, \dots, k} \sum_{j=1}^m v_j^\ell$ with $v_j^\ell = 2$ if $j \in M_\ell$ and 0 otherwise, for $\ell = 1, \dots, k$. Consequently, $v(S) = 2 \max_{\ell=1, \dots, k} |S \cap M_\ell|$. We assume each item $j \in [m]$ is available independently with probability $q_j = 1/k$.

If the bidder can deviate depending on the set of available items, a social optimum b^* is obvious – he considers the group ℓ^* with the maximum number of available items and bids $b_j^* = 1$ for all $j \in M_{\ell^*}$ and 0 otherwise. This way he always obtains a set S of items that maximizes the valuation. Furthermore, this is also the best-response since he obtains the maximum valuation at minimum required payment, and the marginal utility of every obtained item is 1. In contrast, we show that every *oblivious* deterministic best-response bid b allows to recover at most a small fraction of the above described optimum. Thus, even the price of anarchy for pure equilibria cannot be bounded by the smoothness guarantee.

In the optimum b^* , the bidder gets all available items from the group with the maximum number of available items. The number of available items in a group follows a binomial distribution $B(k, 1/k)$. This scenario is almost identical to throwing k balls uniformly at random into k bins and recording the maximum number of balls in any bin. Now in each bin k balls appear independently at random with probability $1/k$ each, and an almost identical analysis implies

$$\mathbf{E}[v(S(b^*))] = \Theta\left(\frac{\log k}{\log \log k}\right).$$

Now consider an oblivious deterministic best-response b . The valuation function v treats all items of a group in a symmetric way and all groups in a symmetric way. Let us denote by r_ℓ the number of items $j \in M_\ell$ with $b_j = 1$. For a fixed vector r , expected value and payments are the same no matter on which particular items $j \in M_\ell$ a bid $b_j = 1$ is placed. For any two groups M_ℓ and $M_{\ell'}$, the expected valuation and payments remain the same if we change the bid to have $b_j = 1$ for r_ℓ items in $M_{\ell'}$ and $r_{\ell'}$ items in M_ℓ . Moreover, the expected payment depends only on $\sum_\ell r_\ell$. Now suppose there are two groups M_ℓ and $M_{\ell'}$ such that $r_\ell, r_{\ell'} \leq k/2$. This bidding strategy is obviously dominated by any bid that bids 1 on $r_\ell + r_{\ell'}$ items in M_ℓ and none in $M_{\ell'}$. In conclusion,

w.l.o.g. we can assume that $r_1 \geq r_2 \geq r_3 \geq \dots \geq r_k$ and there is k' such that $r_\ell \geq k/2$ for $\ell = 1, \dots, k' - 1$, $r_{k'-1} \geq r_{k'} \geq 0$ and $r_\ell = 0$ for $\ell = k' + 1, \dots, k$.

We show that every oblivious best-response b has $\mathbf{E}[v(S(b))] = O(1)$. Let $p(b)$ denote the total payments, X_j denote the event that item j is available, and $Y_\ell = \sum_{j \in M_\ell} X_j$ the number of available items in group M_ℓ , for all $\ell = 1, \dots, k$. Note that

$$\begin{aligned} \mathbf{E}[v(S(b)) - p(b)] &= \mathbf{E} \left[\max_{\ell=1, \dots, k'} \left(\sum_{j \in M_\ell, b_j=1} 2X_j \right) - \sum_{j \in [m], b_j=1} X_j \right] \\ &\leq 2\mathbf{E} \left[\max_{\ell=1, \dots, k'} Y_\ell \right] - \frac{k' - 1}{2} . \end{aligned}$$

Further, for any $d = 1, \dots, k$ we can use Chernoff bounds to see

$$\Pr \left[\max_{\ell=1, \dots, k'} Y_\ell \geq d \right] = 1 - (1 - \Pr[Y_1 \geq d])^{k'} \leq 1 - (1 - e^{-d-1}/d^d)^{k'} \leq \min\{1, k' e^{d-1}/d^d\} .$$

Hence,

$$\mathbf{E} \left[\max_{\ell=1, \dots, k'} Y_\ell \right] = \sum_{d=1}^k \Pr \left[\max_{\ell=1, \dots, k'} Y_\ell \geq d \right] \leq \sum_{d=1}^k \min\{1, (k'/e) \cdot (e/d)^d\} \leq \frac{3 \log k'}{\log \log k'} + \frac{1}{ek'} .$$

Thus, $\mathbf{E}[v(S(b)) - p(b)] < (6 \log k')/(\log \log k') + 6/(ek') - (k' - 1)/2$, which is positive only for $k' \leq 34$. Every bid b with $k' \geq 35$ is dominated by b' with $b'_j = 0$ for all $j \in [m]$. Hence, for a best-response b we have $\mathbf{E}[v(S(b))] < 17$. This proves the theorem. \square

5.5 Applications beyond Auctions

Our results have interesting implications beyond mechanisms that incorporate standard auction formats. A very intriguing one is channel allocation in wireless networks. The overall problem is to maximize the utilization of a wireless channel while avoiding interference. To this end, the following game was defined in [4]: Each player i corresponds to a pair of a sender s_i and a receiver r_i located in a metric space. The transmission from s_i to r_i is successful if the signal-to-interference-plus-noise ratio (SINR) is high enough. This means that the incoming interference from senders transmitting simultaneously plus ambient noise is by a factor smaller than the intended signal. Formally, transmission i is successful if

$$\frac{\frac{p}{d(s_i, r_i)^\alpha}}{\sum_{j \in S \setminus \{i\}} \frac{p}{d(s_j, r_i)^\alpha} + \nu} \geq \beta ,$$

where $d(s_k, r_l)$ denotes the distance between sender s_k and receiver r_l in the metric space, $p > 0$ is the (fixed) power level, $S \subseteq [n]$ is the set of simultaneous transmissions; $\alpha > 0$, $\beta > 0$, and $\nu \geq 0$ are constants.

To derive a game, each player has two strategies b_i : either he decides to transmit or not to. The best possible outcome is a successful transmission. An unsuccessful transmission is the worst possible outcome. Due to the energy consumption, it is considered to be

even worse than not transmitting at all. This is reflected in the following utility function.

$$u_i(b) = \begin{cases} 1, & \text{if } b_i = 1 \text{ and } i \text{ is successful against } b_{-i}, \\ -1, & \text{if } b_i = 1 \text{ and } i \text{ is not successful against } b_{-i}, \\ 0, & \text{if } b_i = 0 . \end{cases}$$

The price of anarchy of this game is constant [9]. In every coarse correlated equilibrium, the expected number of successful transmissions is only a constant smaller than the maximum possible number of simultaneous successful transmissions.

Quite surprisingly, this game corresponds to a smooth mechanism as follows. Each player decides whether to transmit or not; a player always has valuation 2 for making a successful transmission. However, whenever making a transmission (successful or not), the bidder has to pay 1. This is comparable to an all-pay auction, where each bidder has to pay his bid, regardless of whether he wins the respective item.

Theorem 5.13. *The mechanism representing the channel-allocation game is weakly $(1, \mu_1, \mu_2)$ -smooth for $\mu_1 = O(1)$ and $\mu_2 = 0$.*

Proof. Let $S \subseteq N$ be a maximum set of players that can transmit simultaneously. Define b' by setting $b'_i = 1$ for $i \in S$ and $b'_i = 0$ for $i \notin S$. That is, $u_i(b'_i, b_{-i}) = 0$ for all $i \notin S$. Consider some bid vector b , let T be the set of players making a transmission attempt. Note that by our definition $\sum_i p_i(b) = |T|$.

Furthermore, $i \in S$ is successful under (b'_i, b_{-i}) if and only if

$$\frac{\frac{p}{d(s_i, r_i)^\alpha}}{\sum_{j \in T \setminus \{i\}} \frac{p}{d(s_j, r_j)^\alpha} + \nu} \geq \beta ,$$

for which it is sufficient to have

$$\sum_{j \in T \setminus \{i\}} \frac{d(s_i, r_i)^\alpha}{d(s_j, r_j)^\alpha} + \frac{d(s_i, r_i)^\alpha}{p} \nu < \frac{1}{\beta} ,$$

which is equivalent to

$$\sum_{j \in T \setminus \{i\}} a_{j,i} < 1 \quad \text{where } a_{j,i} = \min \left\{ 1, \frac{1}{\frac{1}{\beta} - \frac{d(s_i, r_i)^\alpha}{p} \nu} \frac{d(s_i, r_i)^\alpha}{d(s_j, r_j)^\alpha} \right\} .$$

This implies $u_i(b'_i, b_{-i}) \geq 1 - 2 \sum_{j \in T \setminus \{i\}} a_{j,i}$. Taking the sum over all $i \in S$, we get

$$\sum_{i \in S} u_i(b'_i, b_{-i}) \geq |S| - 2 \sum_{j \in T} \sum_{i \in S \setminus \{j\}} a_{j,i} .$$

Lemma 11 in [9] shows that $\sum_{i \in S \setminus \{j\}} a_{j,i} \leq C$ for some constant C because $S \setminus \{j\}$ is a feasible set. This gives us

$$\sum_{i \in N} u_i(b'_i, b_{-i}) = \sum_{i \in S} u_i(b'_i, b_{-i}) \geq |S| - 2 \sum_{j \in T} C = |S| - 2C \sum_{i \in N} p_i(b) . \quad \square$$

By applying our composition theorems, we obtain a constant price of anarchy for oblivious learning in this game even when we have multiple channels with fully independent or everybody-or-nobody availability. This simplifies and generalizes an approach based on sleeping expert learning in [25]. Furthermore, our analysis can also be conducted similarly for other interference models with a bounded independence condition, see [25, 26] for a discussion.

5.6 Extension to Changing Unit-Demand Functions

We now consider a case in which valuations change over time rather than the supply. In particular, we consider a unit-demand setting (see Definition 3.5). We assume that each of the $v_{i,j}$ is an independent random variable in which constantly many outcomes have a positive probability. So, for a fixed player i , the valuation is defined such that for $k = 1, \dots, K$ we let $v_{i,j} = v_{i,j}^{(k)}$ with probability $q_{i,j}^{(k)}$, $\sum_{k=1}^K q_{i,j}^{(k)} = 1$. Without loss of generality, let $v_{i,j}^{(1)} \geq v_{i,j}^{(2)} \geq \dots \geq v_{i,j}^{(K)}$.

To apply availability-oblivious learning, player i now makes K copies of each item j . The k th copy of item j has value $v_{i,j}^{(k)}$, and it is available whenever $v_{i,j} \geq v_{i,j}^{(k)}$. Note that, when restricting the consideration to only the most valuable item, we can equivalently assume that availabilities of items are drawn independently with probability $q_{i,j}^{(k)} / \sum_{k'=k}^K q_{i,j}^{(k')}$ for the k th copy of item j .

By the same argument as in Section 5.2, we then have

$$\begin{aligned} \sum_i \mathbf{E} [u_i(b'_i, b_{-i})] &= \sum_i \mathbf{E} [v_i(f(b'_i, b_{-i}))] - \mathbf{E} [p_i(b'_i, b_{-i})] \\ &\geq \left(1 - \frac{1}{e}\right) \sum_i \mathbf{E} [v_i(f(b_i^{A_i}, b_{-i}))] - \mathbf{E} [p_i(b_i^{A_i}, b_{-i})] , \end{aligned}$$

when comparing the availability-oblivious deviation b'_i with the availability-aware ones $b_i^{A_i}$.

Therefore, if each single mechanism is weakly (λ, μ_1, μ_2) -smooth, the price of anarchy for oblivious learning is at most $e/(e-1) \cdot (\mu_2 + \max(1, \mu_1))/\lambda$.

PART II

Combinatorial Secretary Problems with Ordinal Information

This part is the result of close collaboration with Martin Hoefer. It is based on an article that appeared in the *Proceedings of the 44th International Colloquium on Automata, Languages and Programming (ICALP) 2017*, pages 133:1–133:14, in July 2017 [52]. A full version is available at <http://arxiv.org/abs/1702.01290>.



CHAPTER 6

Introduction to Part II

The secretary problem is a classic approach to model online decision making under uncertain input. The interpretation is that a firm needs to hire a secretary. There are n candidates and they arrive sequentially in random order for an interview. Following an interview, the firm learns the value of the candidate, and it has to make an immediate decision about hiring him before seeing the next candidate(s). If the candidate is hired, the process is over. Otherwise, a rejected candidate cannot be hired at a later point in time. The optimal algorithm is a simple greedy rule that rejects all candidates in an initial learning phase. In the following acceptance phase, it hires the first candidate that is the best among all the ones seen so far. It manages to hire the best candidate with optimal probability that tends to $1/e$ when $n \rightarrow \infty$. Notably, it only needs to know if a candidate is the best seen so far, but no exact numerical values.

Since its introduction [35], the secretary problem has attracted a huge amount of research interest. Recently, a variety of combinatorial extensions have been studied in the computer science literature [11], capturing a variety of fundamental online allocation problems in networks and markets, such as network design [63], resource allocation [58], medium access in networks [45], or competitive admission processes [22]. Prominently, in the matroid secretary problem [12], the elements of a weighted matroid arrive in uniform random order (e.g., weighted edges of an undirected graph G). The goal is to select a max-weight independent set of the matroid (e.g., a max-weight forest of G). The popular *matroid secretary conjecture* claims that for all matroids, there exists an algorithm with a constant *competitive ratio*, i.e., the expected total weight of the solution computed by the algorithm is at least a constant fraction of the total weight of the optimum solution. Despite much progress on special cases, the conjecture remains open. Beyond matroids, online algorithms for a variety of combinatorial secretary problems with downward-closed structure have recently been studied (e.g., matching [63, 58], independent set [45], linear packing problems [59] or submodular versions [39, 60]).

The best known algorithms for matroid or matching secretary problems rely heavily on knowing the exact weight structure of elements. They either compute max-weight solutions to guide the admission process or rely on advanced bucketing techniques to group elements based on their weight. For a decision maker, in many applications it can be quite difficult to determine an exact cardinal preference for each of the incoming candidates. In contrast, in the original problem, the optimal algorithm only needs *ordinal information* about the candidates. This property provides a much more robust guarantee, since the numerical values can be arbitrary, as long as they are consistent with the preference order.

In Part II of this thesis, we study algorithms for combinatorial secretary problems that rely only on ordinal information. We assume that there is an unknown value for each element, but our algorithms only have access to the *total order* of the elements arrived so

far, which is consistent with their values. We term this the *ordinal model*; as opposed to the *cardinal model*, in which the algorithm learns the exact values. We show bounds on the competitive ratio, i.e., we compare the quality of the computed solutions to the optima in terms of the exact underlying but unknown numerical values. Consequently, competitive ratios for our algorithms are robust guarantees against uncertainty in the input. Our approach follows a recent line of research by studying the potential of algorithms with ordinal information to approximate optima based on numerical values [7, 6, 1, 20].

6.1 Our Contribution

We first point out that many algorithms proposed in the literature continue to work in the ordinal model. In particular, a wide variety of algorithms for variants of the matroid secretary problem with constant competitive ratios continue to obtain their guarantees in the ordinal model (see Table 6.1 for an overview). This shows that many results in the literature are much stronger than they claim to be, since the algorithms require significantly less information. Notably, the algorithm of [13] extends to the ordinal model and gives a ratio of $O(\log^2 r)$ for general matroids, where r is the rank of the matroid. In contrast, the improved algorithms with ratios of $O(\log r)$ and $O(\log \log r)$ [12, 64, 37] are not applicable in the ordinal model.

In Chapter 7 we extend a result of [39] for matroids to the ordinal model: The reduction from submodular to linear matroid secretary can be done with ordinal information on marginal weights of the elements.

Main Result 6. *Whenever there is an algorithm that solves the matroid secretary problem in the ordinal model on some matroid class and has a competitive ratio of α , there is also an algorithm for the submodular matroid secretary problem in the ordinal model on the same matroid class with a competitive ratio of $O(\alpha^2)$.*

The ratio can be shown to be better if the linear algorithm satisfies some further properties. Moreover, we consider the importance of knowing the weights, ordering, and structure of the domain.

Main Result 7. *For algorithms that have complete ordinal information but cannot learn the specific matroid structure, we show a lower bound of $\Omega(\sqrt{n}/(\log n))$, even for partition matroids, where n is the number of elements in the ground set.*

This bound contrasts the $O(\log^2 r)$ -competitive algorithm and indicates that learning the matroid structure is crucial in the ordinal model. Moreover, it contrasts the cardinal model, where thresholding algorithms yield $O(\log r)$ -competitive algorithms without learning the matroid structure.

In Chapter 8 we obtain new algorithms for the ordinal model for several combinatorial secretary problems. For online bipartite matching we give an algorithm that is $2e$ -competitive. We also extend this result to online packing LPs with at most d non-zero entries per variable. Here we obtain an $O(d^{(B+1)/B})$ -competitive algorithm, where B is a tightness parameter of the constraints. Another extension is matching in general graphs, for which we give a 8.78 -competitive algorithm. Finally, we give an $O(\alpha_1^2)$ -competitive algorithm for the online weighted independent set problem in graphs, where α_1 is the

local independence number of the graph. For example, for the prominent case of unit-disk graphs, $\alpha_1 = 5$ and we obtain a constant-competitive algorithm.

Main Result 8. *We give algorithms with small competitive ratios for the secretary versions of (bipartite) matching, packing and weighted independent set in the ordinal model.*

6.2 Related Work

Our work is partly inspired by [7, 8], who study ordinal approximation algorithms for classical optimization problems. They design constant-factor approximation algorithms for matching and clustering problems with ordinal information and extend the results to truthful mechanisms. Our approach here differs due to online arrival. Anshelevich et al. [6] examine the quality of randomized social choice mechanisms when agents have metric preferences but only ordinal information is available to the mechanism. Previously, [1, 20] studied ordinal measures of efficiency in matchings, for instance the average rank of an agent’s partner.

The literature on the secretary problem is too broad to survey here. We only discuss directly related work on online algorithms for combinatorial variants. Cardinal versions of these problems have many important applications in ad-auctions and item allocation in online markets [46]. For multiple-choice secretary, where we can select any k candidates, there are algorithms with ratios that are constant and asymptotically decreasing in k [61, 10]. More generally, the matroid secretary problem has attracted a large amount of research interest [12, 21, 64, 37], and the best-known algorithm in the cardinal model has ratio $O(\log \log r)$. For results on specific matroid classes, see the overview in Table 6.1. Extensions to the submodular version are treated in [13, 39].

Another prominent domain is online bipartite matching, in which one side of the graph is known in advance and the other arrives online in random order, each vertex revealing all incident weighted edges when it arrives [63]. In this case, there is an optimal algorithm with ratio e [58]. Moreover, our work is related to Göbel et al. [45] who study secretary versions of maximum independent set in graphs with bounded inductive independence number ρ . They derive an $O(\rho^2)$ -competitive algorithm for unweighted and an $O(\rho^2 \log n)$ -competitive algorithm for weighted independent set.

Matroid	general	k-uniform	graphic	cographic	transversal	laminar	regular
Ratio	$O(\log^2 r)$	$1 + O(\frac{e}{\sqrt{1/k}})$	$2e$	$3e$	16	$3\sqrt{3}e$	$9e$
Reference	[13]	[35, 10] [61]	[63]	[83]	[29]	[56]	[30]

Table 6.1: Existing algorithms for matroid secretary problems that provide the same guarantee in the ordinal model.

In addition, algorithms have been proposed for further variants of the secretary problem, e.g., the temp secretary problem (candidates hired for a fixed duration) [40], parallel secretary (candidates interviewed in parallel) [38], or local secretary (several firms and limited feedback) [22]. For these variants, some existing algorithms (e.g., for the temp secretary problem in [40]) directly extend to the ordinal model. In general, however, the restriction to ordinal information poses an interesting challenge for future work in these domains.

6.3 Notation and Preliminaries

In the typical problem we study, there is a set E of elements arriving sequentially in random order. The algorithm knows $n = |E|$ in advance. It must accept or reject an element before seeing the next element(s). There is a set $\mathcal{S} \subseteq 2^E$ of *feasible solutions*. \mathcal{S} is downward-closed, i.e., if $S \in \mathcal{S}$, then $S' \in \mathcal{S}$ for every $S' \subseteq S$. The goal is to accept a feasible solution that maximizes an objective function f . In the *linear version*, each element has a *value* or *weight* w_e , and $f(S) = \sum_{e \in S} w_e$. In the *submodular version*, f is submodular and $f(\emptyset) = 0$ (see Definition 3.2).

In the linear ordinal model, the algorithm only sees a strict total order over the elements seen so far that is consistent with their weights (ties are broken arbitrarily). For the submodular version, we interpret the value of an element as its marginal contribution to a set of elements. In this case, our algorithm has access to an *ordinal oracle* $\mathcal{O}(S)$. For every subset S of arrived elements, $\mathcal{O}(S)$ returns a total order of arrived elements consistent with their marginal values $f(e|S) = f(S \cup \{e\}) - f(S)$.

Given this information, we strive to design algorithms that will have a small competitive ratio $f(S^*)/\mathbf{E}[f(S_{alg})]$. Here S^* is an optimal feasible solution and S_{alg} the solution returned by the algorithm. Note that S_{alg} is a random variable due to random-order arrival and possible internal randomization of the algorithm.

In the matroid secretary problem, the pair $\mathcal{M} = (E, \mathcal{S})$ is a matroid. We summarize in Table 6.1 some of the existing results for classes of the (linear) problem that transfer to the ordinal model. The algorithms for all restricted matroid classes other than the graphic matroid assume a priori complete knowledge of the matroid – only weights are revealed online. The algorithms do not use cardinal information, their decisions are based only on ordinal information. As such, they translate directly to the ordinal model. Notably, the algorithm from [13] solves even the general submodular matroid secretary problem in the ordinal model.

6.3.1 Yao's Principle

Yao's principle is a commonly used technique for giving lower bound proofs for randomized algorithms. It states that the expected cost of a randomized algorithm on the worst case input, is no better than a worst-case random probability distribution of the deterministic algorithm which performs best for that distribution. Thus, to establish a lower bound on the performance of randomized algorithms, it suffices to find an appropriate distribution of difficult inputs, and to prove that no deterministic algorithm can perform well against that distribution. We present here the formal statement.

Theorem 6.1 (Yao's principle [88]). *Consider a problem over the inputs \mathcal{X} , and let \mathcal{A} be the set of all possible deterministic algorithms that correctly solve the problem. For any algorithm $a \in \mathcal{A}$ and input $x \in \mathcal{X}$, let $c(a, x) \geq 0$ be the cost of algorithm a run on input x .*

Let p be a probability distributions over the algorithms in \mathcal{A} , and let A denote a random algorithm chosen according to p . Let q be a probability distribution over the inputs \mathcal{X} , and let X denote a random input chosen according to q . Then,

$$\max_{x \in \mathcal{X}} \mathbf{E} [c(A, x)] \geq \min_{a \in \mathcal{A}} \mathbf{E} [c(a, X)] \quad ,$$

i.e., the worst-case expected cost of the randomized algorithm is at least the cost of the best deterministic algorithm against input distribution q .

CHAPTER 7

Matroids

7.1 Submodular Matroids

We start our analysis by showing that – in addition to algorithms for special cases mentioned in the previous chapter – a powerful technique for submodular matroid secretary problems [39] can be adjusted to work even in the ordinal model. More formally, in this section we show that there is a reduction from submodular matroid secretary problems with ordinal information (SMSPO) to linear matroid secretary problems with ordinal information (MSPO). The reduction uses Greedy (Algorithm 1) as a subroutine and interprets the marginal value when added to the greedy solution as the value of an element. These values are then forwarded to whichever algorithm (termed Linear) that solves the linear version of the problem. In the ordinal model, we are unable to see the exact marginal values. Nevertheless, we manage to construct a suitable ordering for the forwarded elements. Consequently, we can apply algorithm Linear as a subroutine to obtain a good solution for the ordinal submodular problem.

Let $\mathcal{M} = (E, \mathcal{S})$ be the matroid, f the submodular function, and E the ground set of elements. The marginal contribution of element u to set M is denoted by $f(u|M) = f(M \cup \{u\}) - f(M)$. Since f can be non-monotone, Greedy in the cardinal model also checks if the marginal value of the currently best element is positive. While we cannot explicitly make this check in the ordinal model, note that $f(u|M) \geq 0 \iff f(M \cup \{u\}) \geq f(M) = f(M \cup \{u'\})$ for every $u' \in M$. Since the ordinal oracle includes the elements of M in the ordering of marginal values, there is a way to check positivity even in the ordinal model. Therefore, our results also apply to non-monotone functions f .

A potential problem with Algorithm 2 is that we must compare marginal contributions of different elements w.r.t. different sets. We can resolve this issue by following the steps of the Greedy subroutine that tries to add new elements to the greedy solution computed on the sample. We use this information to construct a correct ordering over the marginal contributions of elements that we forward to Linear.

Lemma 7.1. *Let us denote by s_u the step of Greedy in which the element u is accepted when applied to $M + u$. Then $s_{u_1} < s_{u_2}$ implies $f(u_1|M_{u_1}) \geq f(u_2|M_{u_2})$.*

Proof. First, note that $M_{u_1} \subset M_{u_2}$ when $s_1 > s_2$. We denote by m_{u_1} the element of M that would be taken in step s_{u_1} if u_1 would not be available. Then we know that $f(u_1|M_{u_1}) \geq f(m_{u_1}|M_{u_1})$. Furthermore, since $s_1 < s_2$, $f(m_{u_1}|M_{u_1}) \geq f(u_2|M_{u_1})$. Lastly, by using submodularity, we know that $f(u_2|M_{u_1}) \geq f(u_2|M_{u_2})$. \square

When $s_{u_1} = s_{u_2}$, then $M_{u_1} = M_{u_2}$ so the oracle provides the order of marginal values. Otherwise, the lemma yields the ordinal information. Thus, we can construct an ordering for the elements that are forwarded to Linear that is consistent with their marginal values

Algorithm 1: Greedy [39]**Input** : ground set E **Output** : independent set M

-
- 1 Let $M \leftarrow \emptyset$ and $E' \leftarrow E$.
 - 2 **while** $E' \neq \emptyset$ **do**
 - 3 Let $u \leftarrow \max_{u'} f(u'|M)$ and $E' \leftarrow E' \setminus \{u\}$
 - 4 **if** $(M \cup \{u\}$ independent in $\mathcal{M}) \wedge (f(u|M) \geq 0)$ **then** add u to M
-

Algorithm 2: Online(p) algorithm [39]**Input** : $n = |E|$, size of the ground set**Output** : independent set $Q \cap N$

-
- 1 Choose X from the binomial distribution $B(n, 1/2)$.
 - 2 Reject the first X elements of the input. Let L be the set of these elements.
 - 3 Let M be the output of Greedy on the set L .
 - 4 Let $N \leftarrow \emptyset$.
 - 5 **for** each element $u \in E \setminus L$ **do**
 - 6 Let $w(u) \leftarrow 0$.
 - 7 **if** u accepted by Greedy applied to $M \cup \{u\}$ **then**
 - 8 With probability p do the following:
 - 9 Add u to N .
 - 10 Let $M_u \subseteq M$ be the solution of Greedy immediately before it adds u to it.
 - 11 $w(u) \leftarrow f(u|M_u)$.
 - 12 Pass u to Linear with weight $w(u)$.
 - 13 **return** $Q \cap N$, where Q is the output of Linear.
-

in the cardinal model. Hence, the reduction can be applied in the ordinal model, and all results from [39] continue to hold. We mention only the main theorem. It implies constant ratios for all problems in Table 6.1 in the ordinal submodular version.

Theorem 7.2. *Given an arbitrary algorithm Linear for MSPO that is α -competitive on a matroid class, there is an algorithm for SMSPO with competitive ratio is at most $24\alpha(3\alpha + 1) = O(\alpha^2)$ on the same matroid class. For SMSPO with monotone f , it can be improved to $8\alpha(\alpha + 1)$.*

7.2 A Lower Bound

Another powerful technique in the cardinal model is thresholding, where we first sample a constant fraction of the elements to learn their weights. Based on the largest weight observed, we pick a threshold and accept subsequent elements greedily if they exceed the threshold. This approach generalizes the classic algorithm [35] and provides logarithmic ratios for many combinatorial domains [12, 63, 45, 22]. Intuitively, these algorithms *learn the weights but not the structure*.

We show that this technique does not easily generalize to the ordinal model. The

algorithms with small ratios in the ordinal model rely heavily on the matroid structure. Indeed, in the ordinal model we show a polynomial lower bound for algorithms in the matroid secretary problem that learn the ordering but not the structure. Formally, we slightly simplify the setting as follows. The algorithm receives the global ordering of all elements in advance. It determines (possibly at random) a threshold position in the ordering. Then elements arrive and are accepted greedily if ranked above the threshold. Note that the algorithm does not use sampling, since in this case the only meaningful purpose of sampling is learning the structure. We call this a *structure-oblivious* algorithm.

For worst-case bounds, we can restrict our attention to instances where all elements have cardinal weights in $\{0, 1\}$. These instances always result in the worst competitive ratio, as shown in the following lemma.

Lemma 7.3. *By converting an arbitrary weighted instance to an instance with weights in $\{0, 1\}$, the competitive ratio between the optimum solution and the solution computed by an algorithm based on ordinal information can only deteriorate.*

Proof. Without loss of generality, we assume that all elements of the original instance have distinct weights. We denote the elements chosen in the optimal solution by a_1^*, \dots, a_k^* and the elements chosen by the algorithm by b_1, \dots, b_m . The numbering respects the ordering of weights, i.e., $a_1^* \succ a_2^* \succ \dots \succ a_k^*$ and $b_1 \succ b_2 \succ \dots \succ b_m$. The competitive ratio is

$$\frac{\text{OPT}}{\text{ALG}} = \frac{w(a_1^*) + \dots + w(a_k^*)}{w(b_1) + \dots + w(b_m)}.$$

This ratio can only increase if we change the weight of all elements that appear after a_k^* in the global ordering to 0. This effectively shortens the set of elements with a contribution chosen by the algorithm to b_1, \dots, b_ℓ , for some suitable $\ell \leq m$. Furthermore, we change the weights of all elements between a_i^* and a_{i+1}^* by decreasing them to a_{i+1}^* . We now denote the elements that the algorithm chose by c_1, \dots, c_ℓ , since their weights might have changed. Both of these changes do not influence OPT, but they reduce the weight of the solution returned by the algorithm. We continue converting the instance, by focusing on $w(a_k^*)$. Then,

$$\begin{aligned} \frac{\text{OPT}}{\text{ALG}} &= \frac{w(a_1^*) + \dots + w(a_k^*)}{w(b_1) + \dots + w(b_m)} \leq \frac{w(a_1^*) + \dots + w(a_k^*)}{w(b_1) + \dots + w(b_\ell)} \\ &\leq \frac{w(a_1^*) + \dots + w(a_k^*)}{w(c_1) + \dots + w(c_\ell)} = \frac{A + w(a_k^*)}{B + r \cdot w(a_k^*)}, \end{aligned}$$

where $A = w(a_1^*) + \dots + w(a_{k-1}^*)$, B is the sum of the weights of all elements that the algorithm chose which are not equal to $w(a_k^*)$ in the altered instance and $r \in \mathbb{N}_0$.

Taking the derivative for $w(a_k^*)$,

$$\frac{d}{d(w(a_k^*))} \left(\frac{A + w(a_k^*)}{B + r \cdot w(a_k^*)} \right) = \frac{B - r \cdot A}{(B + r \cdot w(a_k^*))^2},$$

we either decrease $w(a_k^*)$ to 0 or raise it to $w(a_{k-1}^*)$ (depending what makes the ratio increase, i.e., the sign of the derivative). We continue this procedure until all weights of the instance are equal to either to a_1^* or 0. Note that these changes preserve the global ordering. W.l.o.g., we can finally set $w(a_1^*) = 1$.




Figure 7.1: Values for the family of instances described in the proof of Theorem 7.4, where the position of the “valuable edges” is denoted by the thick segment.

Note that we increased the ratio between the solution of the algorithm and the optimal solution for the original weights, when applying the transformed weights. Note that none of these transformations change the decisions of the algorithm. In contrast, the optimum solution for the transformed weights can only become better, which even further deteriorates the competitive ratio. \square

Theorem 7.4. *Every structure-oblivious randomized algorithm has a competitive ratio of at least $\Omega(\sqrt{n}/(\log n))$.*

Proof. We give a distribution of such instances on which every deterministic algorithm has a competitive ratio of $\Omega(\sqrt{n}/(\log n))$. Using Yao’s principle (see Subsection 6.3.1), this shows the claimed result for randomized algorithms.

All instances in the distribution are based on a graphic matroid (in fact, a partition matroid) of the following form. There is a simple path of $1 + k$ segments. The edges in each segment have weight of 0 or 1. We call the edges with value 1 in the last k segments the “valuable edges”. The total number of edges is the same in each instance and equals $n + 1$. All edges in the first segment have value 1 and there is exactly one edge of value 1 in all other segments (that being the aforementioned valuable edges). In the first instance there are in total $k + 1$ edges of value 1 (meaning that there is only one edge in the first segment). In each of the following instances this number is increased by k (in the i -th instance there are $(i - 1) \cdot k + 1$ edges in the first segment) such that the last instance has only edges with value 1 (there are $n - k + 1$ edges in the first segment). The zero edges are always equally distributed on the last k path segments. The valuable edges are lower in the ordering than any non-valuable edge with value 1 (see Figure 7.1). Each of the instances appears with equal probability of $\frac{k}{n}$ (see Figure 7.2 for one example instance).

A deterministic algorithm picks a threshold at position i . The expected value of the solution is

$$\mathbb{E}[w(S_{alg})] \leq 1 + \frac{k}{n} \sum_{\ell=1}^{\frac{i}{k}} \frac{k}{\ell} \leq 1 + \frac{k^2}{n} \log \frac{i}{k} \leq \frac{k^2}{n} \log \frac{n}{k} + 1 ,$$

where \log denotes the natural logarithm and the expression results from observing that the algorithm cannot obtain more than a value of 1 if its threshold i falls above the




Figure 7.2: One instance from the family described in the proof of Theorem 7.4.

valuable 1's. Otherwise it gets an additional fraction of k , depending on how close the threshold is positioned to the valuable 1's. For instance, if the threshold is set between 1 and k positions below the valuable 1's, the algorithm will in expectation select edges of total value of at least $1 + k/2$. This follows from the random arrival order of the edges and the fact that the ratio of valuable to non-valuable edges that the algorithm is ready to accept is at least 1 : 2. Furthermore, we see that for this distribution of instances the optimal way to set a deterministic threshold is at the lowest position. Using $k = \sqrt{n}$, a lower bound on the competitive ratio is

$$\frac{k}{\frac{k^2}{n} \log \frac{n}{k} + 1} = \frac{n}{k \log \frac{n}{k} + \frac{n}{k}} = \Omega\left(\frac{\sqrt{n}}{\log n}\right). \quad \square$$

CHAPTER 8

Matching, Packing and Independent Set

8.1 Bipartite Matching

In this section, we study online bipartite matching. The vertices on the right side of the graph (denoted by R) are static and given in advance. The vertices on the left side (denoted by L) arrive sequentially in a random order. Every edge $e = (r, \ell) \in R \times L$ has a non-negative weight $w(e) \geq 0$. In the cardinal model, each vertex of L reveals upon arrival the weights of all incident edges. In the ordinal model, we are given a total order on all edges that have arrived so far, consistent with their weights. Before seeing the next vertex of L , the algorithm has to decide to which vertex $r \in R$ (if any) it wants to match the current vertex ℓ . A match that is formed cannot be revoked. The goal is to maximize the total weight of the matching.

The algorithm for the cardinal model in [58] achieves an optimal competitive ratio of e . However, this algorithm heavily exploits cardinal information by repeatedly computing max-weight matchings for the edges seen so far. For the ordinal model, our Algorithm 3 below obtains a competitive ratio of $2e$. While similar in spirit, the main difference is that we rely on a greedy matching algorithm, which is based solely on ordinal information. It deteriorates the ratio only by a factor of 2.

Lemma 8.1. *Let the random variable A_v denote the contribution of the vertex $v \in L$ to the output, i.e., weight assigned to v in M . Let $w(M^*)$ denote the value of the maximum-weight matching in G . For $\ell \in \{\lceil \frac{n}{e} \rceil, \dots, n\}$,*

$$\mathbb{E}[A_\ell] \geq \frac{\lfloor \frac{n}{e} \rfloor}{\ell - 1} \cdot \frac{w(M^*)}{2n}.$$

Proof. We first show that $e^{(\ell)}$ has a significant expected weight. Then we bound the probability of adding $e^{(\ell)}$ to M .

In step ℓ , $|L'| = \ell$ and the algorithm computes a greedy matching $M^{(\ell)}$ on $G[L' \cup R]$. The current vertex ℓ can be seen as selected uniformly at random from L' , and L' can be seen as selected uniformly at random from L . Therefore, $\mathbb{E}[w(M^{(\ell)})] \geq \frac{\ell}{n} \cdot \frac{w(M^*)}{2}$ and $\mathbb{E}[w(e^{(\ell)})] \geq \frac{w(M^*)}{2n}$. Here we use that a greedy matching approximates the optimum by at most a factor of 2 [5].

Edge $e^{(\ell)}$ can be added to M if r has not been matched already. The vertex r can be matched only when it is in $M^{(k)}$. The probability of r being matched in step k is at most $\frac{1}{k}$ and the order of the vertices in steps $1, \dots, k-1$ is irrelevant for this event.

$$\Pr[r \text{ unmatched in step } \ell] = \Pr \left[\bigwedge_{k=\lceil n/e \rceil}^{\ell-1} r \notin e^{(k)} \right] \geq \prod_{k=\lceil n/e \rceil}^{\ell-1} \frac{k-1}{k} = \frac{\lceil \frac{n}{e} \rceil - 1}{\ell - 1}$$

Algorithm 3: Bipartite Matching**Input** : vertex set R and cardinality $n = |L|$ **Output** : matching M

- 1 Let L' be the first $\lfloor \frac{n}{e} \rfloor$ vertices of L , and $M \leftarrow \emptyset$
- 2 **for** each $\ell \in L \setminus L'$ **do**
- 3 $L' \leftarrow L' \cup \{\ell\}$
- 4 $M^{(\ell)} \leftarrow$ greedy matching on $G[L' \cup R]$
- 5 Let $e^{(\ell)} \leftarrow (\ell, r)$ be the edge assigned to ℓ in $M^{(\ell)}$
- 6 **if** $M \cup \{e^{(\ell)}\}$ is a matching **then** add $e^{(\ell)}$ to M

We now know that $\Pr [M \cup e^{(\ell)} \text{ is a matching}] \geq \frac{\lfloor n/e \rfloor}{\ell-1}$. Using this and $\mathbb{E}[w(e^{(\ell)})] \geq \frac{w(M^*)}{2n}$, the lemma follows. \square

Theorem 8.2. *Algorithm 3 for bipartite matching is $2e$ -competitive.*

Proof. The weight of matching M can be obtained by summing over random variables A_ℓ .

$$\mathbb{E}[w(M)] = \mathbb{E} \left[\sum_{\ell=1}^n A_\ell \right] \geq \sum_{\ell=\lfloor n/e \rfloor}^n \frac{\lfloor n/e \rfloor}{\ell-1} \cdot \frac{w(M^*)}{2n} = \frac{\lfloor n/e \rfloor}{2n} \sum_{\ell=\lfloor n/e \rfloor}^{n-1} \frac{1}{\ell} \cdot w(M^*)$$

Since $\frac{\lfloor n/e \rfloor}{n} \geq \frac{1}{e} - \frac{1}{n}$ and $\sum_{\ell=\lfloor n/e \rfloor}^{n-1} \frac{1}{\ell} \geq \ln \frac{n}{\lfloor n/e \rfloor} \geq 1$, it follows that

$$\mathbb{E}[w(M)] \geq \left(\frac{1}{e} - \frac{1}{n} \right) \cdot \frac{w(M^*)}{2} . \quad \square$$

Here we assumed to have access to ordinal preferences over all the edges in the graph. Note that the same approach works if the vertices provide correlated (ordinal) preference lists consistent with the edge weights, for every vertex from R and every arrived vertex from L . In this case, the greedy algorithm can still be implemented by iteratively matching and removing a pair that mutually prefers each other the most, and it provides an approximation guarantee of 2 for the max-weight matching (see, e.g., [5]). In contrast, if we receive only preference lists for vertices on one side, there are simple examples that establish super-constant lower bounds on the competitive ratio¹.

In the submodular version of the offline problem, the natural greedy algorithm gives a 3-approximation [43]. It builds the matching by greedily adding an edge that maximizes the marginal improvement of f , which is the information delivered by the ordinal oracle. When using this algorithm as a subroutine for the bipartite matching secretary problem, the resulting procedure achieves a 12-approximation in the submodular case [60]. Since greedy works based on an ordinal oracle, the algorithm can be applied to give the same ratio in the ordinal model.

¹Consider a bipartite graph with two nodes on each side (named A,B and 1,2). If we only know that both A and B prefer 1 to 2, the ratio becomes at least 2 even in the offline case. Similar examples imply that the (offline) ratio must grow in the size of the graph.

Algorithm 4: Packing LP

Input : capacities \mathbf{b} , total number of requests n , probability $p = \frac{e(2d)^{1/B}}{1+e(2d)^{1/B}}$

Output : assignment vector \mathbf{y}

- 1 Let L' be the first $p \cdot n$ requests, and $\mathbf{y} \leftarrow \mathbf{0}$
- 2 **for** each $j \notin L'$ **do**
- 3 $L' \leftarrow L' \cup \{j\}$
- 4 $\mathbf{x}^{(L')} \leftarrow$ greedy assignment on the LP for L'
- 5 $\mathbf{y}_j \leftarrow \mathbf{x}_j^{(L')}$
- 6 **if** $\neg(\mathbf{A}(\mathbf{y}) \leq \mathbf{b})$ **then** $\mathbf{y}_j \leftarrow \mathbf{0}$

8.2 Packing

Our results for bipartite matching can be extended to online packing LPs of the form $\max \mathbf{c}^\top \mathbf{x}$ s.t. $\mathbf{A}\mathbf{x} \leq \mathbf{b}$ and $\mathbf{0} \leq \mathbf{x} \leq \mathbf{1}$, which model problems with m resources and n requests arriving in random order. Each resource $i \in [m]$ has a capacity b_i that is known in advance. The number of requests is also known in advance. Every online request comes with a set of options, where each option has its profit and resource consumption. Once a request arrives, the coefficients of its variables are revealed and the assignment to the variables has to be determined.

Formally, request $j \in [n]$ corresponds to variables $x_{j,1}, \dots, x_{j,K}$ that represent K options. Each option $k \in [K]$ contributes with profit $c_{j,k} \geq 0$ and has resource consumption $a_{i,j,k} \geq 0$ for resource i . In the ordinal model, we assume access to the global order of the arrived elements by their $c_{j,k}$ values. Overall, at most one option can be selected, i.e., there is a constraint $\sum_{k \in [K]} x_{j,k} \leq 1, \forall j \in [n]$. The objective is to maximize total profit while respecting the resource capacities. The offline problem is captured by the following linear program:

$$\begin{aligned} \max \quad & \sum_{j \in [n]} \sum_{k \in [K]} c_{j,k} x_{j,k} \quad \text{s.t.} \quad \sum_{j \in [n]} \sum_{k \in [K]} a_{i,j,k} x_{j,k} \leq b_i & i \in [m] \\ & \sum_{k \in [K]} x_{j,k} \leq 1 & j \in [n] \end{aligned}$$

As a parameter, we denote by d the maximum number of non-zero entries in any column of the constraint matrix \mathbf{A} , for which by definition $d \leq m$. We compare the solution to the fractional optimum, which we denote by \mathbf{x}^* . The competitive ratio will be expressed in terms of d and the capacity ratio $B = \min_{i \in [m]} \left\lfloor \frac{b_i}{\max_{j \in [n], k \in [K]} a_{i,j,k}} \right\rfloor$.

Kesselheim et al. [58] propose an algorithm that heavily exploits cardinal information – it repeatedly solves an LP-relaxation and uses the solution as a probability distribution over the options. Instead, our Algorithm 4 for the ordinal model is based on greedy assignments in terms of profits $c_{j,k}$. More specifically, the greedy assignment considers variables $x_{j,k}$ in non-increasing order of $c_{j,k}$. It sets a variable to 1 if this does not violate the capacity constraints, and to 0 otherwise.

Theorem 8.3. *Algorithm 4 for online packing LPs is $O(d^{(B+1)/B})$ -competitive.*

The proof will be based on the following lemma.

Lemma 8.4. *Let the random variable A_ℓ denote the contribution of request \mathbf{x}_ℓ to the output and $\mathbf{c}^\tau \mathbf{x}^*$ the value of the optimal fractional solution. For requests \mathbf{x}_j such that $\mathbf{x}_j \in \{pn + 1, \dots, n\}$, it holds that*

$$\mathbb{E}[A_j] \geq \left(1 - d \cdot \left(\frac{e(1-p)}{p}\right)^B\right) \frac{\mathbf{c}^\tau \mathbf{x}^*}{(d+1)n}.$$

Proof. If $x_{j,k}^{(L')} = 1$, then as in the proof of Lemma 8.1, we get $\mathbb{E}[\mathbf{c}_j \mathbf{x}_j^{(L')}] = \mathbb{E}[c_{j,k} x_{j,k}^{(L')}] \geq \frac{\mathbf{c}^\tau \mathbf{x}^*}{(d+1)n}$, where the expectation is taken over the choice of the set L' and the choice of the last vertex in the order of arrival.

The algorithm sets \mathbf{y}_j to $\mathbf{x}_j^{(L')}$ only if the capacity constraints can be respected. For the sake of analysis, we assume that the algorithm only sets \mathbf{y}_j to $\mathbf{x}_j^{(L')}$ if every capacity constraint b_i that $\mathbf{x}_j^{(L')}$ affects ($x_{j,k}^{(L')} = 1$ and $a_{i,j,k} \neq 0$) is affected by at most $B - 1$ previous requests. We bound the probability of a capacity constraint b_i being affected in any preceding step $s \in \{pn + 1, \dots, j - 1\}$, for a fixed i :

$$\begin{aligned} \Pr[b_i \text{ affected by } x_{s,k'}^{(L')} = 1] &\leq \sum_{\mathbf{x}_{j'} \in \{1, \dots, s\}} \Pr[(\mathbf{x}_{j'} \text{ is last in the order}) \wedge (a_{i,j',k'} \neq 0)] \\ &\leq \frac{1}{s} \sum_{\mathbf{x}_{j'} \in \{1, \dots, s\}} \Pr[a_{i,j',k'} \neq 0] \leq \frac{B}{s}, \end{aligned}$$

where the last step follows from \mathbf{y} being a feasible solution throughout the run of the algorithm. Now, we bound the probability of not being able to set \mathbf{y}_j to $\mathbf{x}_j^{(L')}$:

$$\begin{aligned} \Pr[b_i \text{ is affected at least } B \text{ times}] &\leq \sum_{\substack{C \subseteq \{pn+1, \dots, j-1\}, \\ |C|=B}} \left(\prod_{s \in C} \frac{B}{s} \right) \\ &\leq \binom{(1-p)n}{B} \cdot \left(\frac{B}{pn} \right)^B \\ &\leq \left(\frac{(1-p)e}{p} \right)^B, \end{aligned}$$

so the probability of succeeding in setting \mathbf{y}_j to $\mathbf{x}_j^{(L')}$ is

$$\Pr[\mathbf{A}\mathbf{y} \leq \mathbf{b}] \geq 1 - d \cdot \left(\frac{(1-p)e}{p} \right)^B,$$

because we can do a union bound over all b_i that are affected by $\mathbf{x}_j^{(L')}$ and there are at most d such, since that is the maximal number of non-zero entries in any column of the constraint matrix \mathbf{A} .

Combining this with the inequality regarding the expected contribution of $\mathbf{x}_j^{(L')}$, we get the claimed result. \square

Algorithm 5: General Matching**Input** : vertex set V and cardinality $n = |V|$ **Output** : matching M

- 1 Let R be the first $\lfloor \frac{n}{2} \rfloor$ vertices of V
- 2 Let L' be the further $\lfloor \frac{n}{2e} \rfloor$ vertices of V , and $M \leftarrow \emptyset$
- 3 **for** each $\ell \in V \setminus L'$ **do**
- 4 $L' \leftarrow L' \cup \{\ell\}$
- 5 $M^{(\ell)} \leftarrow$ greedy matching on $G[L' \cup R]$
- 6 Let $e^{(\ell)} \leftarrow (\ell, r)$ be the edge assigned to ℓ in $M^{(\ell)}$
- 7 **if** $M \cup \{e^{(\ell)}\}$ is a matching **then** add $e^{(\ell)}$ to M

Proof of Theorem 8.3. Using Lemma 8.4, we get

$$\begin{aligned}
\mathbb{E}[\mathbf{c}^\tau \mathbf{y}] &= \sum_{\ell=pn+1}^n \mathbb{E}[A_\ell] \\
&\geq \sum_{\ell=pn+1}^n \left(1 - d \cdot \left(\frac{e(1-p)}{p} \right)^B \right) \frac{\mathbf{c}^\tau \mathbf{x}^*}{(d+1)n} \\
&= \frac{\mathbf{c}^\tau \mathbf{x}^*}{d+1} \cdot \frac{1}{1 + e(2d)^{1/B}} \cdot \left(1 - d \cdot \left(\frac{1}{(2d)^{1/B}} \right)^B \right) \\
&\geq \frac{\mathbf{c}^\tau \mathbf{x}^*}{2(d+1)(1 + 2ed^{1/B})} . \quad \square
\end{aligned}$$

Note that Theorem 8.3 contains the one-sided b -hypermatching problem as a special case. For the even more special case of $b = 1$ in the one-sided hypermatching, an algorithm was given in [63], which also works in the ordinal model. Our ratio in this special case is similar, but our approach extends to arbitrary capacities $b \geq 1$.

8.3 Matching in General Graphs

Here we study the case when vertices of a general undirected graph arrive in random order. In the beginning, we only know the number n of vertices. Each edge in the graph has a non-negative weight $w(e) \geq 0$. Each vertex reveals the incident edges to previously arrived vertices and their weights (cardinal model), or we receive a total order over all edges among arrived vertices that is consistent with the weights (ordinal model). An edge can be added to the matching only in the round in which it is revealed. The goal is to construct a matching with maximum weight.

We can tackle this problem by prolonging the sampling phase and dividing the vertices into “left” and “right” vertices. Algorithm 5 first samples $n/2$ vertices. These are assigned to be the set R , corresponding to the static side of the graph in bipartite matching. The remaining vertices are assigned to be the set L . The algorithm then proceeds by sampling a fraction of the vertices of L , forming a set L' . The remaining steps are exactly the same as in Algorithm 3.

Theorem 8.5. *Algorithm 5 for matching in general graphs is $12e/(e+1)$ -competitive, where $12e/(e+1) < 8.78$.*

The proof is based on the following lemma.

Lemma 8.6. *Let the random variable A_ℓ denote the contribution of vertex $\ell \in [n/2 + n/(2e)]$ to the output, i.e., the weight of the edge assigned to ℓ in M . Then,*

$$\mathbb{E}[A_\ell] \geq \frac{\lceil \frac{n}{2} + \frac{n}{2e} \rceil}{\ell - 1} \cdot \frac{1}{2} \cdot \frac{w(M^*)}{n}.$$

Proof. The proof is similar to the one of Lemma 8.1, with the additional observation that each edge is available with probability $\frac{1}{2}$. It is available only if the incident vertices are assigned to different sides of the bipartition. \square

Proof of Theorem 8.5. We now use the lemma to bound as follows:

$$\begin{aligned} \mathbb{E}[w(M)] &= \mathbb{E}\left[\sum_{\ell=1}^n A_\ell\right] \geq \sum_{\ell=\lceil n/2+n/(2e) \rceil}^n \frac{\lceil n/2 + n/(2e) \rceil}{\ell - 1} \cdot \frac{1}{2} \cdot \frac{w(M^*)}{n} \\ &= \frac{\lceil n/2 + n/(2e) \rceil}{n} \cdot \frac{w(M^*)}{2} \cdot \sum_{\ell=\lceil n/2+n/(2e) \rceil}^{n-1} \frac{1}{\ell} \\ &\geq \left(\frac{1}{2}\left(1 + \frac{1}{e}\right) - \frac{1}{n}\right) \cdot \frac{w(M^*)}{2} \cdot \frac{1}{3}. \end{aligned} \quad \square$$

8.4 Independent Set and Local Independence

In this section, we study maximum independent set in graphs with bounded local independence number. The set of elements are the vertices V of an underlying undirected graph G . Each vertex has a weight $w_v \geq 0$. We denote by $N(v)$ the set of direct neighbors of vertex v . Vertices arrive sequentially in random order and reveal their position in the order of weights of vertices seen so far. The goal is to construct an independent set of G with maximum weight. The exact structure of G is unknown, but we know that G has a bounded local independence number α_1 .

Definition 8.7 (Local independence number). *An undirected graph G has local independence number α_1 if for each node v , the cardinality of every independent set in the neighborhood $N(v)$ is at most α_1 .*

We propose Algorithm 6, which is inspired by the Sample-and-Price algorithm for matching in [63]. Note that Göbel et al. [45] construct a more general approach for graphs with bounded *inductive* independence number ρ . However, they only obtain a ratio of $O(\rho^2 \log n)$ for the weighted version, where a competitive ratio of $\Omega(\log n / \log^2 \log n)$ cannot be avoided, even in instances with constant ρ . These algorithms rely on ρ -approximation algorithms for the offline problem that crucially exploit cardinal information.

Similar to the analysis in [63], we reformulate Algorithm 6 into an equivalent approach termed ‘‘Simulate’’ (Algorithm 7). Given the same arrival order, the same vertices are

Algorithm 6: Independent Set in Graphs with Bounded Local Independence Number

Input : $n = |G|$, $p = \sqrt{\alpha_1/(\alpha_1 + 1)}$
Output : independent set of vertices S

- 1 Set $k \leftarrow \text{Binom}(n, p)$, $S \leftarrow \emptyset$
- 2 Reject first k vertices of G , denote this set by G'
- 3 Build a maximal independent set of vertices from G' greedily, denote this set by M_1
- 4 **for** each $v \in G \setminus G'$ **do**
- 5 $w^* \leftarrow \max\{w \mid \mathcal{N}(v) \cap M_1\}$
- 6 **if** $(v > w^*) \wedge (S \cup \{v\} \text{ independent set})$ **then** add v to S

Algorithm 7: Simulate

Input : $n = |G|$, $p = \sqrt{\alpha_1/(\alpha_1 + 1)}$
Output : independent set of vertices S

- 1 Sort all vertices in G in non-increasing order of value
- 2 Initialize $M_1, M_2 \leftarrow \emptyset$
- 3 **for** each $v \in G$ in sorted order **do**
- 4 **if** $M_1 \cup \{v\}$ independent set **then**
- 5 flip a coin with probability p of heads
- 6 **if** heads **then** $M_1 \leftarrow M_1 \cup \{v\}$; **else** $M_2 \leftarrow M_2 \cup \{v\}$
- 7 $S \leftarrow M_2$
- 8 **for** each $w \in S$ **do**
- 9 **if** w has neighbors in S **then** remove w and all his neighbors from S

in the sample. Algorithm 7 drops all vertices from S that have neighbors in S while Algorithm 6 keeps one of them. Hence, $\mathbf{E}[w(S_{Alg6})] \geq \mathbf{E}[w(S_{Alg7})]$. In what follows, we analyze the performance of Algorithm 7. The first lemma follows directly from the definition of the local independence number.

Lemma 8.8. $\mathbf{E}[w(M_1)] \geq p \cdot \frac{w(S^*)}{\alpha_1}$, where $\alpha_1 \geq 1$ is the local independence number of the graph G .

Lemma 8.9. $\mathbf{E}[|\mathcal{N}(v) \cap M_2| \mid v \in M_2] \leq \frac{\alpha_1(1-p)}{p}$.

Proof. Let us denote by X_u^1 and X_u^2 the indicator variables for the events $u \in M_1$ and $u \in M_2$ respectively. Then,

$$\begin{aligned} \mathbf{E}[|\mathcal{N}(v) \cap M_2| \mid v \in M_2] &= \mathbf{E}\left[\sum_{u \in \mathcal{N}(v)} X_u^2 \mid v \in M_2\right] = \sum_{u \in \mathcal{N}(v)} \mathbf{E}[X_u^2 \mid v \in M_2] \\ &= \frac{1-p}{p} \sum_{u \in \mathcal{N}(v)} \mathbf{E}[X_u^1 \mid v \in M_2] \leq \frac{1-p}{p} \cdot \alpha_1. \quad \square \end{aligned}$$

Theorem 8.10. *Algorithm 7 for weighted independent set is $O(\alpha_1^2)$ -competitive, where α_1 is the local independence number of the graph G .*

Proof. By using Markov's inequality and Lemma 8.9,

$$\begin{aligned} \Pr [|\mathcal{N}(v) \cap M_2| \geq 1 \mid v \in M_2] &\leq \alpha_1 \cdot (1-p)/p \\ \text{and } \Pr [|\mathcal{N}(v) \cap M_2| < 1 \mid v \in M_2] &> 1 - (\alpha_1(1-p)/p) . \end{aligned}$$

Thus, we can conclude that

$$\mathbf{E}[w(S)] \geq \left(1 - \alpha_1 \cdot \frac{1-p}{p}\right) \cdot \mathbf{E}[w(M_2)] \geq \left(1 - \alpha_1 \cdot \frac{1-p}{p}\right) \cdot \frac{1-p}{\alpha_1} \cdot w(S^*) .$$

The ratio is optimized for $p = \sqrt{\frac{\alpha_1}{\alpha_1+1}}$, which proves the theorem. \square

As a prominent example, $\alpha_1 = 5$ in the popular class of unit-disk graphs. In such graphs, our algorithm yields a constant competitive ratio for online independent set in the ordinal model.

CHAPTER 9

Conclusion and Open Problems

In Part I of this thesis, we have studied two extensions of the smoothness framework. Namely, we were interested in introducing two aspects: risk-averse players and mechanism availability.

Risk Aversion In Chapter 4, we give bounds on the price of anarchy of Bayes-Nash and (coarse) correlated equilibria for smooth mechanisms that allow participation of risk-averse agents. We identify a sufficient and necessary condition that smooth mechanisms have to fulfill in order to achieve constant price of anarchy bounds in the presence of risk-averse players. Our main results lie within the concave utility function model.

An interesting further step would be to analyze risk aversion in the same concave utility function setting but with social welfare being redefined as revenue plus sum of the risk-averse utilities, but where we apply an inverse of the respective concave function to each of the summands. Such approach makes the normalization assumption obsolete and might be a yet better way to model risk aversion. Of course, what also remains open are settings that lie outside of the smoothness framework.

Simultaneous Composition with Varying Availability In Chapter 5, we have studied an oblivious variant for no-regret learning in repeated games with incomplete information and proved a composition theorem for smooth mechanisms. The bounds show that even if bidders apply learning algorithms independently of their types, they can still obtain outcomes that approximate the optimal social welfare within a small ratio.

Our primary motivation were changes over time on the supply side. That is, bidders value items the same at all times but are constrained when they can buy them. A different interpretation that leads to the same model is when bidders value items differently from time to time. Here the valuation for a bundle has the special structure that it is given by the value of a fixed submodular function evaluated on the intersection of this bundle with a random set.

There is potential to generalize the type-oblivious approach to other interesting settings where types are not only connected to availability of mechanisms but capture more general conditions for the bidder. For example, one could consider general independent types, where the complete availability-vector of a single bidder is drawn from a bidder-specific distribution, and for each bidder this is done independently. In Section 5.6, we give a partial answer to this direction of investigation and show how our techniques can be extended to the case of simultaneous single-item auctions with unit-demand valuations.

Lastly, within our scenario with admission we succeeded to show composition theorems for fully independent and “everybody or nobody” availability. It would be interesting

to see in which way we can change the degree of correlation among bidders and/or mechanisms and maintain a small price of anarchy. Previous results on channel access in wireless networks [25] indicate that with arbitrary correlation among bidders, the price of anarchy increases to $\Theta(n)$, even for the constantly smooth all-pay auctions outlined in Section 5.5 and even if availability remains independent among mechanisms. Again, it would be interesting to characterize the domains for which a small price of anarchy can be shown.

Ordinal Secretary Problems In Part II of this thesis, we study algorithms for combinatorial secretary problems that rely only on ordinal information. We show bounds on the competitive ratio, i.e., we compare the quality of the computed solutions to the optima in terms of the exact underlying but unknown numerical values. Consequently, competitive ratios for our algorithms are robust guarantees against uncertainty in the input.

Continuing this line of work, it would be interesting to further explore the differences between the cardinal and ordinal secretary scenarios. For instance, one could try to get better lower bounds for the ordinal model. As of yet, this matter is not well understood even in the simplest cases, such as in weighted bipartite matching.

Furthermore, the authors in [82] study secretary problems in which the algorithm is restricted to an arbitrary downward-closed set system. The broadest domain we considered in the ordinal model are matroid constraints. It would be of interest to see whether the ordinal model can also yield results for this more general set of constraints.

Finally, a more general research direction would be to explore different ordinal models and try to quantify the relation of the amount of information obtainable from the model and the achievable approximation ratios.

List of Figures

4.1	Normalized risk-averse utility function (bold) and quasilinear utility function for a fixed allocation \mathbf{x} and varying payment p_i	22
7.1	Values for the family of instances described in the proof of Theorem 7.4, where the position of the “valuable edges” is denoted by the thick segment.	70
7.2	One instance from the family described in the proof of Theorem 7.4.	71

List of Tables

6.1 Existing algorithms for matroid secretary problems that provide the same guarantee in the ordinal model.	63
--	----

List of Algorithms

1	Greedy [39]	68
2	Online(p) algorithm [39]	68
3	Bipartite Matching	74
4	Packing LP	75
5	General Matching	77
6	Independent Set in Graphs with Bounded Local Independence Number .	79
7	Simulate	79

List of Acronyms

VCG	Vickrey-Clark-Groves mechanism (a generic truthful mechanism for achieving a socially-optimal solution)
GSP	generalized second-price auction (an auction system for selling multiple items in one round of bids that is not truthful)
PoA	Price of Anarchy (the worst possible ratio between the optimal expected welfare and the expected welfare at equilibrium)
XOS	fractionally subadditive (the maximum of several additive set functions)
CE	correlated equilibrium (a distribution over action profiles such that after a profile is drawn, playing his part of the profile is a best response for a player conditioned on seeing his part, given that everyone else will play according to the profile)
CCE	coarse correlated equilibrium (a distribution over action profiles such that after a profile is drawn, playing his part of the profile is a best response for a player, given that everyone else will play according to the profile)
BNE	Bayes-Nash equilibrium (a strategy profile that maximizes the expected payoff for each player, given their type and given that everyone else will play according to the profile)

Bibliography

- [1] D. Abraham, R. Irving, T. Kavitha, and K. Mehlhorn. Popular matchings. *SIAM J. Comput.*, 37(4):1030–1045, 2007.
- [2] S. Agrawal. *Optimization under Uncertainty: Bounding the Correlation Gap*. PhD thesis, Stanford University, 2011.
- [3] S. Agrawal, Y. Ding, A. Saberi, and Y. Ye. Correlation robust stochastic optimization. In *Proc. 21st Symp. Discr. Algorithms (SODA)*, pages 1087–1096, 2010.
- [4] M. Andrews and M. Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In *Proc. 28th IEEE Conf. Computer Communications (INFOCOM)*, pages 1332–1340, 2009.
- [5] E. Anshelevich and M. Hoefer. Contribution games in networks. *Algorithmica*, 63(1–2):51–90, 2012.
- [6] E. Anshelevich and J. Postl. Randomized social choice functions under metric preferences. In *Proc. 25th Intl. Joint Conf. Artif. Intell. (IJCAI)*, pages 46–59, 2016.
- [7] E. Anshelevich and S. Sekar. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. In *Proc. 13th Conf. Artificial Intelligence (AAAI)*, pages 390–396, 2016.
- [8] E. Anshelevich and S. Sekar. Truthful mechanisms for matching and clustering in an ordinal world. In *Proc. 12th Conf. Web and Internet Economics (WINE)*, pages 265–278, 2016.
- [9] E. I. Asgeirsson and P. Mitra. On a game theoretic approach to capacity maximization in wireless networks. In *Proc. 30th IEEE Conf. Computer Communications (INFOCOM)*, pages 3029–3037, 2011.
- [10] M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg. A knapsack secretary problem with applications. In *Proc. 10th Workshop Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 16–28, 2007.
- [11] M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg. Online auctions and generalized secretary problems. *SIGecom Exchanges*, 7(2), 2008.
- [12] M. Babaioff, N. Immorlica, and R. Kleinberg. Matroids, secretary problems, and online mechanisms. In *Proc. 18th Symp. Discrete Algorithms (SODA)*, pages 434–443, 2007.
- [13] M. Bateni, M. Hajiaghayi, and M. Zadimoghaddam. Submodular secretary problem and extensions. *ACM Trans. Algorithms*, 9(4):32, 2013.

- [14] D. Bernoulli. Exposition of a new theory on the measurement of risk. *Commentaries of the Imperial Academy of Science of Saint Petersburg*, 1738. Republished *Econometrica* 22(1), 23–36 (1954).
- [15] A. Bhargat, T. Chakraborty, and S. Khanna. Mechanism design for a risk averse seller. In *Proc. 8th Intl. Workshop Internet & Network Economics (WINE)*, pages 198–211, 2012.
- [16] K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proc. 22nd Symp. Discrete Algorithms (SODA)*, pages 700–709, 2011.
- [17] A. Blum and Y. Mansour. Learning, regret minimization, and equilibria. In N. Nisan, É. Tardos, T. Roughgarden, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 4. Cambridge University Press, 2007.
- [18] Y. Cai and C. Papadimitriou. Simultaneous bayesian auctions and computational complexity. In *Proc. 15th Conf. Econom. Comput. (EC)*, pages 895–910, 2014.
- [19] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R. P. Leme, and É. Tardos. On the efficiency of equilibria in generalized second price auctions. *J. Econom. Theory*, 156:343–388, 2015.
- [20] D. Chakraborty and C. Swamy. Welfare maximization and truthfulness in mechanism design with ordinal preferences. In *Proc. 5th Symp. Innovations in Theoret. Computer Science (ITCS)*, pages 105–120, 2014.
- [21] S. Chakraborty and O. Lachish. Improved competitive ratio for the matroid secretary problem. In *Proc. 23rd Symp. Discrete Algorithms (SODA)*, pages 1702–1712, 2012.
- [22] N. Chen, M. Hoefer, M. Künnemann, C. Lin, and P. Miao. Secretary markets with local information. In *Proc. 42nd Intl. Coll. Automata, Languages and Programming (ICALP)*, volume 2, pages 552–563, 2015.
- [23] G. Christodoulou, A. Kovács, and M. Schapira. Bayesian combinatorial auctions. In *Proc. 35th Intl. Coll. Autom. Lang. Program. (ICALP)*, pages 820–832, 2008.
- [24] G. Christodoulou, A. Kovács, A. Sgouritsa, and B. Tang. Tight bounds for the price of anarchy of simultaneous first price auctions. *ACM Trans. Econom. Comput.*, 4(2):9, 2016.
- [25] J. Dams, M. Hoefer, and T. Kesselheim. Sleeping experts in wireless networks. In *Proc. 27th Symp. Distrib. Comput. (DISC)*, pages 344–357, 2013.
- [26] J. Dams, M. Hoefer, and T. Kesselheim. Jamming-resistant learning in wireless networks. In *Proc. 41st Intl. Coll. Automata, Languages and Programming (ICALP)*, volume 2, pages 447–458, 2014.
- [27] C. Daskalakis and V. Syrgkanis. Learning in auctions: Regret is hard, envy is easy. In *Proc. 57th Symp. Foundations of Computer Science (FOCS)*, pages 219–228, 2016.

-
- [28] N. Devanur, J. Morgenstern, V. Syrgkanis, and M. Weinberg. Simple auctions with simple strategies. In *Proc. 16th Conf. Econom. Comput. (EC)*, pages 305–322, 2015.
- [29] N. Dimitrov and G. Plaxton. Competitive weighted matching in transversal matroids. *Algorithmica*, 62(1-2):333–348, 2012.
- [30] M. Dinitz and G. Kortsarz. Matroid secretary for regular and decomposable matroids. *SIAM J. Comput.*, 43(5):1807–1830, 2014.
- [31] S. Dobzinski, H. Fu, and R. D. Kleinberg. On the complexity of computing an equilibrium in combinatorial auctions. In *Proc. 26th Symp. Discrete Algorithms (SODA)*, pages 110–122, 2015.
- [32] S. Dobzinski and R. P. Leme. Efficiency guarantees in auctions with budgets. In *Proc. 41st Intl. Coll. Autom. Lang. Program. (ICALP)*, pages 392–404, 2014.
- [33] S. Dughmi and Y. Peres. Mechanisms for risk averse agents, without loss. *CoRR*, abs/1206.2957, 2012.
- [34] P. Dütting, T. Kesselheim, and É. Tardos. Mechanism with unique learnable equilibria. In *Proc. 15th Conf. Econom. Comput. (EC)*, pages 877–894, 2014.
- [35] E. Dynkin. The optimum choice of the instant for stopping a Markov process. In *Sov. Math. Dokl*, volume 4, pages 627–629, 1963.
- [36] M. Feldman, H. Fu, N. Gravin, and B. Lucier. Simultaneous auctions are (almost) efficient. In *Proc. 45th Symp. Theory of Computing (STOC)*, pages 201–210, 2013.
- [37] M. Feldman, O. Svensson, and R. Zenklusen. A simple $O(\log \log(\text{rank}))$ -competitive algorithm for the matroid secretary problem. In *Proc. 26th Symp. Discrete Algorithms (SODA)*, pages 1189–1201, 2015.
- [38] M. Feldman and M. Tennenholtz. Interviewing secretaries in parallel. In *Proc. 13th Conf. Electronic Commerce (EC)*, pages 550–567, 2012.
- [39] M. Feldman and R. Zenklusen. The submodular secretary problem goes linear. In *Proc. 56th Symp. Foundations of Computer Science (FOCS)*, pages 486–505, 2015.
- [40] A. Fiat, I. Gorelik, H. Kaplan, and S. Novgorodov. The temp secretary problem. In *Proc. 23rd European Symp. Algorithms (ESA)*, pages 631–642, 2015.
- [41] A. Fiat and C. H. Papadimitriou. When the players are not expectation maximizers. In *Proc. 3rd Intl. Symp. Algorithmic Game Theory (SAGT)*, pages 1–14, 2010.
- [42] G. Fibich, A. Gaviious, and A. Sela. All-pay auctions with risk-averse players. *Int. J. Game Theory*, 34(4):583–599, 2006.
- [43] M. Fisher, G. Nemhauser, and L. Wolsey. An analysis of approximations for maximizing submodular set functions-II. In *Polyhedral combinatorics*, pages 73–87. Springer, 1978.

- [44] H. Fu, J. Hartline, and D. Hoy. Prior-independent auctions for risk-averse agents. In *Proc. 14th Conf. Electr. Commerce (EC)*, pages 471–488. ACM, 2013.
- [45] O. Göbel, M. Hoefer, T. Kesselheim, T. Schleiden, and B. Vöcking. Online independent set beyond the worst-case: Secretaries, prophets and periods. In *Proc. 41st Intl. Coll. Automata, Languages and Programming (ICALP)*, volume 2, pages 508–519, 2014.
- [46] M. Hajiaghayi, R. Kleinberg, and D. Parkes. Adaptive limited-supply online auctions. In *Proc. 5th Conf. Electronic Commerce (EC)*, pages 71–80, 2004.
- [47] G. Hardy, J. Littlewood, and G. Pólya. Inequalities. reprint of the 1952 edition. cambridge mathematical library, 1988.
- [48] J. D. Hartline, V. Syrgkanis, and É. Tardos. No-regret learning in repeated Bayesian games. *CoRR*, abs/1507.00418, 2015.
- [49] A. Hassidim, H. Kaplan, Y. Mansour, and N. Nisan. Non-price equilibria in markets of discrete goods. In *Proc. 12th Conf. Electr. Commerce (EC)*, pages 295–296, 2011.
- [50] M. Hoefer, T. Kesselheim, and B. Kodric. Smoothness for simultaneous composition of mechanisms with admission. In *Proc. 12th Intl. Conf. Web and Internet Economics (WINE)*, pages 294–308, 2016.
- [51] M. Hoefer, T. Kesselheim, and B. Vöcking. Truthfulness and stochastic dominance with monetary transfers. *ACM Trans. Econom. Comput.*, 4(2):11:1–11:18, 2016.
- [52] M. Hoefer and B. Kodric. Combinatorial secretary problems with ordinal information. In *Proc. 44th Intl. Coll. Autom. Lang. Program. (ICALP)*, pages 133:1–133:14, 2017.
- [53] D. Hoy. The concavity of atomic splittable congestion games with non-linear utility functions. In *EC 2012, Workshop on Risk Aversion in Algorithmic Game Theory and Mechanism Design*, 2012.
- [54] D. Hoy, N. Immorlica, and B. Lucier. On-demand or spot? selling the cloud to risk-averse customers. In *Proc. 12th Intl. Conf. Web and Internet Economics (WINE)*, pages 73–86, 2016.
- [55] A. Hu, S. A. Matthews, and L. Zou. Risk aversion and optimal reserve prices in first-and second-price auctions. *J. Econom. Theory*, 145(3):1188–1202, 2010.
- [56] P. Jaillet, J. Soto, and R. Zenklusen. Advances on matroid secretary problems: Free order model and laminar case. In *Proc. 16th Intl. Conf. Integer Programming and Combinatorial Optimization (IPCO)*, pages 254–265, 2013.
- [57] T. Kesselheim and B. Kodric. Price of anarchy for mechanisms with risk-averse agents. In *Proc. 45th Intl. Coll. Autom. Lang. Program. (ICALP)*, pages 155:1–155:14, 2018.
- [58] T. Kesselheim, K. Radke, A. Tönnis, and B. Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *Proc. 21st European Symp. Algorithms (ESA)*, pages 589–600, 2013.

-
- [59] T. Kesselheim, K. Radke, A. Tönnis, and B. Vöcking. Primal beats dual on online packing LPs in the random-order model. In *Proc. 46th Symp. Theory of Computing (STOC)*, pages 303–312, 2014.
- [60] T. Kesselheim and A. Tönnis. Submodular secretary problems: Cardinality, matching, and linear constraints. In *Proc. 20th Workshop Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 16:1–16:22, 2017.
- [61] R. Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *Proc. 16th Symp. Discrete Algorithms (SODA)*, pages 630–631, 2005.
- [62] B. S. Klose and P. Schweinzer. Auctioning risk: The all-pay auction under mean-variance preferences. 2014.
- [63] N. Korula and M. Pál. Algorithms for secretary problems on graphs and hypergraphs. In *Proc. 36th Intl. Coll. Automata, Languages and Programming (ICALP)*, pages 508–520, 2009.
- [64] O. Lachish. $O(\log \log \text{rank})$ competitive ratio for the matroid secretary problem. In *Proc. 55th Symp. Foundations of Computer Science (FOCS)*, pages 326–335, 2014.
- [65] B. Lucier and A. Borodin. Price of anarchy for greedy auctions. In *Proc. 21st Symp. Discr. Algorithms (SODA)*, pages 537–553, 2010.
- [66] T. Lykouris, V. Syrgkanis, and É. Tardos. Learning and efficiency in games with dynamic population. In *Proc. 27th Symp. Discrete Algorithms (SODA)*, pages 120–129, 2016.
- [67] H. M. Markowitz. *Portfolio selection: efficient diversification of investments*, volume 16. Yale university press, 1968.
- [68] H. M. Markowitz. Mean-variance approximations to expected utility. *Europ. J. Oper. Res.*, 234(2):346–355, 2014.
- [69] A. Mas-Colell, M. D. Whinston, J. R. Green, et al. *Microeconomic theory*, volume 1. Oxford university press New York, 1995.
- [70] E. Maskin and J. Riley. Optimal auctions with risk averse buyers. *Econometrica*, pages 1473–1518, 1984.
- [71] T. Mathews and B. Katzman. The role of varying risk attitudes in an auction with a buyout option. *J. Econom. Theory*, 27(3):597–613, 2006.
- [72] S. A. Matthews. Selling to risk averse buyers with unobservable tastes. *J. Econom. Theory*, 30(2):370–400, 1983.
- [73] R. Meir and D. C. Parkes. Playing the wrong game: Smoothness bounds for congestion games with behavioral biases. *SIGMETRICS Performance Evaluation Review*, 43(3):67–70, 2015.
- [74] R. Myerson. Optimal auction design. *Math. Oper. Res.*, 6:58–73, 1981.

- [75] E. Nikolova and N. E. S. Moses. A mean-risk model for the traffic assignment problem with stochastic travel times. *Oper. Res.*, 62(2):366–382, 2014.
- [76] E. Nikolova and N. E. S. Moses. The burden of risk aversion in mean-risk selfish routing. In *Proc. 16th Conf. Econom. Comput. (EC)*, pages 489–506, 2015.
- [77] R. Paes Leme and É. Tardos. Pure and bayes-nash price of anarchy for generalized second price auction. In *Proc. 51st Symp. Foundations of Computer Science (FOCS)*, pages 735–744, 2010.
- [78] G. Piliouras, E. Nikolova, and J. S. Shamma. Risk sensitivity of price of anarchy under uncertainty. *ACM Trans. Econom. Comput.*, 5(1):5:1–5:27, 2016.
- [79] T. Roughgarden. The price of anarchy in games of incomplete information. In *Proc. 13th Conf. Electr. Commerce (EC)*, pages 862–879, 2012.
- [80] T. Roughgarden. Barriers to near-optimal equilibria. In *Proc. 55th Symp. Foundations of Computer Science (FOCS)*, pages 71–80, 2014.
- [81] T. Roughgarden. Intrinsic robustness of the price of anarchy. *J. ACM*, 62(5):32, 2015.
- [82] A. Rubinstein. Beyond matroids: secretary problem and prophet inequality with general constraints. In *Proc. 48th Symp. Theory of Computing (STOC)*, pages 324–332, 2016.
- [83] J. Soto. Matroid secretary problem in the random-assignment model. *SIAM J. Comput.*, 42(1):178–211, 2013.
- [84] M. Sundararajan and Q. Yan. Robust mechanisms for risk-averse sellers. In *Proc. 11th Conf. Econom. Comput. (EC)*, pages 139–148, 2010.
- [85] V. Syrgkanis and É. Tardos. Composable and efficient mechanisms. In *Proc. 45th Symp. Theory of Computing (STOC)*, pages 211–220, 2013.
- [86] J. von Neumann and O. Morgenstern. *Theory of games and economic behavior*. Princeton, New Jersey: Princeton University Press. XVIII, 625 p., 1944.
- [87] Q. Yan. Mechanism design via correlation gap. In *Proc. 22nd Symp. Discr. Algorithms (SODA)*, pages 710–719, 2011.
- [88] A. C. Yao. Probabilistic computations: Toward a unified measure of complexity (extended abstract). In *Proc. 18th Symp. Foundations of Computer Science (FOCS)*, pages 222–227, 1977.