

HABILITATION THESIS

**Primal-Dual Approaches in Online
Algorithms, Algorithmic Game Theory
and Online Learning**

Author:
NGUYỄN KIM Thắng

June 7th 2019

“Mathematics is the art of giving the same name to different things.”

Henri Poincaré

“If people do not believe that mathematics is simple, it is only because they do not realize how complicated life is.”

John von Neumann

Abstract

Habilitation

Primal-Dual Approaches in Online Algorithms, Algorithmic Game Theory and Online Learning

by NGUYỄN KIM THẮNG

Primal-dual is an elegant and powerful method in optimization and in the design of algorithms. The main idea of the method is to construct feasible primal and dual solutions interactively and an algorithm, together with the analysis, are derived naturally from the primal-dual interaction. In this thesis, we present primal-dual approaches as unified techniques in order to study and build connections between the domains of Online Algorithms, Algorithmic Game Theory and Online Learning.

Primal-duale est une méthode élégante et puissante en optimisation et en algorithmique. La méthode consiste à établir de manière interactive des solutions primales et duales, puis un algorithme, ainsi que son analyse, sont guidés naturellement par l'interaction primal-duale. Dans cette habilitation, nous présentons les approches primal-duales comme techniques unifiées afin d'étudier et de développer des liens entre les domaines de l'algorithmique en ligne, de la théorie des jeux algorithmiques et de l'apprentissage en ligne.

Primal-dual là một phương pháp đẹp và hiệu quả trong tối ưu và trong thiết kế các thuật toán. Ý tưởng chính của phương pháp là xây dựng các nghiệm đối ngẫu và sau đó, thuật toán, cùng với phân tích, sẽ dần dần hình thành từ những tương tác đối ngẫu. Trong luận án này, chúng tôi trình bày các phương pháp đối ngẫu như những kỹ thuật thống nhất để nghiên cứu và kết nối các lĩnh vực: Thuật toán trực tuyến, Lý thuyết trò chơi thuật toán và Học máy trực tuyến.

Acknowledgements

Family is the most important part in everyone's life and I am not an exception. This thesis is dedicated to my parents, my wife Hương and our sons Tiên and Minh. They are a source of inspiration for me.

A special thank to Christoph Dürr. Christoph was my PhD advisor and we become close colleague and friend. By his simplicity, Christoph implicitly teaches me a lot and discussions with him (scientific and non-scientific) are always lively.

Research cannot be carry on without collaborators/friends. I am happy to be collaborating with many colleagues who not only are experts in their domains but also have excellent human quality. This thesis would have been impossible without discussions/exchanging with them. Thank all of you: Eric Angel, Spyros Angelopoulos, Evaripidis Bampis, Nikhil Bansal, Vincent Chau, Johanne Cohen, Vincent Cohen-Addad, Marek Chrobak, Carola Doerr, Bruno Escoffier, Laurent Gourvès, Chien-Chung Huang, Giuseppe F. Italiano, Lukasz Jez, Claire Mathieu, Jérôme Monnot, Stefano Moretti, Giorgio Lucarelli, Yannis Manoussakis, Benjamin Mosely, Nguyễn Việt Hưng, Fanny Pascual, Phạm Hồng Phong, Damien Regnault, Olivier Spanjaard, Shikha Singh, Abhinav Srivastav, Denis Trystram, Oscar C. Vasquez.

Contents

Abstract	iii
Acknowledgements	v
1 Introduction	1
1.1 The Primal-Dual Method in Algorithm Design and in Machine Learning	1
1.2 Research Directions	2
1.3 Primal-Dual Approach	4
1.3.1 Configuration Linear Programs	4
1.3.2 Smoothness and concavity properties	5
1.3.3 Primal-dual framework	9
1.4 Contributions to Online Algorithms	10
1.4.1 Notations and Preliminaries	10
1.4.2 Primal-Dual Algorithms for A General Class	10
1.4.3 Primal-Dual Approach for 0 – 1 Covering Problems	12
1.4.4 Primal-Dual Approach for 0 – 1 Packing Problems	13
1.4.5 Related work	14
1.5 Contributions to Algorithmic Game Theory	14
1.5.1 Notations and Preliminaries	15
1.5.2 Games in Full-Information Settings	17
1.5.3 Games in Incomplete-Information Settings	18
1.5.4 Related works	20
1.6 Contributions to Efficient Online Learning Algorithms	21
1.6.1 Definitions and Framework	22
1.6.2 Application to Fictitious Play in Smooth Auctions	22
1.6.3 Application to Revenue Maximization in Multi-Dimensional Environments	23
1.6.4 Related Work	24
2 Online Primal-Dual Algorithms with Configuration Linear Programs	27
2.1 Primal-Dual Algorithm for A General Class of Problems	27
2.1.1 A Greedy Algorithm	27
2.1.2 Applications	29
2.1.2.1 Minimum Power Survival Network Routing	30
2.1.2.2 Online Vector Scheduling	30
2.1.2.3 Online Energy-Efficient Scheduling	32
2.1.2.4 Online Prize Collecting Energy-Efficient Scheduling	33
2.1.2.5 Facility Location with Client-Dependent Facility Cost	35
2.2 Primal-Dual Framework for 0 – 1 Covering Problems	37
2.2.1 Algorithm for Fractional Covering	37
2.2.2 Applications	41
2.2.2.1 Polynomials with non-negative coefficients.	42
2.2.2.2 Beyond convex functions.	42

2.2.2.3	Submodular functions.	43
2.3	Primal-Dual Framework for Packing Problems	44
2.3.1	Algorithm for Fractional Packing	44
2.3.2	Applications to online submodular maximization	47
3	Game Efficiency through Linear Programming Duality	49
3.1	Smooth Games under the Lens of Duality	49
3.2	Congestion Games	51
3.2.1	Atomic Congestion Games	51
3.2.2	Nonatomic Congestion Games	53
3.2.2.1	Efficiency of Non-Atomic Congestion Games	54
3.2.2.2	Resource Augmentation in Non-Atomic Congestion Games	56
3.2.3	Splittable Congestion Games	57
3.3	Efficiency of Smooth Auctions in Welfare Maximization	60
3.4	Simultaneous Item-Bidding Auctions	62
3.4.1	Connection between Primal-Dual and Non-Smooth Techniques	63
3.4.2	Connection between Primal-Dual and No-Envy Learning	65
3.5	Sequential Auctions	68
3.5.1	Sequential Second Price Auctions in Sponsored Search	68
3.5.2	Sequential First Price Auctions in Matching Markets	72
4	Efficient Online Learning Algorithms and Auction Design	77
4.1	Framework of Online Learning	77
4.1.1	Regret of (λ, μ) -Concave Functions	77
4.1.2	Multilinear Extension of Discretization	79
4.2	Applications to Fictitious Play in Smooth Auctions	81
4.3	Online Simultaneous Second-Price Auctions with Reserve Prices	84
5	Conclusion	89
A	Smoothness Parameters	91
A.1	Smoothness Parameters of Polynomials	91
B	Concavity Parameters	93
B.1	Concavity Parameters of Multilinear Extension of Submodular Functions	93

To my family: bố, mẹ, Hương, Tiến và Minh.

Chapter 1

Introduction

Research has its own beauty: the beauty of diversity and also of unicity. A research domain witnesses continually the beauty in its development: interesting notions, results have been flourished at the early stage (stage of diversity), then deep methods, techniques have been established at a more mature stage (stage of unicity). These methods/techniques lead to further achievements, connections to other domains (higher stage of diversity) and so on. Beauty lies in the eyes of the beholder. My research, following this beauty perception, is the quest of the *diversity*, the *unicity* and the *simplicity* in algorithms.

Among others, the algorithmic methods based on duality [33] have particularly attracted by their conceptual simplicity, their wide applicability (diversity) and the fact that they provide a general recipe (unicity) for classes of problems. My research is devoted to the development of primal-dual methods in Online Algorithms, Algorithmic Game Theory and Online Learning and to explore connections between these domains using the primal-dual approach as a versatile tool.

1.1 The Primal-Dual Method in Algorithm Design and in Machine Learning

The primal-dual method has been widely used in the design and analysis of algorithms. The main idea of the method is to construct feasible primal and dual solutions interactively and an algorithm is naturally derived from the primal-dual interaction. The method has several advantages: (i) it provides conceptually simple algorithms; (ii) the designed algorithms are usually faster than the ones relying on directly solving linear/mathematical programs; and (iii) the analysis of the algorithm performance interactively comes with the primal-dual construction.

Combinatorial Optimization. Primal-dual algorithms have been designed for various problems such as linear programming problems, network flow problems, shortest $s - t$ path problems, and many others. Primal-dual algorithms start with a dual feasible solution and use dual information to infer a primal, possibly infeasible, solution. If the primal solution is infeasible, the dual solution is modified and so on. For example, the shortest $s - t$ path problem can be solved by an algorithm that greedily increases dual variables until the corresponding primal solution becomes feasible [76]. For an overview of the primal-dual method in combinatorial optimization, one can see the book of Papadimitriou and Steiglitz [104].

Approximation Algorithms. The primal-dual method for approximation algorithms is a generalization of the primal-dual method used for linear programming and combinatorial optimization problems. The first use of the primal-dual method in approximation algorithms is due to Bar-Yehuda and Even [19] who gave an algorithm for the vertex cover problem. Primal-dual approximation algorithms were aroused by work on the generalized Steiner tree problem, especially the 2-approximation primal-dual algorithm of Agrawal et al. [1]. The use of linear programming and LP duality was made explicit by Goemans and Williamson [64]. Subsequently, the primal-dual method has been extensively studied for approximation algorithms (surveys in [127, 126]).

Online Algorithms. The first explicit primal-dual approach in online algorithm is given by Buchbinder and Naor [34] who presented general algorithms based on the multiplicative weights update method for covering/packing problems. Work on primal-dual online algorithms was revived by Buchbinder and Naor's framework. Several uses of the primal-dual method and competitive algorithms for online algorithms then followed; for example the generalized caching problem [17, 16], the ad-auction maximization [35] and others. We refer the reader to the survey [32]. Recently, Azar et al. [11] have generalized Buchbinder and Naor's framework for a class of convex functions with monotone gradient.

Online Learning (Online Convex Optimization). The primal-dual method has been used in machine learning in different forms. The multiplicative weights update method, which can be viewed as a primal-dual method (by Buchbinder and Naor [32]), is indeed a ubiquitous meta-algorithm in computation and learning [7]. This method has first been used in online learning by Littlestone and Warmuth [91] to derive the Weighted Majority algorithm for the problem of prediction from expert advice. Besides, the multiplicative weights update method belongs to the class of first-order methods in optimization.

Among the first-order methods, mirror descent is the one which explicitly uses the interaction between solutions in primal and dual spaces. Hazan and Kale [72] showed that the mirror descent method is equivalent to the *Follow-the-Regularized-Leader* (FTRL), an algorithm introduced by Shalev-Shwartz and Singer [115] for the regret minimization problem.

1.2 Research Directions

Online Primal-Dual Algorithms for Non-Convex Problems. One real-world phenomenon, known as the *economy of scale*, consists in sub-linear growth of cost as a function of the amount of used resources. This happens in many scenarios in which one gets a discount when buying resources in bulk. A representative setting is the extensively-studied domain of sub-modular optimization. Another phenomenon, known as the *diseconomy of scale*, is that the cost grows super-linearly in the quantity of resources used. An illustrative example for this phenomenon is the energy cost of computation where the cost grows super-linearly, typically as a convex function. The diseconomy of scale has been widely studied in the domain of convex optimization [29]. However, in many settings, the costs are the mix of both phenomena and the objective functions are indeed non-convex. Non-convex objective functions appear in various problems in both theory and practice, ranging from scheduling, sensor energy management, to influence and revenue maximization, and facility location.

Such problems call for the design of algorithms with performance guarantee for non-convex objective functions.

Convex objectives have been extensively studied in recent years where the convexity was crucial for the analyses. As mentioned earlier, Azar et al. [11] have generalized the Buchbinder and Naor’s framework for a class of convex functions with monotone gradient. However, problems with non-convex objectives resist against current approaches and non-convexity represents a strong barrier in optimization in general and in online algorithms in particular. Hence, designing competitive algorithms for non-convex problems represents an interesting and important challenge.

Primal-Dual Approach in Algorithmic Game Theory. Algorithmic Game Theory — a domain at the intersection of Game Theory and Algorithms — has been extensively studied in the last two decades. In a game, the price of anarchy (PoA) [84] is defined as the worst ratio between the cost of a Nash equilibrium and that of an optimal solution. The PoA is now considered as standard and is the most popular measure to characterize the inefficiency of Nash equilibria — solutions to games — in the same spirit of the approximation ratio in Approximation Algorithms (the price of being restricted to polynomial running time) and the competitive ratio in Online Algorithms (the price of limited knowledge about the future).

Mathematical programming in general and linear programming in particular are powerful tools in many research fields. Among others, linear programming has a tremendous impact on the design of algorithms. Linear programming and duality play crucial and fundamental roles in several elegant methods such as the primal-dual and the dual-fitting ones for Approximation Algorithms [127] and online primal-dual framework [32] in Online Algorithms. Given the tremendous impact of tools from mathematical programming in the design of algorithms and the similarity of the notions of PoA, approximation and competitive ratios, it is intriguing and also desirable to develop a framework based on duality to study the efficiency of games.

Connections between Online Algorithms, Online Learning and Algorithmic Game Theory. The success of machine learning in a wide range of applications opens many exciting directions, in particular exploring the interaction between Algorithms and Machine Learning. While primal-dual is an elegant tool developed in algorithmic community, the mirror descent approach is widely studied in learning community. Primal-dual methods in Online Algorithms and mirror descent methods in Online Learning are recently revealed to be the same thing but have been viewed by different lenses in different communities. Recently, Buchbinder et al. [39] have shown an unified framework for the methods and its connection by the mean of regularization technique [37]. Furthermore, Bubeck et al. [31] have achieved a breakthrough (on the k -server problem) in Online Algorithms using powerful tools in Online Learning. Exploring the interaction between Online Algorithms and Online Learning is a promising research direction.

A challenge in Online Learning is the design of efficient (polynomial time) algorithms with performance guarantees in adversarial non-stationary, non-stochastic environments by applying optimization methods that learn from experience and observations. Hazan and Koren [74] have proved that designing such algorithms is not possible in general adversarial environments. However, efficient online learning

may be achievable in well-structured settings with regularity conditions. Characterizing conditions, or in general discovering the hidden regularity, under which efficient online learning algorithms exist is a major research agenda in online learning. Among others, studying learning dynamics and designing efficient learning algorithms in games are of particular interests as described in the book of Cesa-Bianchi and Lugosi [42], which served as an inspiration to the entire field of learning in games, and the current work in data-driven mechanism design [52, 113].

1.3 Primal-Dual Approach

In this section, we present the key notions and our approach on the development of primal-dual methods along these directions.

1.3.1 Configuration Linear Programs

Linear program based techniques are powerful tools in theoretical computer science in general and in approximation and online algorithms in particular. Given an optimization problem, one first attempt consists in formalizing the problem as an integer/linear program. The most common and natural way to construct a formulation for a problem is to consider the optimization procedure locally. That is, variables represent local decisions of an algorithm (for example, if some object is selected, if some edge in a graph is used, ect) and the aggregation of these decisions (variables) through constraints of the formulation forms a solution. Another way to derive a formulation for a problem is to see the optimization procedure globally. Now, one considers all feasible solutions and each variable, associated to a solution, represents a global decision of an algorithm. In other words, variables indicate which solution is chosen in order to optimize the objective of the problem. The latter formulation, if it is linear, is called *configuration linear program* (LP).

The first crucial step for any LP-based approach, including the primal-dual method, is to derive an LP formulation with reasonably small *integrality gap*. The latter is defined as the worst ratio between the optimal integer solution of the formulation and the optimal solution with the integrality conditions relaxed. The quest for formulations with small integrality gap and for techniques to reduce this gap is a major direction in the area of approximation algorithms. Different techniques have been successfully developed: the lift-and-project methods (Sum-of-Squares hierarchy, LP-hierarchies), knapsack-inequalities, etc. For this purpose, in our approach we consider configuration LPs constructed as follows.

First consider a natural (local) formulation of a given optimization problem. As we are interested in non-linear cost functions, it is not surprising that the natural relaxation suffers from large integrality gap. Then, introduce an exponential number of new variables that represent global decisions of an algorithm and constraints that make connection between new variables (in global view) to original variables (in local view). Specifically, the new variables represent which outcome of the problem is selected. Moreover, the new constraints guarantee that (1) the problem admits exactly one outcome; and (2) if a local decision is made then this local decision must be a component of the outcome. The new LP captures both local and global decisions of an algorithm and it indeed reduces substantially the integrality gap. The construction of configuration LP has been observed and presented by Makarychev

and Sviridenko [94] where they study the problem of minimizing the objective of form $f(u) = u^\alpha$ for constant α .

Example 1 (Network Routing) Consider a network $G(V, E)$ and a set of k routing requests (s_i, t_i) for $1 \leq i \leq k$. The cost $f_e : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ of a link e is $f_e(u)$ where u is the quantity of flow passing through the link. The objective is to route integrally one unit flow from s_i to t_i for every i with minimum total cost.

We illustrate the issue of the integrality gap even on a very simple setting. Assume that the network consists of two nodes s, t and m parallel links connecting s and t . The cost of each link e is $f_e(u) = u^2$ for every e where u is the quantity of flow passing through link e . The goal is to route integrally one unit from s to t with minimum cost. The optimal integer solution has cost 1 by routing one unit flow through an arbitrary link. However, the following natural relaxation has optimal fractional solution of cost $m \cdot (1/m)^2 = 1/m$ which leads to an integrality gap of m .

$$\min \sum_{e=1}^m u_e^2 \quad \text{s.t.} \quad \sum_{e=1}^m u_e = 1, \quad u_e \geq 0 \quad \forall e.$$

Now we show the construction of configuration LPs from the natural LP (for the general setting). Let \mathcal{S}_i be the set of paths connection s_i to t_i . Let x_{ij} be a 0-1 variable such that $x_{ij} = 1$ if one routes the unit flow from s_i to t_i through the path $P_{ij} \in \mathcal{S}_i$. Besides, for every subset $A \subseteq \{1, 2, \dots, k\}$ and for every link e , let $z_{eA} = 1$ if the set of requests using link e is exactly A (and $z_{eA} = 0$ otherwise).

$$\begin{aligned} \min \quad & \sum_{e,A} f_e(|A|) z_{e,A} \\ \sum_{j: P_{ij} \in \mathcal{S}_i} x_{ij} &= 1 & \forall i \\ \sum_A z_{e,A} &= 1 & \forall e \\ \sum_{A: i \in A} z_{e,A} &= \sum_{j: e \in P_{ij}} x_{ij} & \forall i, e \\ x_{ij}, z_{e,A} &\in \{0, 1\} & \forall i, j, e, A \end{aligned}$$

The first constraint ensures that every request i has to be routed by some path from s_i to t_i . This is a constraint in the natural relaxation. The other constraints relate (local) variables x to (global) variables z . The second constraint stands for the fact that for every link e , there is a subset of requests, possibly empty, using e . The third constraint means that if a link e is used by the request i (by choosing some path P_{ij} such that $e \in P_{ij}$) then the set of requests using link e must contain i . We will see later that the configuration LP, obtained by relaxing the integral constraint of x and z , has significantly smaller integrality gap compared to the natural one and for a large class of cost functions, it gives the optimal relaxation.

1.3.2 Smoothness and concavity properties

In our approach, we characterize the performance of algorithms based on properties of the cost functions, called *smoothness*. The notion together with its name are

inspired by the smoothness framework introduced by Roughgarden [108] in the context of algorithmic game theory in order to characterize the price of anarchy for large classes of games. We first recall the smoothness definition of Roughgarden [108] in game theory.

Consider a game with n players in which each player i selects a strategy s_i from a set \mathcal{S}_i for $1 \leq i \leq n$. These strategies form a *strategy profile* $\mathbf{s} = (s_1, \dots, s_n)$. The cost $C_i(\mathbf{s})$ of player i is a real function of the strategy profile \mathbf{s} . A game with a joint cost objective function $C(\mathbf{s}) = \sum_{i=1}^n C_i(\mathbf{s})$ is (λ, μ) -smooth if for every two outcomes \mathbf{s} and \mathbf{s}^* , it holds that

$$\sum_{i=1}^n C_i(\mathbf{s}_i^*, \mathbf{s}_{-i}) \leq \lambda \cdot C(\mathbf{s}^*) + \mu \cdot C(\mathbf{s})$$

The price of anarchy of a game can be bounded, using the smoothness notion, by the following term

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

which is proved to be optimal for many games (for example, the celebrated congestion games) [108]. Rapidly, the smoothness framework became popular and has been generalized and widely used to bound the price of anarchy of games in complete and incomplete information environments [109, 119].

Inspired by the generality and applicability of the smoothness framework in algorithmic game theory, we define properties of cost functions which will be used to characterize the performance of our algorithms. Through these notions, we show an interesting connection between online algorithms and algorithmic game theory.

Smoothness Property

We start with a notion of *smoothness* of functions, which is the closest to the definition in [108].

Definition 1.1 Let \mathcal{N} be a ground set of elements. A set function $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$ is (λ, μ) -smooth if for any set $A = \{a_1, \dots, a_n\} \subseteq \mathcal{N}$ and any collection $B_1 \subseteq B_2 \subseteq \dots \subseteq B_n \subseteq B \subseteq \mathcal{N}$, the following inequality holds.

$$\sum_{i=1}^n [f(B_i \cup a_i) - f(B_i)] \leq \lambda f(A) + \mu f(B)$$

A set of cost functions $\{f_e : e \in \mathcal{E}\}$ is (λ, μ) -smooth if every function f_e is (λ, μ) -smooth.

Let us explain intuitively the meaning of the definition. Imagine that B_i is the current solution of an algorithm at step i . Then, the inequality means the following: if the total marginal increase by following the choice/strategy a_i at step i (the left-hand side) can be bounded by a combination of the algorithm cost (the second term of the right-hand side) and the cost of an adversary (potentially the set A of all strategies a_i 's), then the algorithm is competitive. As we will see later, the competitive ratio will be $\lambda/(1 - \mu)$. Indeed, given a (λ, μ) -smooth function, the quantity $\frac{\lambda}{1 - \mu}$ informally measures how far the function is from being linear.

Example 2

1. If \mathcal{N} consists of real numbers and f is an affine linear function ($f(A) = \sum_{a \in A} a$) then it is $(1, 0)$ -smooth since

$$\sum_{i=1}^n [f(B_i \cup a_i) - f(B_i)] = \sum_{a \in A} a = 1 \cdot f(A) + 0 \cdot f(B)$$

2. If \mathcal{N} consists of real numbers and $f(A) = g(\sum_{a \in A} a)$ where g is a polynomial of degree k with non-negative coefficients, then f is $(\Theta(k^{k-1}), \frac{k-1}{k})$ -smooth. That is due to the following inequality proved in [45].

$$\sum_{i=1}^n \left[\left(b_i + \sum_{j=1}^i a_j \right)^k - \left(\sum_{j=1}^i a_j \right)^k \right] \leq \Theta(k^{k-1}) \cdot \left(\sum_{i=1}^n b_i \right)^k + \frac{k-1}{k} \cdot \left(\sum_{i=1}^n a_i \right)^k$$

Local Smoothness Property

We extend the notions of smoothness in order to design algorithms in more complex settings. Given two vectors \mathbf{x} and \mathbf{y} in $[0, 1]^n$, denote $\mathbf{x} \vee \mathbf{y}$ the vector such that its component at coordinate $1 \leq i \leq n$ is $\max\{x_i, y_i\}$. For minimization problems, we consider the following notion of *min-locally-smooth*.

Definition 1.2 A differentiable function $F : [0, 1]^n \rightarrow \mathbb{R}^+$ is (λ, μ) -min-locally-smooth if for any set $S \subseteq \{1, \dots, n\}$, and for all vectors $\mathbf{x}^e \in [0, 1]^n$ with $1 \leq e \leq n$, the following inequality holds.

$$\sum_{e \in S} \nabla_e F(\mathbf{x}^e) \leq \lambda F(\mathbf{1}_S) + \mu F(\mathbf{x}) \quad (1.1)$$

where $\mathbf{x} := \vee_{e \in S} \mathbf{x}^e$, meaning that $x_i = \max_e \{x_i^e\}$ for every coordinate $1 \leq i \leq n$.

If the gradient $\nabla F(\mathbf{x})$ is non-decreasing, we only need a simpler version. We say that a differentiable function $F : [0, 1]^n \rightarrow \mathbb{R}^+$ with *monotone gradient* is (λ, μ) -min-locally-smooth if for any set $S \subseteq \{1, \dots, n\}$, and for any vector $\mathbf{x} \in [0, 1]^n$, the following inequality holds.

$$\sum_{e \in S} \nabla_e F(\mathbf{x}) \leq \lambda F(\mathbf{1}_S) + \mu F(\mathbf{x}) \quad (1.2)$$

The local smoothness has similar meaning as the standard smoothness notion. Considering Inequality (1.2) and imagine that \mathbf{x} is the current solution of an algorithm. Then, if the *local* increase of the objective function F (the left-hand side) at the current solution in any direction (including the direction chosen by an adversary) can be bounded by a combination of the current cost (the second term of the right-hand side) and the cost of an adversary (the first term of the right-hand side), then the algorithm is competitive. The competitive ratio will be determined as a function of λ and μ .

An advantage of min-local-smoothness over smoothness is that the former allows us to analyze the algorithm performance *locally*, which in several settings, captures precisely the nature of the corresponding problem.

Example 3

1. If $F(A) = g(\sum_{a \in A} a)$ where g is a polynomial of degree k with non-negative coefficients, then F is also $(\Theta(k^{k-1}), \frac{k-1}{k})$ -min-locally-smooth.

2. If F is the multilinear extension of a submodular function f then F is $(\frac{1}{1-\kappa}, 0)$ -min-locally-smooth (proof can be found in Proposition 2.8) where κ is the *total curvature* [47] of f over an universe \mathcal{E} is defined as

$$\kappa_f = 1 - \min_{e \in \mathcal{E}} \frac{f(\mathbf{1}_{\mathcal{E}}) - f(\mathbf{1}_{\mathcal{E} \setminus \{e\}})}{f(\mathbf{1}_{\{e\}})}$$

For maximization problems, we introduce another appropriate notion of local-smoothness.

Definition 1.3 A differentiable function $F : [0, 1]^n \rightarrow \mathbb{R}^+$ is (λ, μ) -max-locally-smooth if for any set $S \subset \mathcal{E}$, and for any vectors $\mathbf{x}^e \in [0, 1]^n$, the following inequality holds:

$$\sum_{e \in S} \nabla_e F(\mathbf{x}^e) \geq \lambda F(\mathbf{1}_S) - \mu F(\mathbf{x})$$

where $\mathbf{x} := \bigvee_{e \in S} \mathbf{x}^e$, meaning that $x_i = \max_e \{x_i^e\}$ for any coordinate $1 \leq i \leq n$.

This definition is different to the min-local-smoothness. On one hand, it is due to different natures of minimization and maximization problems. On the other hand, in non-convex problems only weak duality holds while strong duality does not. So informally, there is no symmetry between primal and dual. Specifically, in linear programming, the dual of the dual is the primal while this property does not hold in a non-convex setting.

Example 4 Let F be the multilinear extension of a submodular function f .

1. If f is monotone then F is $(1, 1)$ -max-locally-smooth.
2. If f is monotone then F is $(1/3, 1)$ -max-locally smooth.

The proof can be found in Lemma 2.8.

Concavity Property

As mentioned earlier in Section 1.2, characterizing conditions under which efficient online learning algorithms exist is a major research agenda in online learning. In this direction, we introduce a regularity condition which is crucial in our approach in order to design efficient online learning algorithm and in analyze dynamics of games. The regularity notion generalizes the standard notion of concavity and it also has the flavour of the smoothness concept.

Definition 1.4 A function F is (λ, μ) -concave if for all vectors \mathbf{x} and \mathbf{x}^* ,

$$\langle \nabla F(\mathbf{x}), \mathbf{x}^* - \mathbf{x} \rangle \geq \lambda F(\mathbf{x}^*) - \mu F(\mathbf{x}).$$

Example 5

1. If F is concave then it is $(1, 1)$ -concave (according to the definition above).
2. If F is the multilinear extension of a monotone submodular function then F is $(1, 2)$ -concave (Lemma B.1).

1.3.3 Primal-dual framework

The primal-dual method typically consists of three steps:

1. formulation of the given problem as a mathematical program,
2. construction of the primal and dual variables,
3. proving primal-dual feasibility and bounding the primal/dual objectives.

In the following, we describe these steps in our framework.

Formulations

In our approach, we consider systematically the configuration LPs constructed from natural relaxations of given problems in order to reduce the integrality gap. To the best of our knowledge, the configuration LPs have been used only in the offline setting and the approach is to round an optimal fractional solution to an integer one and to bound the approximation ratio. The first encountered difficulty of this approach is that a configuration LP has exponential size, so one has to look for a separating oracle (for the dual LP) in order to compute an optimal fractional solution. Finding separating oracles is in general far from trivial and represents an obstacle in using configuration LPs to design performant algorithms. Moreover, solving an LP is intrinsically offline and it is not suitable for studying problems in online settings.

Primal-dual has several advantages in order to bypass this difficulty. First, one does not have to compute an optimal fractional solution. Second, the configuration formulation admits a *sparsity* property: the number of variables is exponential but for an integer solution, only one configuration variable equals 1 (that corresponds to the chosen solution of an algorithm) and all others equals 0. Intuitively, this sparsity property explains why primal-dual is appropriate to study configuration LP. Formally, although there are exponentially many variables in a configuration LP, at any time in our primal-dual framework, one needs to maintain only a feasible primal solution and a small number of meaningful dual variables with non-zero values.

Construction of Primal-Dual Variables

In the direction of designing competitive algorithms for non-convex problems, we consider a primal-dual updating procedure which is inspired by the frameworks of Buchbinder and Naor [33] for linear objectives and Azar et al. [11] for convex objectives with monotone gradients. The procedure is essentially the multiplicative weights update. However, the crucial distinguishing point of our algorithm compared to the ones in [11] is that they use the gradient of the objective function at the current primal solution to define a multiplicative update whereas our update is guided by the gradient of the multilinear extension of the objective function. This multiplicative update, together with the configuration LPs and the notions of smoothness, enable us to derive competitive algorithms for convex objective functions whose gradients are not necessarily monotone and more generally, for non-convex objectives.

In the direction of analyzing the price of anarchy, the dual variables of configuration LPs have a surprising connection with the concept of Nash equilibria in games. Specifically, if dual variables are defined as the cost of players and the social cost in a game (up to some constant factors) then the dual constraints represent *exactly* the definitions of Nash equilibria and of smooth games. Indeed, this observation is the

starting point to the introduction of the smoothness notions in Section 1.3.2 in order to analyze the algorithm performances. Based on this observation, the construction of dual variables in our framework is guided by economic concepts. This allows us to derive other interesting connections and improved PoA bounds in games.

In the direction of designing efficient online learning algorithms, the updating procedure is carried out by the mirror descent method. We consider a discretization of the objective function and the multilinear extension of this discretization. Then, building on salient ideas in the above directions, we update the solutions in the direction of the gradient of this multilinear extension. The discretization guarantees the efficiency of algorithms and the multilinear-extension-based updates have several useful properties which can be exploited using the machinery developed in our framework.

Feasibility and Performance Guarantee

The notions of smoothness/concavity are particularly useful in our approach. Specifically, the dual feasibility is proved naturally using these notions. The performance of an algorithm and the efficiency of a game are then characterized by the smoothness/concavity parameters of cost functions. The use of the smoothness/concavity notions have several advantages: (i) it avoids the cumbersome technical details in the analysis as well as in the assumptions of objective functions; (ii) it reduces the analysis of bounding the competitive ratios to determining the smoothness/concavity parameters.

1.4 Contributions to Online Algorithms

We systematically use the primal-dual framework to design competitive algorithms for problems with non-convex objective. We first illustrate our approach by considering a model which captures several well-studied problems. Then, we study general online problems with covering and packing constraints.

1.4.1 Notations and Preliminaries

In an online problem, requests arrive over time and at any time an algorithm needs to make an irrevocable decision to satisfy the current request without the full knowledge about the future. An algorithm is *r-competitive* if for any request sequence, the algorithm cost (or gain) is at most (at least, resp.) r times that of the optimal solution (which has the full knowledge on the request sequence).

1.4.2 Primal-Dual Algorithms for A General Class

Model. In the model, there is a set of resources \mathcal{E} and requests arrive online. At the arrival of request i , a set of feasible strategies (actions) \mathcal{S}_i to satisfy request i is revealed. Each strategy $s_{ij} \in \mathcal{S}_i$ consists of a subset of resources in \mathcal{E} . Each resource e is associated to an *arbitrary* non-negative non-decreasing cost function f_e and the cost induced by resource e depends on the set of requests using e . The cost of a solution is the total cost of resources, i.e., $\sum_e f_e(A_e)$ where A_e is the set of requests using resource e . The goal is design an algorithm that upon the arrival of each request, selects a feasible strategy for the request while maintaining the cost of the overall solution as small as possible.

Result. Following the primal-dual approach based on configuration LPs and the smoothness notion (Definition 1.1), we derive a simple competitive algorithm which achieves optimal bounds in several applications.

Theorem 1.1 *Assume that all resource cost functions are (λ, μ) -smooth for some parameters $\lambda > 0, \mu < 1$. Then there exists a greedy $\frac{\lambda}{1-\mu}$ -competitive algorithm for the general problem.*

Applications.

We show the applicability of the theorem by deriving competitive algorithms for several problems in online setting, such as MINIMUM POWER SURVIVAL NETWORK ROUTING, VECTOR SCHEDULING, ENERGY-EFFICIENT SCHEDULING, PRIZE COLLECTING ENERGY-EFFICIENT SCHEDULING, SUBSPACE APPROXIMATION, NON-CONVEX FACILITY LOCATION. We describe some of the most representative ones below.

In ONLINE ENERGY-EFFICIENT SCHEDULING, one has to process jobs on unrelated machines. Each job has a released date, a deadline and a processing volume. Each job has to be assigned and fully processed in a machine between its release date and deadline. An algorithm can choose appropriate speed to process jobs and that incur energy cost. The objective is to minimize the total energy cost. No result has been known for this problem in parallel machine environments. Among others, a difficulty is the construction of formulation with bounded integrality gap. We notice that for this problem, Gupta et al. [67] gave a primal-dual competitive algorithm for a single machine. However, their approach cannot be used for unrelated parallel machines due to the large integrality gap of their formulation. For this problem, we present competitive algorithms for *arbitrary* energy cost functions beyond the convexity property. Note that the convexity of cost functions is a crucial property employed in previous work. If the energy cost have typical form $f(x) = x^\alpha$ (as the function of the speed x) for some parameter α then the competitive ratio of our algorithm is $O(\alpha^\alpha)$. This competitive ratio is optimal up to a constant factor.

In ONLINE NON-CONVEX FACILITY LOCATION, clients arrive online and have to be assigned to facilities. The cost of a facility consists of a fixed opening cost and a serving cost, which is an arbitrary monotone function depending on the number of clients assigned to the facility. The objective is to minimize the total client-facility connection cost and the facility cost. This problem is related to the capacitated network design and energy-efficient routing problems [6, 87]. In the latter, given a graph and a set of connectivity demands, the cost of each edge is *uniform* and given by $c + f^\alpha$ if $f > 0$ and 0 if $f = 0$, where c is a fixed cost for every edge and f is the total of flow passing through the edge. (Here uniformity means the cost functions are the same for every edge.) The objective is to minimize the total cost while satisfying all connectivity demands. Antoniadis et al. [6], Krishnaswamy et al. [87] have provided online/offline algorithms with poly-logarithmic guarantees. It is an intriguing open questions (originally raised in [4]) to design a poly-logarithmic competitive algorithm for non-uniform cost functions. The ONLINE NON-CONVEX FACILITY LOCATION can be seen as a step towards this goal. In fact, the former can be considered as the connectivity problem on a simple depth-2-graph and the cost functions are now non-uniform.

Using our primal-dual framework, we derive a $O(\log n + \frac{\lambda}{1-\mu})$ -competitive algorithm if the cost function is (λ, μ) -smooth. Specifically, the algorithm is inspired by a primal-dual algorithm [58] in the classic setting and our configuration LP-based approach. In particular, for the problem with non-uniform cost functions such as

$c_i + w_i f_i^\alpha$ where c_i, w_i are parameters depending on facility i and f_i is the number of clients assigned to facility i , the algorithm yields a competitive ratio of $O(\log n + \alpha^\alpha)$.

Besides, the algorithm mentioned in Theorem 1.1 can be used in the offline setting. Restricted to the class of polynomials with non-negative coefficients, our algorithm yields the competitive ratio of $O(\alpha^\alpha)$ while the best-known approximation ratio is $B_\alpha \approx (\frac{\alpha}{\log \alpha})^\alpha$ [94]. Our greedy algorithm is light-weight and much simpler and faster than that in [94] which requires solving an LP of exponential size and rounding fractional solutions. Hence, our algorithm can also be used to design approximation algorithms if one looks for the tradeoff between the simplicity and the performance guarantee.

1.4.3 Primal-Dual Approach for 0 – 1 Covering Problems

0 – 1 Covering Problems. Let \mathcal{E} be a set of n resources and let $f : \{0, 1\}^n \rightarrow \mathbb{R}^+$ be an *arbitrary* monotone cost function. Let $x_e \in \{0, 1\}$ be a variable indicating whether resource e is selected. If a resource e is selected then it remains selected until the end. The covering constraints $\sum_e a_{i,e} x_e \geq 1$ for every i are revealed one-by-one and at any step, one needs to maintain a feasible integer solution \mathbf{x} . The goal is to design an algorithm that minimizes $f(\mathbf{x})$ subject to the online covering constraints and $x_e \in \{0, 1\}$ for every e such that variables can only be changed from 0 to 1.

In order to design algorithms for this class of problem, we consider the multilinear extension of function f .

Definition 1.5 Given $f : \{0, 1\}^n \rightarrow \mathbb{R}^+$, its multilinear extension $F : [0, 1]^n \rightarrow \mathbb{R}^+$ is defined as

$$F(\mathbf{x}) := \sum_S \prod_{e \in S} x_e \prod_{e \notin S} (1 - x_e) \cdot f(\mathbf{1}_S)$$

where $\mathbf{1}_S$ is the characteristic vector of S (i.e., the e^{th} -component of $\mathbf{1}_S$ equals 1 if $e \in S$ and equals 0 otherwise).

An alternative way to define F is to set $F(\mathbf{x}) = \mathbb{E}[f(\mathbf{1}_T)]$ where T is a random set such that a resource e appears in T with probability x_e . Note that $F(\mathbf{1}_S) = f(\mathbf{1}_S)$.

Our algorithm, as well as the one in [11] for convex with monotone gradients and the recent algorithm for ℓ_k -norms [101], are extensions of the Buchbinder-Naor primal-dual framework [33]. A distinguishing point of our algorithm compared to the ones in [11, 101] relies on the multiplicative update, which is crucial in online primal-dual methods. The approaches in [11, 101] use the gradient $\nabla f(\mathbf{x})$ at the current primal solution \mathbf{x} to define a multiplicative update for the primal. In our approach, we multiplicatively update the primal by some parameter related to the gradient of the multilinear extension $\nabla F(\mathbf{x})$. This parameter is always maintained to be at least $\nabla F(\mathbf{x})$ and in case $\nabla F(\mathbf{x})$ is non-decreasing, the parameter is indeed equal to $\nabla F(\mathbf{x})$. This multiplicative update, together with the configuration LPs and the notion of local smoothness, enable us to derive a competitive algorithm for convex objective functions whose gradients are not necessarily monotone and more generally, for non-convex objectives.

Building on our approach, we derive a primal-dual algorithm in which the competitive ratio is determined in terms of the min-locally-smoothness parameters (Definition 1.2).

Theorem 1.2 Let F be the multilinear extension of the objective cost f and d be the maximal row sparsity of the constraint matrix, i.e., $d = \max_i |\{a_{i,e} : a_{i,e} > 0\}|$. Assume that F is

$(\lambda, \frac{\mu}{\ln(1+2d^2)})$ -min-locally-smooth for some parameters $\lambda > 0$ and $\mu < 1$. Then there exists a $O(\frac{\lambda}{1-\mu} \cdot \ln d)$ -competitive algorithm for the fractional covering problem.

Applications.

We apply our algorithm to several widely-studied classes of functions in optimization.

First, for the class of non-negative polynomials of degree k , the algorithm yields a $O((k \log d)^k)$ -competitive fractional solution that matches a result in [11].

Second, beyond convexity, we consider a natural class of non-convex cost functions which represent a typical behaviour of resources in serving demand requests. Non-convexity represents a strong barrier in optimization in general and in the design of algorithms in particular. We show that our algorithm is competitive for this class of functions.

Finally, we illustrate the applicability of our algorithm to the class of submodular functions. We make a connection between the local-smooth parameters to the concept of *total curvature* κ of submodular functions. The total curvature has been widely used to determine both upper and lower bounds on the approximation ratios for many submodular and machine learning problems [47, 65, 12, 124, 80, 117]. We show that our algorithm yields a $O(\frac{\log d}{1-\kappa})$ -competitive fractional solution for the problem of minimizing a submodular function under covering constraints. To the best of our knowledge, the submodular minimization under general covering constraints has not been studied in the online setting.

1.4.4 Primal-Dual Approach for 0 – 1 Packing Problems

0 – 1 Packing Problems. Let \mathcal{E} be a set of n resources and let $f : \{0, 1\}^n \rightarrow \mathbb{R}^+$ be an arbitrary cost function. Let $x_e \in \{0, 1\}$ be a variable indicating whether resource e is selected. The packings constraints $\sum_e b_{ie} x_e \leq 1$ for every i are given in advance and resources e are revealed online one-by-one. At any time, one needs to maintain a feasible integer solution x . The goal is to design an algorithm that maximizes $f(x)$ subject to the online packing constraints and $x_e \in \{0, 1\}$ for every e .

Following the salient ideas for fractional covering problems, we derive an algorithm with the competitive ratio determined in terms of the max-locally-smoothness parameters (Definition 1.3).

Theorem 1.3 *Let F be the multilinear extension of the objective cost f . Denote the row sparsity $d := \max_i |\{b_{ie} : b_{ie} > 0\}|$ and $\rho := \max_i \max_{e, e': b_{ie'} > 0} b_{ie} / b_{ie'}$. Assume that F is (λ, μ) -max-locally-smooth for some parameters $\lambda > 0$ and $\mu < 1$. Then there exists a $O(\frac{2\ln(1+d\rho)+\mu}{\lambda})$ -competitive algorithm for the fractional packing problem.*

Note that when f is a linear function, the smooth parameters are $\lambda = 1$ and $\mu = 0$. In this case, the performance guarantee is the same (up to a constant factor) to that of maximizing a linear function under packing constraints [33] and therefore asymptotically optimal.

Applications

We consider applications to online submodular maximization problems. Submodular functions are interesting since they are neither convex nor concave. Besides,

submodular maximization constitutes a major research agenda in optimization, machine learning and has been widely studying. However, in the context of online algorithms, not much has been known especially for submodular maximization with constraints. Designing competitive algorithms for online submodular maximization has been identified as an important direction in the recent survey [85]. Buchbinder et al. [36] have studied the online problem of maximizing the sum of weighted rank functions subject to matroid constraints. The objective here is a particular submodular function. The authors give an algorithm with competitive ratio depending logarithmically on the numbers of elements and on weighted rank functions. In another approach, Buchbinder et al. [38] have considered submodular optimization with preemption, where one can reject previously accepted elements, and have given constant competitive algorithms for unconstrained and knapsack-constraint problems.

We show that there exists an algorithm which outputs competitive fractional solutions for online submodular maximization with packing constraints. The competitive ratio is $O(\log(1 + d\rho))$ which is independent of the submodular objective. Note that using the online contention resolution rounding schemes [57], one can obtain randomized algorithms for several specific constraint polytopes, for example, knapsack polytopes, matching polytopes and matroid polytopes.

1.4.5 Related work

In this section we summarize further related work which are not mentioned earlier.

Primal-dual methods have been shown to be powerful tools in online computation. Buchbinder and Naor [33] presented a primal-dual method for linear programs with packing/covering constraints. Their method unifies several previous potential-function-based analyses and give a principled approach to design and analyze algorithms for problems with linear relaxations. Convex objective functions have been extensively studied in online settings in recent years, in areas such as energy-efficient scheduling [3, 103, 50, 78, 10], paging [96], network routing [67], combinatorial auctions [28, 77] and matching [51]. Recently, Azar et al. [11] gave an unified framework for covering/packing problems with convex objectives whose gradients are monotone. Consequently, improved algorithms have been derived for several specific problems. The above approaches rely crucially on the convexity of cost functions. Specifically, the construction of dual programs is based on convex conjugates and Fenchel duality for primal convex programs. Very recently, Nagarajan and Shen [101] have considered objective functions as the of sum of ℓ_k -norms. This class of functions does not fall into the framework developed in [11], since the gradients are not necessarily monotone. Nagarajan and Shen [101] proved that the algorithm presented in [11] yields a nearly tight $O(\log d + \log \frac{\max a_{ij}}{\min a_{ij}})$ -competitive ratio where a_{ij} 's are entries in the covering matrix. Using these approaches, it is not clear how to design competitive algorithms for *non-convex* functions or even for other convex functions with non-monotone gradient. A distinguishing point of our approach is that it gives a framework to study non-convex cost functions.

1.5 Contributions to Algorithmic Game Theory

The goal of this section is to present a unified primal-dual approach to analyze efficiency of games. Recall that the approach consists of associating a game to an optimization problem and formulate a corresponding configuration integer program to the problem. Next consider the linear program by relaxing the integer constraints

and its dual LP. Then given a Nash equilibrium, construct dual variables in such a way that one can relate the dual objective to the cost of the Nash equilibrium. Intuitively, one dual constraint corresponds exactly to the definition of Nash equilibrium and the other constraint settles the PoA bounds. The PoA is then bounded by the primal objective (essentially, the cost of the Nash equilibrium) and the dual objective (a lower bound of the optimum cost by weak duality).

We show the potential and the wide applicability of the approach throughout various results in the contexts of complete and incomplete-information environments, from the settings of congestion games to welfare maximization. The approach allows us to unify several previous results and establish new ones beyond the current techniques. It is worthy to note that the analyses are simple and are guided by dual LP in the same spirit as primal-dual methods for designing algorithms. Moreover, under the lens of LP duality, the notion of smooth games in both full-information settings [108] and incomplete-information settings [109, 119], the recent notion of no-envy learning [49] and the new notion of dual smooth (in this paper) can be naturally derived, which lead to the optimal bounds of the PoA of several games.

1.5.1 Notations and Preliminaries

Nash Equilibria and Smooth Games

In a game, there are n players and each player i selects a strategy s_i from a set \mathcal{S}_i for $1 \leq i \leq n$ and this forms a *strategy profile* $\mathbf{s} = (s_1, \dots, s_n)$. The cost $C_i(\mathbf{s})$ of player i is a function of the strategy profile \mathbf{s} . A *pure Nash equilibrium* is a strategy profile \mathbf{s} such that no player can decrease its cost via a unilateral deviation; that is, for every player i and every strategy $s'_i \in \mathcal{S}_i$,

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

where \mathbf{s}_{-i} denotes the strategies chosen by all players other than i in \mathbf{s} . The notion of Nash equilibrium is extended to the following more general equilibrium concepts.

A *mixed Nash equilibrium* [102] of a game is a product distribution $\sigma = \sigma_1 \times \dots \times \sigma_n$ where σ_i is a probability distribution over the strategy set of player i such that no player can decrease its expected cost under σ via a unilateral deviation:

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

for every i and $s'_i \in \mathcal{S}_i$, where σ_{-i} is the product distribution of all $\sigma_{i'}$'s other than σ_i .

A *correlated equilibrium* [8] of a game is a joint probability distribution σ over the strategy profile of the game such that

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

for every i and $s_i, s'_i \in \mathcal{S}_i$.

Finally, a *coarse correlated equilibrium* [98] of a game is a joint probability distribution σ over the strategy profiles of the game such that

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

for every i and $s'_i \in \mathcal{S}_i$.

These notions of equilibria are presented in the order from the least to the most general ones and a notion captures the previous one as a strict subset.

Let $C : \mathcal{S}_1 \times \dots \times \mathcal{S}_n \rightarrow \mathbb{R}$ be the social cost of a game. Then the *price of anarchy* (PoA) of a minimization game is defined as $\sup_{\mathbf{s}} C(\mathbf{s}) / \text{OPT}$ where the supremum is taken all Nash equilibria \mathbf{s} and OPT is the optimal social cost. The definition extends to the notions of mixed PoA, coarse PoA and coarse correlated PoA.

Definition 1.6 ([108]) A game with a joint cost objective function $C(\mathbf{s}) = \sum_{i=1}^n C_i(\mathbf{s})$ is (λ, μ) -smooth if for every two outcomes \mathbf{s} and \mathbf{s}^* ,

$$\sum_{i=1}^n C_i(\mathbf{s}_i^*, \mathbf{s}_{-i}) \leq \lambda \cdot C(\mathbf{s}^*) + \mu \cdot C(\mathbf{s})$$

The *robust price of anarchy* of a game G is defined as

$$\rho(G) := \inf \left\{ \frac{\lambda}{1 - \mu} : \text{the game is } (\lambda, \mu)\text{-smooth where } \mu < 1 \right\}$$

Mechanism Design and Smooth Auctions

In a general mechanism design setting, each player i has a set of actions \mathcal{A}_i for $1 \leq i \leq n$. Given an action $a_i \in \mathcal{A}_i$ chosen by each player i for $1 \leq i \leq n$, which lead to the action profile $\mathbf{a} = (a_1, \dots, a_n) \in \mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$, the auctioneer decides an outcome $o(\mathbf{a})$ among the set of feasible outcomes \mathcal{O} . Each player i has a *valuation* (or *type*) v_i taking values in a parameter space \mathcal{V}_i . For each outcome $o \in \mathcal{O}$, player i has *utility* $u_i(o, v_i)$ depending on the outcome of the game and its valuation v_i . Since the outcome $o(\mathbf{a})$ of the game is determined by the action profile \mathbf{a} , the utility of a player i is denoted as $u_i(\mathbf{a}; v_i)$. We are interested in auctions that in general consist of an allocation rule and a payment rule. Given an action profile $\mathbf{a} = (a_1, \dots, a_n)$, the auctioneer decides an allocation and a payment $p_i(\mathbf{a})$ for each player i . Then, the *utility* of player i with valuation v_i , following the quasi-linear utility model, is defined as

$$u_i(\mathbf{a}; v_i) = v_i - p_i(\mathbf{a})$$

The *social welfare* of an auction is defined as the total utility of all participants (the players and the auctioneer):

$$\text{SW}(\mathbf{a}; \mathbf{v}) = \sum_{i=1}^n u_i(\mathbf{a}; v_i) + \sum_{i=1}^n p_i(\mathbf{a})$$

In incomplete-information settings, the valuation v_i of each player is a private information and is drawn independently from a publicly known distribution \mathbf{F} with density function \mathbf{f} . Let $\Delta(\mathcal{A}_i)$ be the set of probability distributions over the actions in \mathcal{A}_i . A strategy of a player is a mapping $\sigma_i : \mathcal{V}_i \rightarrow \Delta(\mathcal{A}_i)$ from a valuation $v_i \in \mathcal{V}_i$ to a distribution over actions $\sigma_i(v_i) \in \Delta(\mathcal{A}_i)$.

Definition 1.7 (Bayes-Nash equilibrium) A strategy profile $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)$ is a Bayes-Nash equilibrium (BNE) if for every player i , for every valuation $v_i \in \mathcal{V}_i$, and for every action $a'_i \in \mathcal{A}_i$:

$$\mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}(v_i)} \left[\mathbb{E}_{\mathbf{a} \sim \boldsymbol{\sigma}(\mathbf{v})} [u_i(\mathbf{a}; v_i)] \right] \geq \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}(v_i)} \left[\mathbb{E}_{\mathbf{a}_{-i} \sim \boldsymbol{\sigma}_{-i}(\mathbf{v}_{-i})} [u_i(a'_i, \mathbf{a}_{-i}; v_i)] \right]$$

For a vector \mathbf{w} , we use \mathbf{w}_{-i} to denote the vector \mathbf{w} with the i -th component removed. Besides, $\mathbf{F}_{-i}(v_i)$ stands for the probability distribution over all players other than i conditioned on the valuation v_i of player i .

The price of anarchy of Bayes-Nash equilibria of an auction is defined as

$$\inf_{F, \sigma} \frac{\mathbb{E}_{v \sim F} [\mathbb{E}_{\mathbf{a} \sim \sigma(v)} [\text{SW}(\mathbf{a}; \mathbf{v})]]}{\mathbb{E}_{v \sim F} [\text{OPT}(\mathbf{v})]}$$

where the infimum is taken over Bayes-Nash equilibria σ and $\text{OPT}(\mathbf{v})$ is the optimal welfare with valuation profile \mathbf{v} .

Smooth auctions have been defined by Roughgarden [109] and Syrgkanis and Tardos [119]. The definitions are slightly different but both are inspired by the original one [108] (Definition 1.6) and all known smoothness-based proofs can be equivalently analyzed by one of these definitions.

Definition 1.8 ([109]) For parameters $\lambda, \mu \geq 0$, an auction is (λ, μ) -smooth if for every valuation profile $\mathbf{v} = (v_1, \dots, v_n)$, there exist action distributions $D_1^*(\mathbf{v}), \dots, D_n^*(\mathbf{v})$ over $\mathcal{A}_1, \dots, \mathcal{A}_n$ such that, for every action profile \mathbf{a} ,

$$\sum_i \mathbb{E}_{a_i^* \sim D_i^*(\mathbf{v})} [u_i(a_i^*, \mathbf{a}_{-i}; v_i)] \geq \lambda \cdot \text{SW}(\mathbf{a}^*; \mathbf{v}) - \mu \cdot \text{SW}(\mathbf{a}; \mathbf{v}). \quad (1.3)$$

Definition 1.9 ([119]) For parameters $\lambda, \mu \geq 0$, an auction is (λ, μ) -smooth if for every valuation profile $\mathbf{v} = (v_1, \dots, v_n)$, there exist action distributions $D_1^*(\mathbf{v}), \dots, D_n^*(\mathbf{v})$ over $\mathcal{A}_1, \dots, \mathcal{A}_n$ such that, for every action profile \mathbf{a} ,

$$\sum_i \mathbb{E}_{a_i^* \sim D_i^*(\mathbf{v})} [u_i(a_i^*, \mathbf{a}_{-i}; v_i)] \geq \lambda \cdot \text{OPT}(\mathbf{v}) - \mu \cdot \text{REV}(\mathbf{a}; \mathbf{v}). \quad (1.4)$$

1.5.2 Games in Full-Information Settings

We first revisit smooth games by the primal-dual approach and show that the primal-dual approach captures the smoothness framework. The smoothness framework was introduced in [108] and quickly became a standard technique. More precisely, Roughgarden [108] proved that every (λ, μ) -smooth game admits the PoA at most $\lambda/(1 - \mu)$. Through the duality approach, we show that in terms of techniques to study the PoA for complete information settings, LP duality and the smoothness framework are exactly the same thing. Specifically, one of the dual constraint corresponds exactly to the definition of smooth games given in [108].

Theorem 1.4 *The primal-dual approach captures the smoothness framework in full information settings.*

Congestion Games

We consider fundamental classes of *congestion games* in which we revisit and unify results for both atomic and non-atomic congestion games and prove the optimal PoA bound of coarse correlated equilibria in splittable congestion games.

Atomic congestion games. In this class, although the PoA bound follows the results for smooth games (Theorem 1.4), we provide another configuration formulation and a similar primal-dual approach. The purpose of this formulation is twofold. First it shows the flexibility of the primal-dual approach. Second, it sets up the ground for an unified approach to other classes of congestion games.

Non-atomic congestion games. In this class, we re-prove the *optimal* PoA bound [112]. Along the line towards the optimal PoA bound for non-atomic congestion games, the equilibrium characterization by a variational inequality is at the core of the analyses [112, 48, 46]. In our proof, we establish the optimal PoA directly by the mean of LP duality. By the LP duality as the unified approach, one can clearly observe that non-atomic setting is a version of the atomic setting in large games (in the sense of [55]) in which each player weight becomes negligible (hence, the PoA of the atomic congestion games tend to that of non-atomic ones). Besides, an advantage with LP approaches is that one can benefit from powerful techniques that have been developing for linear programming. Concretely, using the general framework on resource augmentation and primal-dual recently presented [92], we manage to recover and extend a resource augmentation result related to non-atomic setting [111].

Theorem 1.5 *In every non-atomic congestion game, for any constant $r > 0$, the cost of an equilibrium is at most $1/r$ the optimum of the underlying optimization problem in which each demand is multiplied by a factor $(1 + r)$.*

Splittable congestion games. Roughgarden and Schoppmann [110] has presented a *local* smoothness property, a refinement of the smoothness framework, and proved that every (λ, μ) -local-smooth splittable game admits the PoA $\lambda/(1 - \mu)$. This bound is tight for a large class of scalable cost functions in splittable games and holds for PoA of pure, mixed, correlated equilibria. However, this bound does not hold for coarse correlated equilibria and it remains an intriguing open question raised in [110]. Building upon the resilient ideas of non-atomic and atomic settings, we define a notion, called *dual smoothness*, which is inspired by the dual constraints. This new notion indeed leads to the *tight* PoA bound for coarse correlated equilibria in splittable games for a large class of cost functions; that answers the question in [110]. Not that the matching lower bound is given in [110] and that holds even for pure equilibria.

Definition 1.10 *A cost function $\ell : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is (λ, μ) -dual-smooth if for every vectors $\mathbf{u} = (u_1, \dots, u_n)$ and $\mathbf{v} = (v_1, \dots, v_n)$,*

$$v\ell(u) + \sum_{i=1}^n u_i(v_i - u_i) \cdot \ell'(u) \leq \lambda \cdot v\ell(v) + \mu \cdot u\ell(u)$$

where $u = \sum_{i=1}^n u_i$ and $v = \sum_{i=1}^n v_i$. A splittable congestion game is (λ, μ) -dual-smooth if for every resource e in the game, function ℓ_e is (λ, μ) -dual-smooth.

Theorem 1.6 *The price of anarchy of coarse correlated equilibria of a splittable congestion game G is at most $\inf \lambda/(1 - \mu)$ where the infimum is taken over (λ, μ) such that G is (λ, μ) -dual-smooth. This bound is tight for the class of scalable cost functions.*

1.5.3 Games in Incomplete-Information Settings

We next consider the inefficiency of Bayes-Nash equilibria in the context of welfare maximization in incomplete-information environments.

Smooth Auctions. The notion of smooth auctions in incomplete-information settings, inspired by the original smoothness framework [108], has been introduced

by Roughgarden [109], Syrgkanis and Tardos [119]. This powerful notion has been widely used to study the PoA of Bayes-Nash equilibria (see the recent survey [114]). We show that the primal-dual approach captures the smoothness framework in incomplete-information settings. In other words, the notion of smooth auctions can be naturally derived from dual constraints in the primal-dual approach.

Informal Theorem 1.1 *The primal-dual approach captures the smoothness framework in incomplete-information settings.*

Simultaneous Item-Bidding Auctions: Beyond Smoothness. Many PoA bounds in auctions are settled by smoothness-based proofs. However, there are PoA bounds for auctions proved via non-smooth techniques and these techniques seem more powerful than the smoothness framework in such auctions. Representative examples are the simultaneous first- and second-price auctions where players' valuations are sub-additive. Feldman et al. [54] have proved that the PoA is constant while the smooth argument gives only logarithmic guarantees. We show that in this context, our approach is beyond the smoothness framework and also captures the non-smooth arguments in [54] by re-establishing their results. Specifically, a main step in our analysis — proving the feasibility of a dual constraint — corresponds exactly to a crucial claim in [54]. From this point of view, the primal-dual approach helps to identify the key steps in settling the PoA bounds.

Informal Theorem 1.2 ([54]) *Assume that players have independent distributions over sub-additive valuations. Then, every Bayes-Nash equilibrium of a first-price auction and of a second price auction has expected welfare at least 1/2 and 1/4 of the maximal welfare, respectively.*

Subsequently, we illuminate the potential of the primal-dual approach in formulating new concepts. Concretely, Daskalakis and Syrgkanis [49] have very recently introduced *no-envy learning dynamic* — a novel concept of learning in auctions. Note that when players have fractionally sub-additive (XOS) valuations¹, no-envy outcomes are a relaxation of no-regret outcomes. No-envy dynamics have advantages over no-regret dynamics. In particular, no-envy outcomes maintain the approximate welfare optimality of no-regret outcomes while ensuring the computational tractability. Perhaps surprisingly, there is a connection between the primal-dual approach and no-envy dynamics. Indeed, the latter can be naturally derived from the dual constraints very much in the same way as the smoothness argument is. We show this connection by revisiting the following theorem by the means of the primal-dual approach.

Informal Theorem 1.3 ([49]) *Assume that players have XOS valuations. Then, every no-envy dynamic has the average welfare at least half the expected optimal welfare.*

Sequential Auctions. To illustrate the applicability of the primal-dual approach, we consider thereafter another format of auctions — sequential auctions. In a simple model of sequential auctions, items are sold one-by-one via single-item auctions. Sequential auctions have a long and rich literature [86] and sequentially selling items leads to complex issues in analyzing PoA. Leme et al. [90], Syrgkanis and Tardos

¹A valuation $v(\cdot)$ is XOS if there exists a family of vectors $\mathcal{W} = (w^\ell)_\ell$ where $w^\ell \in \mathbb{R}_+^m$ such that $v(S) = \max_{w^\ell \in \mathcal{W}} \sum_{j \in S} w_j^\ell \forall S \subset [m]$. The class XOS is a subset of sub-additive functions and is a superset of sub-modular functions.

[118] have studied sequential auctions for matching markets and matroid auctions in complete and incomplete-information settings in which at each step, an item is sold via the first-price auctions. In this paper, we consider the sequential auctions for sponsored search via the second-price auctions. Informally, auctioneer sells advertising slots one-by-one in the non-increasing order of click-through-rates (from the most attractive to the least one). At each step, players submit bids for the currently-selling slot and the highest-bid player receives the slot and pays the second highest bid. In the auction, we study the PoA of perfect Bayesian equilibria and show the following improvement over the best-known PoA bound of 2.927 [41] for the sponsored search problem.

Informal Theorem 1.4 *The PoA of sequential second-price auctions for the sponsored search problem is at most 2.*

An observation is that although the behaviour of players in sequential auctions might be complex, the performance guarantee is better than the currently best-known one for simultaneous second price auctions for the sponsored search problem. Consequently, this result shows that the efficiency of sequential auctions is not necessarily worse than the simultaneous ones (and also analyzing sequential auction is not necessarily harder than analyzing simultaneous ones). Moreover, using the primal-dual approach, the proof is fairly simpler than the smoothness-based one.

Building upon the salient ideas for the sponsored search problem, we provide an improved PoA bound of 2 for the matching market problem where the best known PoA bound is $2e/(e-1) \approx 3.16$ due to Syrgkanis and Tardos [118]. That also answers a question raised in [118] whether the PoA in the incomplete-information settings must be strictly larger than the best-known PoA bound (which is 2) in the full-information settings.

Informal Theorem 1.5 *The PoA of sequential first-price auctions for the matching market problem is at most 2.*

1.5.4 Related works

As the main point of the paper is to emphasize the primal-dual approach to study game efficiency, in this section we mostly concentrate on currently existing methods. Results related to specific problems will be summarized in the corresponding sections.

The most closely related to our work is a recent result [88]. In their approach, Kulkarni and Mirrokni [88] considered a convex formulation of a given game and its dual program based on Fenchel duality. Then, given a Nash equilibrium, the dual variables are constructed by relating the cost of the Nash equilibrium to that of the dual objective. In high-level, our approach has the same idea as [88] and both approaches indeed have inspired by the standard primal-dual and dual-fitting in the design of algorithms. Our approach is distinguished to that in [88] in the two following aspects. First, we consider arbitrary (non-decreasing) objective functions and make use of configuration LPs in order to reduce substantially the integrality gap while the approach in [88] needs convex objective functions. In term of approaches based on mathematical programs in approximation algorithms, we came up with stronger formulations than those in [88] — a crucial point toward optimal bounds. Second, we have shown a wide applicability of our approach from full-information environments to incomplete-information ones while the approach in [88] deals only with full-information settings. A question, which has been raised

in a the recent survey [114], is whether the framework in [88] could be extended to incomplete-information settings. Our primal-dual approach gives a positive answer to that question.

The connection between LP duality and the PoA have been previously considered by Nadav and Roughgarden [100] and Bilo [25]. Both papers follow an approach which is different to ours. Roughly speaking, given a game they consider corresponding natural formulations and incorporate the equilibrium constraint directly to the primal (whereas in our approach the equilibrium constraint appears naturally in the dual). This approach suffers the integrality-gap issue when one considers pure Nash equilibria and the objectives are non-linear or non-convex.

For the problems studied in the paper, we systematically strengthen natural LPs by the construction of configuration LPs presented in [94]. Makarychev and Sviridenko [94] propose a scheme that consists in solving the new LPs (with exponential number of variables) and rounding the fractional solutions to integer ones using decoupling inequalities for optimization problems. Instead of rounding techniques, we consider a primal-dual approach which is more adequate to studying game efficiency.

The smoothness framework has been introduced by Roughgarden [108]. This simple, elegant framework gives tight bounds for many classes of games in full-information settings including the celebrated atomic congestion games (and others in [108, 24]). Subsequently, Roughgarden and Schoppmann [110] presented a similar notion, called local-smoothness, to study the PoA of splittable games in which players can split their flow to arbitrarily small amounts and route the amounts in different manners. The local-smoothness is also powerful. It has been used to settle the PoA for a large class of cost functions in splittable games [110] and in opinion formation games [23].

The smoothness framework has been extended to incomplete-information environments by Roughgarden [109], Syrgkanis and Tardos [119]. It has successfully produced tight PoA bounds for several widely-used auction formats. We recommend the reader to a very recent survey [114] for applications of the smoothness framework in incomplete-information settings. However, the smoothness argument has its limit. As mentioned earlier, the most illustrative examples are the simultaneous first and second price auctions where players' valuations are sub-additive. Feldman et al. [54] have proved that the PoA is constant while the smooth argument gives only logarithmic guarantees. An interesting open direction, as raised in [114], is to develop new approaches beyond the smoothness framework.

Linear programming (and mathematical programming in general) has been a powerful tool in the development of game theory. There is a vast literature on this subject. One of the most interesting recent treatments on the role of linear programming in game theory is the book [122]. Vohra [122] revisited fundamental results in mechanism design in an elegant manner by the means of linear programming and duality. It is surprising to see that many results have been shaped nicely by LPs.

1.6 Contributions to Efficient Online Learning Algorithms

We consider the direction of designing efficient learning algorithms and present regularity conditions which enables such algorithms. We establish the applicability of our approach by deriving efficient online learning algorithms in welfare and revenue maximization problems.

1.6.1 Definitions and Framework

General problem. We are given a convex set $\mathcal{K} \subset \mathbb{R}^n$ and a family of gain functions \mathcal{F} . At each time step $t = 1, 2, \dots$, an algorithm chooses $\mathbf{x}^t \in \mathcal{K}$. After the algorithm committed to its choice, an adversary selects a function $f^t \in \mathcal{F} : \mathcal{K} \rightarrow \mathbb{R}$ that subsequently induces the gain $f^t(\mathbf{x}^t)$ for the algorithm. Function f^t is considered as a *black-box* and one can only query its values. The goal is to achieve the total gain approximately close to that obtained by the best decision $\mathbf{x} \in \mathcal{K}$ in hindsight.

We consider the following notion of regret which measures the performance of algorithms.

Definition 1.11 *An algorithm is $(r, R(T))$ -regret if for arbitrary total number of time steps T and for any sequence of cost functions $f_1, \dots, f_T \in \mathcal{F}$,*

$$\sum_{t=1}^T f^t(\mathbf{x}^t) \geq r \cdot \max_{\mathbf{x} \in \mathcal{K}} \sum_{t=1}^T f^t(\mathbf{x}) - R(T).$$

In the paper, we seek algorithms with regret bound $(r, R(T))$ such that $r > 0$ is as large as possible (ideally, close to 1) and $R(T)$ is sublinear as a function of T , i.e., $R(T) = o(T)$. We also call r as the *approximation ratio* of the algorithm.

As mentioned earlier, efficient online algorithms do not exist in general adversarial environments without any structure regularity. In our framework, we introduce the regularity notion of (λ, μ) -concavity (Definition 1.4) which is crucial in our framework in order to design efficient online learning algorithms.

Our approach follows the mirror descent algorithm. Using the latter, one can derive the regret bound as a function of the concavity parameters *assuming* access to the gradient values. However, the main issue of the problem is that functions f^t 's are given as black-boxes and we do not have access to the gradients of the functions. (However, we can query the values of the functions.) We bypass that issue by considering a lattice discretization of the domain $[0, 1]^n$ and the multilinear extension of the functions restricted on the lattice. The multilinear extension approximates the corresponding function when the latter is regular, for example a Lipschitz function. More importantly, the gradient of the multilinear extension can be computed by requesting function values at lattices points. Subsequently, applying the standard technique of mirror descent w.r.t the gradient of the multilinear extension, we can bound the regret of our algorithm.

Theorem 1.7 *Given functions $f^t : [0, 1]^n \rightarrow \mathbb{R}$ for $1 \leq t \leq T$. Let F^t be the multilinear extension of the discretization of f^t based on a lattice \mathcal{L} (defined later). Assume that for every $1 \leq t \leq T$, the multilinear extension F^t is (λ, μ) -concave, $\|\nabla F^t\|_*$ is bounded by L and the Bregman divergence (defined later) based on some α_Φ -strongly convex function Φ is bounded by G^2 . Then, there exists an efficient online randomized algorithm that achieves*

$$\sum_{t=1}^T \mathbb{E}[f^t(\mathbf{x}^t)] \geq \frac{\lambda}{\mu} \cdot \max_{\mathbf{x} \in \mathcal{K}} \sum_{t=1}^T f^t(\mathbf{x}) - O\left(\frac{GL}{\mu} \sqrt{2\alpha_\Phi T}\right)$$

The algorithm makes a polynomial number of value queries to functions f^t 's.

1.6.2 Application to Fictitious Play in Smooth Auctions

We consider adaptive dynamics in auctions. In the setting, there is an underlying auction \mathcal{o} and there are n players, each player i has a set of actions \mathcal{A}_i and a

valuation function v_i taking values in $[0, 1]$ (by normalization). In every time step $1 \leq t \leq T$, each player i selects a strategy which is a distribution in $\Delta(\mathcal{A}_i)$ according to some given adaptive dynamic. After all players committed to their strategies, which results in a strategy profile $\sigma^t \in \Delta(\mathcal{A})$, the auction induces a social welfare $\text{SW}(\sigma, \sigma^t) := \mathbb{E}_{\mathbf{a} \sim \sigma^t} [\text{SW}(\sigma(\mathbf{a}); \mathbf{v})]$. In this setting, we study the total welfare achieved by the given adaptive dynamic and compare it to the optimal welfare. This problem can be cast in the online optimization framework in which at time step t , the player strategy profile corresponds to the decision of the algorithm and subsequently, the gain of the algorithm is the social welfare induced by the auction w.r.t the strategy profile.

Smooth auctions (Definitions 1.9) is an important class of auctions in welfare maximization. The smoothness notion has been introduced in order to characterize the efficiency of (Bayes-Nash) equilibria of auctions. It has been shown that several auctions in widely studied settings are smooth; and many proof techniques analyzing equilibrium efficiency can be reduced to the smooth argument.

It has been proved that if an auction is (λ, μ) -smooth then every Bayes-Nash equilibrium of the auction has expected welfare at least λ/μ fraction of the optimal auction [109, 119]. The performance guarantee holds even for vanishing regret sequences. A sequence of actions profiles $\mathbf{a}^1, \mathbf{a}^2, \dots$, is a *vanishing regret sequence* if for every player i and action a'_i ,

$$\lim_T \frac{1}{T} \sum_{t=1}^T [u_i(a'_i, \mathbf{a}_{-i}^t; v_i) - u_i(\mathbf{a}^t; v_i)] \leq 0.$$

The smoothness framework does not extend to non-vanishing regret dynamics. However, several interesting dynamics are not guaranteed to have the vanishing regret property. In a recent survey, Roughgarden et al. [114] have raised a question whether adaptive dynamics without the vanishing regret condition can achieve approximate optimal welfare. Among others, *fictitious play* [30] is an interesting, widely-studied dynamic which attracts a significant attention in the community.

We consider a continuous version of fictitious play in smooth auctions, called Perturbed Discrete Time Fictitious Play (PDTFP). This dynamic in general does not admit vanishing regret [21, Example 1.2]. We prove that given an offline (λ, μ) -smooth auction, PDTFP dynamic achieves a $\lambda/(1 + \mu)$ fraction of the optimal welfare; thus answering the above open question of Roughgarden et al. [114]. To the best of our knowledge, prior to our work, no such guarantee has been proven in non-stationary, non-stochastic environments.

Theorem 1.8 *If the underlying auction σ is a (λ, μ) -smooth then the PDTFP dynamic achieves $\left(\frac{\lambda}{1+\mu}, R(T)\right)$ -regret where $R(T) = O\left(\frac{G\sqrt{T}}{1+\mu}\right)$ and G is a parameter defined later.*

1.6.3 Application to Revenue Maximization in Multi-Dimensional Environments

We consider online simultaneous second-price auctions with reserve prices. In this setting, there are n bidders² and m items to be sold to these bidders. At every time step $t = 1, 2, \dots, T$, the auctioneer selects reserve prices $r_i^t = (r_{i1}^t, \dots, r_{im}^t)$ for each bidder i where r_{ij}^t is the reserve price of item j for bidder i . Each bidder i for $1 \leq i \leq n$ has a (private) valuation $v_i^t : 2^{[m]} \rightarrow \mathbb{R}^+$ over subsets of items. After the reserve

²in this setting we call players as bidders

prices have been chosen, every bidder i picks a bid vector b_i^t where b_{ij}^t is the bid of bidder i on item j for $1 \leq j \leq m$. Then the auction for each item $1 \leq j \leq m$ works as follows: (1) remove all bidders i with $b_{ij}^t < r_{ij}^t$; (2) run the second price auction on the remaining bidders to determine the winner of item j ; (3) charge the winner of item j the maximum between r_{ij}^t and the second highest bid among non-removed bids b_{ij}^t . The objective of the auctioneer is to achieve the total revenue approximately close to that achieved by the best fixed reserve-price auction.

The second-price auctions with reserve prices in single-parameter environments have been considered in online learning framework by Roughgarden and Wang [113]. They gave a polynomial-time online algorithm that achieves half the revenue of the best fixed reserve-price auction minus a term $O(\sqrt{T} \log T)$ (so their algorithm is $(1/2, O(\sqrt{T} \log T))$ -regret in our terminology).

In this paper, we show that there exists an efficient online learning algorithm for the problem in multi-parameter environments. A distinguishing point of our algorithm compared to [113] is that we consider the Follow-the-Regularized-Leader approach whereas [113] considers the Follow-the-Perturbed-Leader strategy.

Theorem 1.9 *There exist an online polynomial-time algorithm that achieves a regret bound of $(1/2, O(m\sqrt{nmT} \log T))$ for revenue maximization in multi-parameter environments.*

Using our framework, the main task is to design a (λ, μ) -concave offline algorithm. Roughly speaking, the offline algorithm selects a reserve price or the fixed zero reserve price, each with probability $1/2$. This offline algorithm is based on the observation that the maximum between the revenue of a reserve price and that of the fixed zero reserve price gives rise to a $(1, 1)$ -concave function. Therefore, in order to guarantee a competitive revenue, the offline algorithm chooses each with probability $1/2$. Interestingly, the offline algorithm is exactly the one in [113] (for single-parameter environments) which has been proved using a different approach. Note that on the negative side, no efficient algorithm can achieve an approximation ratio better than $\frac{884}{885}$ unless $NP \subseteq RP$ [113].

1.6.4 Related Work

Our work is related to the design of efficient online learning algorithms and the auction design using learning techniques.

Online Learning. Online Learning, or Online Convex Optimization, is an active research domain. In this section, we only summarize works which are directly related to ours. We refer the reader to excellent books [116, 75] and references therein for a more complete overview. The first no-regret algorithm has been given by Hannan [68]. Subsequently, Littlestone and Warmuth [91] and Freund and Schapire [60] gave improved algorithms with regret $\sqrt{\log(|\mathcal{A}|)}o(T)$ where $|\mathcal{A}|$ is the size of the action space. However, these algorithms have running-time $\Omega(|\mathcal{A}|)$ which is exponential in many applications. An intriguing question is whether there exists a no-regret online algorithm with running-time polynomial in $\log(|\mathcal{A}|)$. As mentioned earlier, Hazan and Koren [74] proved that no such algorithm exists in general settings without any assumption on the structure. Characterizing necessary and sufficient conditions of the existence of efficient no-regret online algorithms is a major open question.

In their breakthrough, Kalai and Vempala [81] presented the first efficient online algorithm, called *Follow-the-Perturbed-Leader* (FTPL), for linear objective functions. The strategy consists of adding perturbation to the cumulative gain (payoff) of each

action and then selecting the action with the highest perturbed gain. This strategy has been generalized and successfully applied to several settings [73, 120, 49]. Recently Dudik et al. [52] gave an algorithm called *Generalized Follow-the-Perturbed-Leader* (GFTPL) and derived sufficient conditions for oracle-efficient online learning. As its name suggested, algorithm GFTPL generalizes the Follow-the-Perturbed Leader approach for linear functions [81] and its extension to submodular functions [73], contextual learning [120] and learning in simultaneous second-price auctions [49]. Consequently, Dudik et al. [52] obtained oracle-efficient no-regret algorithms for several classes in auction design and in contextual learning. Note that their approach requires an (exact/approximate) best-response oracle.

In the primal-dual aspect, the FTPL is indeed the *Follow-the-Regularized-Leader* (FTRL), introduced by Shalev-Shwartz and Singer [115], in which the random perturbation can be considered as a regularization. Moreover, Hazan and Kale [72] showed that FTRL is equivalent to Online Mirror Descent. In the paper, we consider the latter to design our online learning algorithm.

Smooth Auctions and Fictitious Play. The smoothness framework has been introduced in order to prove approximation guarantees for equilibria in complete-information [108] and incomplete-information [119, 109] games. Smooth auctions (Definitions 1.8 and 1.9) is a large class of auctions where the price of anarchy can be systematically characterized by the smooth arguments. Many interesting auctions have been shown to be smooth; and the smooth argument is a central proof technique to analyze the price of anarchy. We refer the reader to a recent survey [114] for more details.

The smoothness framework extends to adaptive dynamics with vanishing regret. However, several important dynamics are not guaranteed to have the vanishing regret property, for example the class of fictitious play [30] and other classes of dynamics in [62]. A research agenda, as raised in [114], is to characterize the performance of such dynamics.

Revenue Maximization. Optimal truthful auctions in single-parameter environments are completely characterized by Myerson [99]. A crucial assumption in the Myerson's construction of optimal auctions is the full knowledge of the distribution over bidder valuations. Recently, a major line of research in data-driven mechanism design focus on designing competitive auctions without the full knowledge on the valuation distribution and even in non-stochastic settings. The study of second-price auctions with reserve prices in single-parameter environments are of particular interest since the optimal auctions correspond to second-price auctions with reserve prices when knowledge on the distribution over valuations is available to the auctioneer. That problem and its variants have been considered in [82, 26, 43]. Recently, Roughgarden and Wang [113] gave a polynomial-time online algorithm that achieves $(1/2, O(\sqrt{T}))$ -regret. Subsequently, Dudik et al. [52] showed that the same regret bound can be obtained using their framework. Our approach differs from [113, 52] in that we consider the regularization approach whereas they follow the perturbation approaches.

Chapter 2

Online Primal-Dual Algorithms with Configuration Linear Programs

In this chapter, we present our contribution to online algorithms. We first give a primal-dual algorithm for the general class described in Section 1.4.2. Despite the simplicity of the algorithm, it has indeed various direct and indirect applications (shown in subsequent sections) with optimal bounds. We will go through them as a warm-up to present our approach. Subsequently, we give competitive algorithms for online problems with covering and packing constraints.

2.1 Primal-Dual Algorithm for A General Class of Problems

Recall that the problem consists of a set of resources \mathcal{E} and requests which arrive online. At the arrival of request i , a set of feasible strategies (actions) \mathcal{S}_i to satisfy request i is revealed. Each strategy $s_{ij} \in \mathcal{S}_i$ consists of a subset of resources in \mathcal{E} . Each resource e is associated to a non-negative non-decreasing *arbitrary* cost function $f_e : 2^{\mathcal{E}} \rightarrow \mathbb{R}^+$ and the cost induced by resource e depending on the set of requests using e . The cost of a solution is the total cost of resources, i.e., $\sum_e f_e(A_e)$ where A_e is the set of requests using resource e . The goal is to design an algorithm that upon the arrival of each request, selects a feasible strategy while maintaining the cost of the overall solution as small as possible.

2.1.1 A Greedy Algorithm

Formulation. We consider the formulation for the resource cost minimization problem following the configuration LP construction in [94]. We say that A is a *configuration* associated to resource e if A is a subset of requests using e . Let x_{ij} be a variable indicating whether request i selects strategy (action) $s_{ij} \in \mathcal{S}_i$. For configuration A and resource e , let z_{eA} be a variable such that $z_{eA} = 1$ if and only if for every request $i \in A$, $x_{ij} = 1$ for some strategy $s_{ij} \in \mathcal{S}_i$ such that $e \in s_{ij}$. In other words, $z_{eA} = 1$ iff the set of requests using e is exactly A . We consider the following formulation and the dual of its relaxation.

Primal:

$$\begin{aligned}
 & \min \sum_{e,A} f_e(A) z_{e,A} \\
 & \sum_{j:s_{ij} \in \mathcal{S}_i} x_{ij} = 1 \quad \forall i \\
 & \sum_{A:i \in A} z_{eA} = \sum_{j:e \in s_{ij}} x_{ij} \quad \forall i, e \\
 & \sum_A z_{eA} = 1 \quad \forall e \\
 & x_{ij}, z_{eA} \in \{0, 1\} \quad \forall i, j, e, A
 \end{aligned}$$

Dual:

$$\begin{aligned}
 & \max \sum_i \alpha_i + \sum_e \gamma_e \\
 & \alpha_i \leq \sum_{e:e \in s_{ij}} \beta_{ie} \quad \forall i, j \\
 & \gamma_e + \sum_{i \in A} \beta_{ie} \leq f_e(A) \quad \forall e, A
 \end{aligned}$$

In the primal, the first constraint guarantees that request i selects some strategy $s_{ij} \in \mathcal{S}_i$. The second constraint ensures that if request i selects strategy s_{ij} that contains resource e then in the solution, the set of requests using e must contain i . The third constraint says that in the solution, there is always a configuration associated to resource e .

Algorithm. We first interpret intuitively the dual variables, dual constraints and derive useful observations for a competitive algorithm. Variable α_i represents the increase of the total cost due to the arrival of request i . Variable $\beta_{i,e}$ stands for the marginal cost on resource e if request i uses e . By this interpretation, the first dual constraint clearly indicates the behaviour of an algorithm. That is, if a new request i is released, select a strategy $s_{ij} \in \mathcal{S}_i$ that minimizes the marginal increase of the total cost. Therefore, we deduce the following greedy algorithm.

Let A_e^* be the set of current requests using resource e . Initially, $A_e^* \leftarrow \emptyset$ for every e . At the arrival of request i , select strategy s_{ij}^* that is an optimal solution of

$$\min \sum_{e \in s_{ij}} \left[f_e(A_e^* \cup i) - f_e(A_e^*) \right] \quad \text{over } s_{ij} \in \mathcal{S}_i. \quad (2.1)$$

Although computational complexity is not a main issue for online problems, we notice that in many applications, the optimal solution for this mathematical program can be efficiently computed (for example when f_e 's are convex and \mathcal{S}_i can be represented succinctly in form of a polynomial-size polytope).

Dual variables. Assume that all resource cost f_e are (λ, μ) -smooth for some fixed parameters $\lambda > 0$ and $\mu < 1$. We are now constructing a dual feasible solution. Define α_i as $1/\lambda$ times the optimal value of the mathematical program (2.1). Informally, α_i is proportional to the increase of the total cost due to the arrival of request i . Note that this increase is also called *marginal cost* due to request i . For each resource e and request i , define

$$\beta_{i,e} := \frac{1}{\lambda} \left[f_e(A_{e, \prec i}^* \cup i) - f_e(A_{e, \prec i}^*) \right]$$

where $A_{e, \prec i}^*$ is the set of requests using resource e (due to the algorithm) prior to the arrival of i . In other words, β_{ij} equals $1/\lambda$ times the marginal cost of resource e if i uses e . Finally, for every resource e define the dual variable $\gamma_e := -\frac{\mu}{\lambda} f_e(A_e^*)$ where A_e^* is the set of all requests using e (at the end of the instance).

Lemma 2.1 *The dual variables defined as above are feasible.*

Proof The first dual constraint follows immediately from the definitions of $\alpha_i, \beta_{i,e}$ and the decisions by the algorithm. Specifically, the right-hand side of the constraint represents $1/\lambda$ times the increase cost if the request selects a strategy s_{ij} . This is larger than $1/\lambda$ times the minimum increase cost optimized over all strategies in \mathcal{S}_i , which is α_i .

We now show that the second constraint holds. Fix a resource e and a configuration A . The corresponding constraint reads

$$\begin{aligned} & -\frac{\mu}{\lambda}f_e(A_e^*) + \frac{1}{\lambda} \sum_{i \in A} \left[f_e(A_{e, \prec i}^* \cup i) - f_e(A_{e, \prec i}^*) \right] \leq f_e(A) \\ \Leftrightarrow & \sum_{i \in A} \left[f_e(A_{e, \prec i}^* \cup i) - f_e(A_{e, \prec i}^*) \right] \leq \lambda f_e(A) + \mu f_e(A_e^*). \end{aligned}$$

This inequality is due to the definition of (λ, μ) -smoothness for resource e . Hence, the second dual constraint follows. \square

Theorem 1.1 *Assume that all resource cost functions are (λ, μ) -smooth for some parameters $\lambda > 0, \mu < 1$. Then there exists a greedy $\frac{\lambda}{1-\mu}$ -competitive algorithm for the general problem.*

Proof By the definitions of dual variables, the dual objective is

$$\sum_i \alpha_i + \sum_e \gamma_e = \sum_e \frac{1}{\lambda} f_e(A_e^*) - \sum_e \frac{\mu}{\lambda} f_e(A_e^*) = \frac{1-\mu}{\lambda} \sum_e f_e(A_e^*)$$

Besides, the cost of the solution due to the algorithm is $\sum_e f_e(A_e^*)$. Hence, the competitive ratio is at most $\lambda/(1-\mu)$. \square

2.1.2 Applications

Theorem 1.1 yields simple algorithm with *optimal* competitive ratios for several problems as shown in the following sections. Among others, we give optimal algorithms for energy efficient scheduling problems (in unrelated machine environment) and the facility location with client-dependent cost problem. Prior to our work, no competitive algorithm has been known for the problems. The proofs are now reduced to computing smooth parameters λ, μ that subsequently imply the competitive ratios. We mainly use the following smooth inequality, developed in [45], to derive the explicit competitive bounds in case of non-negative polynomial cost functions. (For completeness, the proof can be found in Appendix A).

Lemma 2.2 *For any sequences of non-negative real numbers $\{a_1, a_2, \dots, a_n\}$ and $\{b_1, b_2, \dots, b_n\}$ and for any polynomial g of degree k with non-negative coefficients, it holds that*

$$\sum_{i=1}^n \left[g\left(b_i + \sum_{j=1}^i a_j\right) - g\left(\sum_{j=1}^i a_j\right) \right] \leq \lambda(k) \cdot g\left(\sum_{i=1}^n b_i\right) + \mu(k) \cdot g\left(\sum_{i=1}^n a_i\right)$$

where $\mu(k) = \frac{k-1}{k}$ and $\lambda(k) = \Theta(k^{k-1})$. The same inequality holds for $\mu(k) = \frac{k-1}{k \ln k}$ and $\lambda(k) = \Theta((k \ln k)^{k-1})$.

2.1.2.1 Minimum Power Survival Network Routing

Problem. In the problem, we are given a graph $G(V, E)$ and requests arrive online. The demand of a request i is specified by a source $s_i \in V$, a sink $t_i \in V$, the load vector $p_{i,e}$ for every edge (link) $e \in E$ and an integer number k_i . At the arrival of request i , one needs to choose k_i edge-disjoint paths connecting s_i to t_i . Request i increases the load $p_{i,e}$ for each edge e used to satisfy its demand. The load ℓ_e of an edge e is defined as the total load incurred by the requests using e . The power cost of edge e with load ℓ_e is $f_e(\ell_e)$. The objective is to minimize the total power $\sum_e f_e(\ell_e)$. Typically $f_e(\ell_e) = c_e \ell_e^{\alpha_e}$ where c_e and α_e are parameters depending on e .

This problem generalizes the MINIMUM POWER ROUTING problem — a variant in which $k_i = 1$ and $p_{i,e} = 1 \forall i, e$ — and the LOAD BALANCING problem — a variant in which $k_i = 1$, all the sources (sinks) are the same $s_i = s_{i'} \forall i, i'$ ($t_i = t_{i'} \forall i, i'$) and every $s_i - t_i$ path has length 2. For the MINIMUM POWER ROUTING in offline setting, Andrews et al. [5] gave a polynomial-time poly-log-approximation algorithm. The result has been improved by Makarychev and Sviridenko [94] who gave an B_α -approximation algorithm. In the online setting, Gupta et al. [67] presented an α^α -competitive online algorithm. For the LOAD BALANCING problem, the currently best-known approximation is B_α due to [94] via their rounding technique based on decoupling inequality. In the online setting, it has been shown that the optimal competitive ratio for the LOAD BALANCING problem is $\Theta(\alpha^\alpha)$ [40].

Contribution. In the problem, a strategy for each request i consists of k_i edge-disjoint paths connecting s_i and t_i . Applying the general framework, we deduce the following greedy algorithm.

Let ℓ_e be the load of edge e . Initially, set $\ell_e \leftarrow 0$ for every edge e . At the arrival of request i , compute a strategy consisting of k_i edge-disjoint paths from s_i and t_i such that the increase of the total cost is minimum. Select this strategy for request i and update ℓ_e .

We notice that computing the strategy for request i can be done efficiently. Given the current loads ℓ_e on every edge e , create a graph H consisting of the same vertices and edges as graph G . For each edge e in graph H , define the capacity to be 1 and the cost on e to be $f_e(p_{i,e} + \ell_e) - f_e(\ell_e)$. Then the computing k_i edge-disjoint paths from s_i and t_i with the minimal marginal cost in G is equivalent to solving a transportation problem in graph H .

Proposition 2.1 *If the congestion costs of all edges are (λ, μ) -smooth then the algorithm is $\lambda / (1 - \mu)$ -competitive. In particular, if $f_e(z) = z^{\alpha_e}$ then the algorithm is $O(\alpha^\alpha)$ -competitive where $\alpha = \max_e \alpha_e$.*

Proof The proposition follows directly from Theorem 1.1 and the particular case is derived additionally by Lemma 2.2. \square

2.1.2.2 Online Vector Scheduling

Problem. In the problem, there are m unrelated machines and jobs arrive online. The load of a job j in machine i is specified by a vector $p_{ij} = \langle p_{ij}(k) : 1 \leq k \leq d \rangle$ where $p_{ij}(k) \geq 0$ and d , a fixed parameter, is the dimension of the vector. At the arrival of a job j , vectors p_{ij} for all i are revealed and job j must be assigned immediately to a machine. Given a job-machine assignment σ , the load in dimension k of machine i is defined as $\ell_{i,\sigma}(k) := \sum_{j:\sigma(j)=i} p_{ij}(k)$ for $1 \leq k \leq d$. The L_α -norm for

$\alpha \geq 1$ in dimension k is $\|\Lambda_\sigma(k)\|_\alpha := \left(\sum_{i=1}^m \ell_{i,\sigma}(k)^\alpha\right)^{1/\alpha}$; and the L_∞ -norm (makespan norm) in dimension k is $\|\Lambda_\sigma(k)\|_\infty := \max_{i=1}^m \ell_{i,\sigma}(k)$. In the L_α -norm, the objective is to find an online assignment σ minimizing $\max_k \|\Lambda_\sigma(k)\|_\alpha$. In the L_∞ -norm, the objective is to find an online assignment σ minimizing $\max_k \|\Lambda_\sigma(k)\|_\infty$. An algorithm is r -competitive for the L_α -norm if it outputs an assignment σ such that for any assignment σ^* , it holds that $\max_k \|\Lambda_\sigma(k)\|_\alpha \leq r \cdot \max_k \|\Lambda_{\sigma^*}(k)\|_\alpha$.

The online vector scheduling is introduced by Chekuri and Khanna [44]. Recently, Im et al. [79] showed an optimal competitive algorithm for this problem. Their analysis is based on a carefully constructed potential function. In the following, we can also derive an optimal algorithm for this problem based on our general framework. The analysis is much simpler and follows directly from Theorem 1.1.

Contribution. In the problem, the set of strategies of a job j is the set of all machines. Applying the general framework, we deduce the following greedy algorithm.

In the L_α -norm objective for $1 \leq \alpha < \infty$, consider the cost function

$$C(\sigma) := \sum_{k=1}^d \left(\sum_{i=1}^m \ell_{i,\sigma}(k)^\alpha \right)^{\frac{\alpha+\log d}{\alpha}}$$

where σ is a job-machine assignment of all jobs released so far.

In the L_∞ -norm objective, consider the cost function

$$C(\sigma) := \sum_{k=1}^d \sum_{i=1}^m \ell_i(k)^{\log m + \log d}.$$

Initially, σ is an empty assignment. At the arrival of j , assign j to machine i^* that minimizes the increase of $C(\sigma)$. Again, the assignment of a job j can be efficiently computed.

Proposition 2.2 ([79]) *For the L_α -norm objective where $\alpha < \infty$, the algorithm is $O(\max\{\alpha, \log d\})$ -competitive. For the L_∞ -norm objective, the algorithm is $O(\log d + \log m)$ -competitive.*

Proof Let σ^* is an optimal assignment for the L_α -norm objective. We have

$$\begin{aligned} \left(\sum_{i=1}^m \ell_{i,\sigma}(k)^\alpha \right)^{\frac{\alpha+\log d}{\alpha}} &\leq \sum_{k=1}^d \left(\sum_{i=1}^m \ell_{i,\sigma}(k)^\alpha \right)^{\frac{\alpha+\log d}{\alpha}} \\ &\leq O(\alpha + \log d)^{\alpha+\log d} \cdot \sum_{k=1}^d \left(\sum_{i=1}^m \ell_{i,\sigma^*}(k)^\alpha \right)^{\frac{\alpha+\log d}{\alpha}} \\ &\leq O(\alpha + \log d)^{\alpha+\log d} \cdot d \cdot \max_{k=1}^d \left(\sum_{i=1}^m \ell_{i,\sigma^*}(k)^\alpha \right)^{\frac{\alpha+\log d}{\alpha}} \end{aligned}$$

In these inequalities, we apply Theorem 1.1 and Lemma 2.2 (note that $C(\sigma)$ is a polynomial of degree $(\alpha + \log d)$). Taking the $(\alpha + \log d)$ th root, the result for L_α -norm objective follows.

For the L_∞ -norm, similarly we have

$$\begin{aligned} \max_{k=1}^d \max_{i=1}^m \ell_{i,\sigma}(k)^{\log m + \log d} &\leq \sum_{k=1}^d \sum_{i=1}^m \ell_{i,\sigma}(k)^{\log m + \log d} \\ &\leq (\log m + \log d)^{\log m + \log d} \cdot \sum_{k=1}^d \sum_{i=1}^m \ell_{i,\sigma^*}(k)^{\log m + \log d} \\ &\leq (\log m + \log d)^{\log m + \log d} \cdot d \cdot m \cdot \max_{k=1}^d \max_{i=1}^m \ell_{i,\sigma^*}(k)^{\log m + \log d}. \end{aligned}$$

Again, taking the $(\log m + \log d)^{\text{th}}$ root, the proposition follows. \square

2.1.2.3 Online Energy-Efficient Scheduling

Problem. Energy-efficient algorithms have received considerable attention and have been widely studied in scheduling. One main direction is to design algorithms towards a more realistic setting — online multiple machine setting [2]. We consider the following energy minimization problem. In the problem, we are given m unrelated machines and a set of jobs. Each job j is specified by its released date r_j , deadline d_j and processing volumes p_{ij} if job j is processed in machine i . We consider *non-migration* schedules; that is, every job j has to be assigned to exactly one machine and is fully processed in that machine during time interval $[r_j, d_j]$. However, jobs can be executed *preemptively*, meaning that a job can be interrupted during its execution and can be resumed later on the same machine. An algorithm can choose appropriate speed $s_i(t)$ for every machine i at any time t in order to complete all jobs. Every machine i has a non-decreasing energy power function $P_i(s_i(t))$ depending on the speed $s_i(t)$. Typically, $P_i(z)$ has form z^{α_i} for constant $\alpha_i \geq 1$ or in a more general context, $P_i(z)$ is assumed to be convex. In the problem, we consider general non-decreasing continuous functions P_i *without* convexity assumption. The objective is to minimize the total energy consumption while completing all jobs. In the *online* setting, jobs arrive over time and the assignment and the schedule have to be done irrevocably.

In the offline setting, for the typical energy function $P(z) = z^\alpha$, the best known algorithms [66, 13] have competitive ratio $O(B_\alpha)$ where B_α is the Bell number. Prior to our work, competitive algorithms for this online problem are known only in the single machine setting and the energy power function $P(z) = z^\alpha$. Specifically, Bansal et al. [15] gave a $2\left(\frac{\alpha}{\alpha-1}\right)^\alpha e^\alpha$ -competitive algorithm. In terms of lower bounds, Bansal et al. [18] showed that no deterministic algorithm has competitive ratio less than $e^{\alpha-1}/\alpha$ for single machine. For unrelated machines, the lower bound $\Omega(\alpha^\alpha)$ follows the construction of Caragiannis [40] for LOAD BALANCING (with L_α -norm) (by considering all jobs have the same span $[r_j, d_j] = [0, 1]$). Kling and Pietrzyk [83] gave a $O(\alpha^\alpha)$ -competitive algorithm in the multi-identical-processor setting in which job migration is allowed. Surprisingly, no competitive algorithm is known in the non-migratory multiple-machine environment, that is in contrast to the similar online problem with objective as the total energy plus flow-time [3]. The main difference here is that for the latter, one can make a tradeoff between energy and flow-time and derive a competitive algorithm whereas for the former, one has to deal directly with a non-linear objective and no LP with relatively small integrality gap was known. We notice that Gupta et al. [67] gave also a primal-dual competitive algorithm for the single machine environment. However, their approach cannot be used for unrelated machines due to the large integrality gap of the formulation.

Contribution. In this problem, the speed of a job can be an arbitrary (non-negative) real number. However, in order to employ tools from linear programming, we consider a discretization of speeds at the price of a small loss in the competitive ratio. Fix an arbitrary constant $\epsilon > 0$ and $\delta > 0$. Define the set of speeds $\mathcal{V} = \{\ell \cdot \epsilon : 0 \leq \ell \leq L\}$ for some sufficiently large L . During a time interval $[t, t + \delta]$, a job can be executed at a speed in \mathcal{V} . As the energy cost functions are continuous, this assumption on the setting worsens the energy cost by at most a factor $(1 + \tilde{\epsilon})$ for arbitrarily small $\tilde{\epsilon}$. Given a job j , the set of feasible strategies \mathcal{S}_j of j contains all feasible non-migratory executions of j on some machine. Specifically, a strategy of job j can be described as the union over all machines i of solutions determined by the following program:

$$\sum_{t=r_j}^{d_j} \delta \cdot v_{ijt} \geq p_{ij} \quad \text{s.t.} \quad v_{ijt} \in \mathcal{V},$$

where in the sum we increment each time t by δ . Here, v_{ijt} stands for the speed of executing job j in machine i at time t .

Applying the general framework, we derive the following algorithm.

Algorithm. Assume that the energy functions are (λ, μ) -smooth. Let u_{it} be the speed of machine i at time t . Initially, set $u_{it} \leftarrow 0$ for every machine i and time t . At the arrival of a job j , compute the minimum energy increase if job j is assigned to machine i and define β_{ij} to be equal to $1/\lambda$ times that minimum energy. It is indeed an optimization problem

$$\min \sum_{r_j}^{d_j} \delta \cdot \left[P_i(u_{it} + v_{ijt}) - P_i(u_{it}) \right] \quad \text{s.t.} \quad \sum_{r_j}^{d_j} \delta \cdot v_{ijt} \geq p_{ij}, \quad v_{ijt} \in \mathcal{V} \quad (2.2)$$

Observe that if P_i is a convex function then it is a convex program and can be solved efficiently. In this case, using the KKT conditions, the optimal solution can be constructed as follows. We initiate a variable v_{ijt} as 0. While $\sum_{r_j}^{d_j} \delta \cdot v_{ijt} < p_{ij}$, i.e., the total volume of job j has not been completed, continue increasing v_{ijt} at $\arg \min_{r_j \leq t \leq d_j} u_{it} + v_{ijt}$. Note that this is exactly algorithm OA in [15] for a single machine. Let $v_{i^*jt}^*$ be an optimal solution to the mathematical program (2.2). Then, assign job j to machine $i^* \in \arg \min_i \beta_{ij}$ and execute j at time t with speed $v_{i^*jt}^*$.

Proposition 2.3 *If the energy cost functions are (λ, μ) -smooth then the algorithm is $(1 + \epsilon)\lambda / (1 - \mu)$ -competitive for arbitrarily small ϵ . In particular, if $P_i(z) = z^{\alpha_i}$ then the algorithm is $(1 + \epsilon)O(\alpha^\alpha)$ -competitive where $\alpha = \max_i \alpha_i$.*

Proof The proposition follows directly from Theorem 1.1. The factor $(1 + \epsilon)$ is due to the discretization of the speeds. In the particular case $P_i(z) = z^{\alpha_i}$, the functions are (λ, μ) -smooth with $\mu = (\alpha - 1)/\alpha$ and $\lambda = O(\alpha^{\alpha-1})$ by Lemma 2.2. The competitive ratio of this case follows. \square

2.1.2.4 Online Prize Collecting Energy-Efficient Scheduling

Problem. We consider the same setting as in the ENERGY MINIMIZATION problem. Additionally, each job j has a penalty π_j . There is no penalty from job j if it is completely executed during $[r_j, d_j]$ in some machine i . Otherwise, if job j is not

completed (even if most volume of job j have been executed) then the algorithm has to pay a penalty π_j . The objective is to minimize the total penalty of uncompleted jobs plus the energy cost.

Contribution. The result does not follow immediately from Theorem 1.1 but the approach is exactly the one for the general framework.

By the same formulation from the previous section, assume that the set of speeds is finite and discrete. The set of feasible strategies \mathcal{S}_j of a job j are all the feasible non-migratory executions of a job j on some machine as defined in the previous section. The sets \mathcal{S}_j 's are also finite and discrete. We say that A is a *configuration* of machine i if it is a schedule of a subset of jobs in i . Specifically, a configuration A consists of tuples (i, j, k) specifying that a job j is assigned to machine i and is executed according to strategy $s_{ijk} \in \mathcal{S}_j$.

We are now formulating a configuration LP for the problem. Let x_{ijk} be a variable indicating whether job j is processed in machine i according to strategy $s_{ijk} \in \mathcal{S}_j$. For configuration A and machine i , let z_{iA} be a variable such that $z_{iA} = 1$ if and only if $x_{ijk} = 1$ for every $(i, j, k) \in A$. In other words, $z_{iA} = 1$ iff A is the solution of the problem restricted on machine i . Let $c_{i,A}$ be the energy cost of configuration A in machine i . We consider the following formulation.

$$\begin{aligned}
 \min \quad & \sum_{i,A} c_{i,A} z_{iA} + \sum_j \left(1 - \sum_{i,k} x_{ijk} \right) \pi_j \\
 & \sum_{i,k} x_{ijk} \leq 1 && \forall j \\
 & \sum_{A:(i,j,k) \in A} z_{iA} = x_{ijk} && \forall i, j, k \\
 & \sum_A z_{iA} = 1 && \forall i \\
 & x_{ijk}, z_{iA} \in \{0, 1\} && \forall i, j, A
 \end{aligned}$$

The first constraint guarantees that a job j can be assigned to at most one machine i and be executed according to at most one feasible strategy. The second constraint ensures that if job j is assigned to machine i and is executed according to strategy $s_{ijk} \in \mathcal{S}_j$ then the configuration corresponding to the solution restricted on machine i must contain (i, j, k) . The third constraint says that there is always a configuration associated to machine i for every i . The dual of the relaxation reads

$$\begin{aligned}
 \max \quad & \sum_j (\pi_j - \alpha_j) + \sum_i \gamma_i \\
 & \alpha_j + \beta_{ijk} \geq \pi_j && \forall i, j, k \\
 & \gamma_i + \sum_{(i,j,k) \in A} \beta_{ijk} \leq c_{iA} && \forall i, A \\
 & \alpha_j \geq 0 && \forall j
 \end{aligned}$$

Greedy Algorithm. Assume that all energy power functions are (λ, μ) -smooth for some fixed parameters λ and μ . At the arrival of job j , compute the minimum energy increase if j is assigned to some machine i . If the minimum energy increase is larger

than $\lambda \cdot \pi_j$ then reject the job. Otherwise, assign and execute j such that the energy increase is minimum.

Proposition 2.4 *Assume that all energy power functions are (λ, μ) -smooth. Then the algorithm is $\lambda/(1 - \mu)$ -competitive.*

Proof We define the dual variables similarly as in the general framework. Let $A_{i, \prec j}^*$ be the configuration of machine i (due to the algorithm) before the arrival of job j . (Initially, $A_{i, \prec 1}^* \leftarrow \emptyset$ for every machine i .) For each machine i and a strategy $s_{jk} \in \mathcal{S}_j$ such that s_{jk} is a schedule of j in machine i , define

$$\beta_{ijk} = \frac{1}{\lambda} \left[c_i(A_{i, \prec j}^* \cup s_{jk}) - c_i(A_{i, \prec j}^*) \right].$$

If s_{jk} is not a schedule of j in machine i then define $\beta_{ijk} = \infty$. Moreover, define

$$\alpha_j = \max \{ \pi_j - \min_{i,k} \beta_{ijk}, 0 \} \quad \text{and} \quad \gamma_i = \frac{\mu}{\lambda} c_i(A_i^*)$$

where A_i^* is the configuration of machine i at the end of the instance (when all jobs have been released).

The variables constitute a dual feasible solution. The first dual constraint follows the definition of α_j . The second dual constraint follows the definition of (λ, μ) -smoothness. Note that that for any configuration A of a machine i (a feasible schedule in machine i), if $(i, j, k) \in A$ then by definition of dual variables, $\beta_{ijk} \neq \infty$.

We are now bounding the dual. The algorithm has the property *immediate-reject*. It means that if the algorithm accepts a job then the job will be completed; and otherwise, the job is rejected at its arrival. By the algorithm, $\alpha_j = 0$ for every rejected job j . Besides, if job j is accepted then $\pi_j - \alpha_j = \beta_{ijk}$ where i is the machine to which job j is assigned and job j is executed according to strategy s_{jk} . Therefore, by the definition of dual variables, $\sum_j (\pi_j - \alpha_j)$, where the sum is taken over accepted jobs j , equals $1/\lambda$ times the total energy consumption. Recall that the total energy consumption of the algorithm is $\sum_i c_i(A_i^*)$. The dual objective is

$$\sum_j (\pi_j - \alpha_j) + \sum_i \gamma_i = \sum_{j: j \text{ rejected}} \pi_j + \frac{1}{\lambda} \sum_i c_i(A_i^*) - \frac{\mu}{\lambda} \sum_i c_i(A_i^*)$$

Moreover, the primal is equal to the total penalty of rejected jobs plus $\sum_i c_i(A_i^*)$. Therefore, the ratio between primal and dual is at most $\lambda/(1 - \mu)$. \square

2.1.2.5 Facility Location with Client-Dependent Facility Cost

Non-Convex Facility Location. In the problem, we are given a metric space (M, d) and clients arrive online. Let N be the set of clients and $n = |N|$. A location $i \in M$ is characterized by a fixed opening cost a_i and an arbitrary non-decreasing serving cost function $f_i : 2^N \rightarrow \mathbb{R}^+$. If a subset S of clients is served by a facility at location i then the facility cost at this location is $a_i + f_i(S)$. At the arrival of a client, an algorithm needs to assign the client to some facility. The goal is to minimize the total cost, which is the total distance from clients to their facilities plus the total facility cost.

Facility Location is one of the most widely studied problems. In the classic version, the facility cost consists only of the opening cost. There is a large literature in

the offline setting. In online setting, Meyerson [97] gave a randomized $O(\frac{\log n}{\log \log n})$ -competitive algorithm. This competitive ratio matches to the randomized lower bound due to Fotakis [59]. For deterministic algorithms, Fotakis [58] first presented a primal-dual $O(\log n)$ -competitive algorithm and subsequently improved it to an optimal $O(\frac{\log n}{\log \log n})$ -competitive algorithm [59]. The online capacitated facility location in which function $f_i(S) = 0$ if $|S| \leq u_i$ for some capacity u_i and $f_i(S) = \infty$ otherwise has been studied in [9]. Using a primal-dual framework for mixed packing and covering constraints, the authors derived a $O(\log m \log mn)$ -competitive algorithm.

Contribution. We derive a competitive algorithm by combining the primal-dual algorithm due to Fotakis [58] for the online (classic) facility location and our primal-dual framework for non-convex functions.

Let x_{ij} and y_i be variables indicating whether client j is assigned to facility i and whether facility i is open, respectively. For subset $S \subset N$, let $z_{i,S}$ be a variable such that $z_{i,S} = 1$ if and only if $x_{ij} = 1$ for every client $j \in S$, and $x_e = 0$ for $j \notin S$. We consider the following formulation and the dual of its relaxation.

Primal:

$$\begin{aligned} \min \quad & \sum_i a_i y_i + \sum_{i,j} d_{ij} x_{ij} + \sum_{i,S} f_i(S) z_{i,S} \\ & \sum_i x_{ij} \geq 1 \quad \forall j \\ & y_i \geq x_{ij} \quad \forall i, j \\ & \sum_{S:j \in S} z_{i,S} = x_{ij} \quad \forall i, j \\ & \sum_S z_{i,S} = 1 \quad \forall i \\ & x_{ij}, z_{i,S} \in \{0, 1\} \quad \forall i, j, S \end{aligned}$$

Dual:

$$\begin{aligned} \max \quad & \sum_j \alpha_j + \sum_i \theta_i \\ & \alpha_j \geq d_{ij} + \beta_{ij} + \gamma_{ij} \quad \forall i, j \\ & \sum_j \beta_{ij} \leq a_i \quad \forall i \\ & \theta_i + \sum_{j \in S} \gamma_{ij} \leq f_i(S) \quad \forall i, S \\ & \alpha_j, \beta_{ij} \geq 0 \quad \forall i, j \end{aligned}$$

Algorithm. Assume that all serving cost f_i are (λ, μ) -smooth for some parameters λ and μ . Intuitively, β_{ij} and γ_{ij} can be interpreted as the contributions of client j to the opening cost and the serving cost at location i . At the arrival of client j , continuously increase α_j . For any facility such that $\alpha_j = d_{ij}$, start increasing β_{ij} . If $\sum_j \beta_{ij} = a_i$ then stop increasing β_{ij} and start increasing γ_{ij} until $\frac{\mu}{\lambda} [f_i(S \cup j) - f_i(S)]$ where S is the current set of clients assigned to i . Assign j to the first facility i such that $\gamma_{ij} = \frac{\mu}{\lambda} [f_i(S \cup j) - f_i(S)]$, i.e., $x_{ij} = 1$. Open i (i.e., $y_i = 1$) if it has not been opened.

Proposition 2.5 *Assume that all serving cost f_i are (λ, μ) -smooth. Then the algorithm is $O(\log n + \frac{\lambda}{1-\mu})$ -competitive.*

Proof We define dual variables similarly as in Theorem 1.1. The α -variables, β -variables and γ -variables are defined in the algorithm. Define θ_i equal $-1/\lambda$ times the (final) serving cost at facility i . Let $\pi(j)$ be the facility to which j is assigned and $\pi(N)$ the set of facilities opened by the algorithm.

The dual variables constitute a feasible solution. The first and second dual constraints are due to the algorithm. Note that by the definition of γ -variables, it always holds that $\gamma_{ij} \leq \frac{\mu}{\lambda} [f_i(S \cup j) - f_i(S)]$ where S is the set of clients assigned to i before

the arrival of j . The last constraint follows the (λ, μ) -smoothness of serving costs. We are now bounding the primal and the dual. We have

$$\begin{aligned} \sum_{i \in \pi(N)} a_i + \sum_j d_{\pi(j),j} &\leq O(\log n) \sum_j \left(\alpha_j - \gamma_{\pi(j),j} \right) \\ \sum_i f_i(\pi^{-1}(i)) &\leq \frac{\lambda}{1-\mu} \left(\sum_j \gamma_{\pi(j),j} + \sum_i \theta_i \right) \end{aligned}$$

where the first inequality is due to Fotakis [58] and the second one follows the definition of dual variables. The proposition follows. \square

2.2 Primal-Dual Framework for 0 – 1 Covering Problems

Consider the following integer optimization problem. Let \mathcal{E} be a set of n resources and let $f : \{0,1\}^n \rightarrow \mathbb{R}^+$ be a monotone cost function. Let $x_e \in \{0,1\}$ be a variable indicating whether resource e is selected. The problem is to minimize $f(\mathbf{x})$ subject to covering constraints $\sum_e a_{i,e} x_e \geq 1$ for every constraint i and $x_e \in \{0,1\}$ for every e . In the online setting, the constraints are revealed one-by-one and at any step, one needs to maintain a feasible integer solution \mathbf{x} .

2.2.1 Algorithm for Fractional Covering

Recall that a differentiable function $F : [0,1]^n \rightarrow \mathbb{R}^+$ is (λ, μ) -min-locally-smooth if for any set $S \subset \mathcal{E}$, and for any vectors $\mathbf{x}^e \in [0,1]^n$ where $e \in \mathcal{E}$, the following inequality holds:

$$\sum_{e \in S} \nabla_e F(\mathbf{x}^e) \leq \lambda F(\mathbf{1}_S) + \mu F(\mathbf{x})$$

where $\mathbf{x} := \bigvee_{e \in S} \mathbf{x}^e$, meaning that $x_{e'} = \max_e \{x_{e'}^e\}$ for any coordinate e' .

Formulation. We say that $S \subset \mathcal{E}$ is a *configuration* if $\mathbf{1}_S$ corresponds to a feasible solution. Let x_e be a variable indicating whether the resource e is used. For configuration S , let z_S be a variable such that $z_S = 1$ if and only if $x_e = 1$ for every resource $e \in S$, and $x_e = 0$ for $e \notin S$. In other words, $z_S = 1$ iff $\mathbf{1}_S$ is the selected solution to the problem. For any subset $A \subset \mathcal{E}$, define $c_{i,A} = \max\{1 - \sum_{e' \in A} a_{i,e'}; 0\}$ and $a_{i,e,A} := \min\{a_{i,e}; c_{i,A}\}$. Denote $b_{i,e,A} = \frac{a_{i,e,A}}{c_{i,A}}$ where $c_{i,A} > 0$. We consider the following formulation and the dual of its relaxation.

Primal:

$$\begin{aligned} \min \sum_S f(\mathbf{1}_S) z_S \\ \sum_{e \notin A} b_{i,e,A} \cdot x_e &\geq 1 & \forall i, A \subset \mathcal{E} \\ \sum_{S: e \in S} z_S &= x_e & \forall e \\ \sum_S z_S &= 1 \\ x_e, z_S &\in \{0,1\} & \forall e, S \end{aligned}$$

Dual:

$$\begin{aligned} \max \sum_{i,A} \alpha_{i,A} + \gamma \\ \sum_i \sum_{A: e \notin A} b_{i,e,A} \cdot \alpha_{i,A} &\leq \beta_e & \forall e \\ \gamma + \sum_{e \in S} \beta_e &\leq f(\mathbf{1}_S) & \forall S \\ \alpha_i &\geq 0 & \forall i \end{aligned}$$

In the primal, the first constraints are knapsack-constraints of the form $\sum_{e \notin A} a_{i,e,A} \cdot x_e \geq c_{i,A}$ corresponding to the given polytope. Note that it is sufficient to consider only constraints with $c_{i,A} > 0$. The second constraint ensures that if a resource e is chosen then the selected solution must contain e . The third constraint says that one solution (configuration) must be selected.

Algorithm. Assume that function $F(\cdot)$ is $(\lambda, \frac{\mu}{4 \ln(1+2d^2)})$ -min-locally smooth. Let d be the maximal number of positive entries in a row, i.e., $d = \max_i |\{a_{ie} : a_{ie} > 0\}|$. Denote $\nabla_e F(\mathbf{x}) = \partial F(\mathbf{x}) / \partial x_e$. Consider the following Algorithm 1 which follows the scheme in [11] with some more subtle steps due to the non-monotone behavior of the gradient. In the algorithm, the current dual variable α increases at constant rate (Step 6) and the update of dual variables β 's is shown in Step 8. If the gradient $\nabla_e F(\mathbf{x})$ at coordinate e is monotone then β_e is set to be $\frac{1}{\lambda} \nabla_e F(\mathbf{x})$. However, in case $\nabla_e F(\mathbf{x})$ decreases, the value of β_e is kept unchanged. The primal update rule follows a multiplicative increase where the increasing rate of x_e is inversely proportional to β_e (Step 9). Finally, using the same idea as in [11], some dual variables α will be decreased in order to maintain the feasibility of our dual solution.

Algorithm 1 Algorithm for Covering Constraints.

- 1: Initially, set $A^* \leftarrow \emptyset$. Intuitively, A^* consists of all resources e such that $x_e = 1$.
- 2: All primal and dual variables are initially set to 0.
- 3: At every step, always maintain $z_S = \prod_{e \in S} x_e \prod_{e \notin S} (1 - x_e)$.
- 4: Upon the arrival of primal constraint $\sum_e a_{k,e} x_e \geq 1$ and the new corresponding dual variable α_k .
- 5: **while** $\sum_{e \notin A^*} b_{k,e,A^*} x_e < 1$ **do** # Increase primal, dual variables
- 6: Increase τ at rate 1 and increase α_{k,A^*} at rate $\frac{1}{\lambda \cdot \ln(1+2d^2)}$.
- 7: **for** $e \notin A^*$ such that $b_{k,e,A^*} > 0$ **do**
- 8: **if** $\beta_e < \frac{1}{\lambda} \nabla_e F(\mathbf{x})$ **then** $\beta_e \leftarrow \frac{1}{\lambda} \nabla_e F(\mathbf{x})$
- 9: Increase x_e according to the following function

$$\frac{\partial x_e}{\partial \tau} \leftarrow \frac{b_{k,e,A^*} \cdot x_e + 1/d}{\lambda \cdot \beta_e}$$

- 10: **end for**
- 11: **if** $x_e = 1$ **then** update $A^* \leftarrow A^* \cup \{e\}$.
- 12: **for** $e \notin A^*$ such that

$$\sum_{i=1}^k \sum_{A: e \notin A} b_{i,e,A} \cdot \alpha_{i,A} \geq \beta_e$$

- 13: **do** # Decrease dual variables
 - 14: Let $m_e^* \leftarrow \arg \max \{b_{i,e,A} \mid \forall A : e \notin A, \forall 1 \leq i \leq k : \alpha_{i,A} > 0\}$.
 - 15: Increase $\alpha_{m_e^*,A}$ continuously at rate $-\frac{b_{k,e,A^*}}{b_{m_e^*,e,A}} \cdot \frac{1}{\lambda \cdot \ln(1+2d^2)}$.
 - 16: **end for**
 - 17: **end while**
-

Dual variables. Variables $\alpha_{i,A}$ and β_e are constructed in the algorithm. Let \mathbf{x} be the current solution of the algorithm. Define $\gamma = -\frac{\mu}{4\lambda \cdot \ln(1+2d^2)} F(\mathbf{x})$. Note that due to the algorithm, $\beta_e \geq \frac{1}{\lambda} \cdot \nabla_e F(\mathbf{x})$.

The following lemma gives a lower bound on x -variables. Remark that the monotonicity of the gradient is crucial in the analysis of [11], in particular to prove the bounds on x -variables. However, by our approach the gradient monotonicity is not needed.

Lemma 2.3 *Let e be an arbitrary resource. At any moment during the execution of the algorithm where the k^{th} request has been released, it always holds that*

$$x_e \geq \frac{1}{\max b_{i,e,A} \cdot d} \left[\exp\left(\frac{\ln(1+2d^2)}{\beta_e} \cdot \sum_{A:e \notin A} \sum_i b_{i,e,A} \cdot \alpha_{i,A}\right) - 1 \right]$$

where denote $\max b_{i,e,A} := \max\{b_{i,e,A} > 0 \mid \forall A : e \notin A, \forall 1 \leq i \leq k : \alpha_{i,A} > 0\}$.

Proof Fix a resource e . We prove the lemma by induction. At the beginning of the instance, while no request has been released yet, both sides of the lemma are 0. Assume that the lemma holds until the arrival of the k^{th} request. Consider a moment τ and let A^* be the current set of resources e' such that $x_{e'} = 1$. If at time τ , $x_e = 1$ then by the algorithm, the set A^* has been updated so that $e \in A^*$. The increasing rates of both sides in the lemma inequality are 0. In the remaining, assume that $x_e < 1$. Recall that by the algorithm, $\beta_e \geq \frac{1}{\lambda} \nabla_e F(\mathbf{x})$. We consider two cases: $\beta_e > \frac{1}{\lambda} \nabla_e F(\mathbf{x})$ and $\beta_e = \frac{1}{\lambda} \nabla_e F(\mathbf{x})$.

Case 1: $\beta_e > \frac{1}{\lambda} \nabla_e F(\mathbf{x})$. In this case, by the algorithm, the value of β_e remains unchanged at time τ . Hence, the derivative of the right hand side of the lemma inequality according to τ is

$$\begin{aligned} & \sum_i \frac{\partial \alpha_{i,A^*}}{\partial \tau} \cdot \frac{b_{i,e,A^*}}{\max b_{i,e,A} \cdot d} \cdot \frac{\ln(1+2d^2)}{\beta_e} \cdot \exp\left(\frac{\ln(1+2d^2)}{\beta_e} \cdot \sum_{A:e \notin A} \sum_i b_{i,e,A} \alpha_{i,A}\right) \\ & \leq \frac{b_{k,e,A^*} \cdot x_e + 1/d}{\lambda \cdot \beta_e} = \frac{\partial x_e}{\partial \tau} \end{aligned}$$

In the inequality, we use the induction hypothesis; $\frac{\partial \alpha_{k,A^*}}{\partial \tau} > 0$ and $\frac{\partial \alpha_{i,A^*}}{\partial \tau} \leq 0$ for $i \neq k$ and $\frac{\partial \beta_e}{\partial \tau} = 0$; and the increasing rate of α_{k,A^*} according to the algorithm. So the rate in the left-hand side is always larger than that in the right-hand side. Moreover, at some steps in the algorithm, α -variables might be decreased while the x -variables are maintained monotone. Hence, the lemma inequality holds.

Case 2: $\beta_e = \frac{1}{\lambda} \nabla_e F(\mathbf{x})$. In this case, by the algorithm, $\frac{1}{\lambda} \nabla_e F(\mathbf{x})$ is locally non-decreasing at τ (since otherwise, β_e is not maintained to be equal to $\frac{1}{\lambda} \nabla_e F(\mathbf{x})$). Therefore, $\frac{\partial \beta_e}{\partial \tau} \geq 0$ and so $\partial(\frac{1}{\beta_e})/\partial \tau \leq 0$. Hence, the derivative of the right hand side of the lemma inequality according to τ is upper bounded by

$$\sum_i \frac{\partial \alpha_{i,A^*}}{\partial \tau} \cdot \frac{b_{i,e,A^*}}{\max b_{i,e,A} \cdot d} \cdot \frac{\ln(1+2d^2)}{\beta_e} \cdot \exp\left(\frac{\ln(1+2d^2)}{\beta_e} \cdot \sum_{A:e \notin A} \sum_i b_{i,e,A} \alpha_{i,A}\right)$$

which is bounded by $\frac{\partial x_e}{\partial \tau}$ by the same argument as the previous case. The lemma follows. \square

Lemma 2.4 *The dual variables defined as above are feasible.*

Proof As long as a primal covering constraint is unsatisfied, the x -variables are always increased. Therefore, at the end of an iteration, the primal constraint is satisfied. Consider the first dual constraint. The algorithm always maintains that $\sum_i \sum_{A:e \notin A} b_{i,e,A} \alpha_{i,A} \leq \beta_e$ (strict inequality happens only if $x_e = 1$). Whenever this inequality is violated then by the algorithm, some α -variables are decreased in such a way that the increasing rate of $\sum_i \sum_{A:e \notin A} b_{i,e,A} \alpha_{i,A}$ is at most 0. Hence, by the definition of β -variables, the first dual constraint holds.

Consider the second dual constraint. Let \mathbf{x} be the current solution of the algorithm. By the algorithm, for each fixed resource e , $\beta_e = \frac{1}{\lambda} \nabla_e F(\mathbf{y}^e)$ for some \mathbf{y}^e where $y_{e'} \leq x_{e'}$ for every resource e' . (Since at some moment, the algorithm increases x_e without increasing β_e for some e .) Moreover, $\mathbf{y} := \bigvee_e \mathbf{y}^e \leq \mathbf{x}$ (meaning that $y_{e'} \leq x_{e'}$ for every e'). By definitions of dual variables, the second dual constraint (after rearranging terms) reads

$$\frac{1}{\lambda} \sum_{e \in S} \nabla_e F(\mathbf{y}^e) \leq F(\mathbf{1}_S) + \frac{\mu}{4\lambda \cdot \ln(1 + 2d^2)} F(\mathbf{x}).$$

Besides, as F is monotone, $F(\mathbf{x}) \geq F(\mathbf{y})$. To prove the above inequality, it is sufficient to prove that

$$\frac{1}{\lambda} \sum_{e \in S} \nabla_e F(\mathbf{y}^e) \leq F(\mathbf{1}_S) + \frac{\mu}{4\lambda \cdot \ln(1 + 2d^2)} F(\mathbf{y}).$$

This inequality is exactly the $(\lambda, \frac{\mu}{4\ln(1+2d^2)})$ -min-local smoothness of F . Hence, the lemma follows. \square

We are now ready to prove the main theorem.

Theorem 1.2 *Let F be the multilinear extension of the objective cost f and d be the maximal row sparsity of the constraint matrix, i.e., $d = \max_i |\{a_{ie} : a_{ie} > 0\}|$. Assume that F is $(\lambda, \frac{\mu}{\ln(1+2d^2)})$ -min-locally-smooth for some parameters $\lambda > 0$ and $\mu < 1$. Then there exists a $O(\frac{\lambda}{1-\mu} \cdot \ln d)$ -competitive algorithm for the fractional covering problem.*

Proof We will bound the increases of the cost and the dual objective at any time τ in the execution of Algorithm 1. Let A^* be the current set of resources e such that $x_e = 1$. The derivative of the objective with respect to τ is:

$$\begin{aligned} \sum_e \nabla_e F(\mathbf{x}) \cdot \frac{\partial x_e}{\partial \tau} &= \sum_{\substack{e: b_{k,e,A^*} > 0 \\ x_e < 1}} \nabla_e F(\mathbf{x}) \cdot \frac{b_{k,e,A^*} \cdot x_e + 1/d}{\lambda \cdot \beta_e} \\ &\leq \sum_{e: b_{k,e,A^*} > 0} \left(b_{k,e,A^*} \cdot x_e + \frac{1}{d} \right) \leq 2. \end{aligned} \quad (2.3)$$

The first inequality follows by $\nabla_e F(\mathbf{x}) \leq \lambda \cdot \beta_e$. The second inequality is due to the definition of d and the fact that $\sum_{e \notin A^*} b_{k,e,A^*} \cdot x_e \leq 1$ always holds during the algorithm.

For a time τ , let $U(\tau)$ be the set of resources e such that $\sum_i \sum_{A:e \notin A} b_{i,e,A} \alpha_{i,A} = \beta_e$ and $b_{k,e,A^*} > 0$. Note that $|U(\tau)| \leq d$ by definition of d . As long as $\sum_{e \notin A^*} b_{k,e,A^*} x_e < 1$, by Lemma 2.3, we have for every $e \in U(\tau)$,

$$\frac{1}{b_{k,e,A^*}} > x_e \geq \frac{1}{\max b_{i,e,A} \cdot d} \left[\exp\left(\ln(1 + 2d^2)\right) - 1 \right].$$

Therefore, $\frac{b_{k,e,A^*}}{\max_i b_{i,e,A}} \leq \frac{1}{2d}$.

We are now bounding the increase of the dual at time τ . The derivative of the dual with respect to τ is:

$$\begin{aligned} \frac{\partial D}{\partial \tau} &= \sum_i \sum_A \frac{\partial \alpha_{i,A}}{\partial \tau} + \frac{\partial \gamma}{\partial \tau} = \sum_i c_{i,A^*} \cdot \frac{\partial \alpha_{i,A^*}}{\partial \tau} + \frac{\partial \gamma}{\partial \tau} \\ &= \frac{1}{\lambda \cdot \ln(1 + 2d^2)} \left(1 - \sum_{e \in U(\tau)} \frac{b_{k,e,A^*}}{b_{m^*,e,A}} \right) - \frac{\mu}{4\lambda \cdot \ln(1 + 2d^2)} \sum_e \nabla_e F(\mathbf{x}) \cdot \frac{\partial x_e}{\partial \tau} \\ &\geq \frac{1}{\lambda \cdot \ln(1 + 2d^2)} \left(1 - \sum_{e \in U(\tau)} \frac{1}{2d} \right) - \frac{\mu}{2\lambda \cdot \ln(1 + 2d^2)} \\ &\geq \frac{1 - \mu}{2\lambda \cdot \ln(1 + 2d^2)} \end{aligned}$$

The third equality holds since α_{k,A^*} is increased and other α -variables in $U(\tau)$ are decreased. The first inequality uses the fact that $\frac{b_{k,e,A^*}}{\max_i b_{i,e,A}} \leq \frac{1}{2d}$ and Inequality (2.3). The last inequality holds since $|U(\tau)| \leq d$. Hence, the competitive ratio is $O\left(\frac{\lambda}{1-\mu} \cdot \ln d\right)$. \square

2.2.2 Applications

In this section, we consider the applications of Theorem 1.2 for classes of cost functions which have been extensively studied in optimization such as polynomials with non-negative coefficients, ℓ_k -norms and submodular functions. We are interested in deriving fractional solutions¹ with performance guarantee. We show that Algorithm 1 with multilinear extension yields competitive fractional solutions for the classes of functions mentioned above and also for some natural classes of non-convex functions.

We first take a closer look to the definition of min-local smoothness. Let F be a multilinear extension of a set function f . By definition of multilinear extension, $F(\mathbf{x}) = \mathbb{E}[f(\mathbf{1}_T)]$ where T is a random set such that a resource e appears in T with probability x_e . Moreover, since F is linear in x_i , we have

$$\begin{aligned} \frac{\partial F}{\partial x_e}(\mathbf{x}) &= F(x_1, \dots, x_{e-1}, 1, x_{e+1}, \dots, x_n) - F(x_1, \dots, x_{e-1}, 0, x_{e+1}, \dots, x_n) \\ &= \mathbb{E} \left[f(\mathbf{1}_{R \cup \{e\}}) - f(\mathbf{1}_R) \right] \end{aligned}$$

where R is a random subset of resources $N \setminus \{e\}$ such that e' is included with probability $x_{e'}$. Therefore, in order to prove that F is (λ, μ) -min-locally-smooth, it is equivalent to show that, for any set $S \subset \mathcal{E}$ and for any vectors $\mathbf{x}^e \in [0, 1]^n$ for $e \in \mathcal{E}$,

$$\sum_{e \in S} \mathbb{E} \left[f(\mathbf{1}_{R^e \cup \{e\}}) - f(\mathbf{1}_{R^e}) \right] \leq \lambda f(\mathbf{1}_S) + \mu \mathbb{E} \left[f(\mathbf{1}_R) \right] \quad (2.4)$$

¹Rounding schemes (in order to obtain integral solution) for concrete problems are problem-specific and are not considered in this section. Several rounding techniques have been shown for different problems, for example in [11] for polynomials with non-negative coefficients, or using online contention resolution schemes for submodular functions [57].

where R^e is a random subset of resources $N \setminus \{e\}$ such that e' is included with probability $x_{e'}^e$ and R is a random subset of resources $N \setminus \{e\}$ such that e' is included with probability $\max_{e \in S} x_{e'}^e$. Note that if ∇F is monotone (in all coordinates) then the following inequality implies Inequality (2.4).

$$\sum_{e \in S} \mathbb{E} \left[f(\mathbf{1}_{R \cup \{e\}}) - f(\mathbf{1}_R) \right] \leq \lambda f(\mathbf{1}_S) + \mu \mathbb{E} \left[f(\mathbf{1}_R) \right] \quad (2.5)$$

2.2.2.1 Polynomials with non-negative coefficients.

Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a polynomial with non-negative coefficients and the cost function $f : \{0, 1\}^n \rightarrow \mathbb{R}^+$ defined as $f(\mathbf{1}_S) = g(\sum_{e \in S} a_e)$ where $a_e \geq 0$ for every e . The following proposition shows that our algorithm yields the same competitive ratio as the one derived in [11] for this class of cost functions. This bound indeed is *tight* [11] (up to a constant factor). Note that, Azar et al. [11] gave randomized algorithms for several problems by rounding their fractional solutions. As one can approach a multilinear extension of any function up to a high precision [123], applying the same rounding schemes in [11] for the corresponding problems based on our fractional solutions, one can obtain randomized algorithms with similar bounds as in [11].

Proposition 2.6 ([11]) *For any convex polynomial function g of degree k , there exists an $O((k \ln d)^k)$ -competitive algorithm for the fractional covering problem.*

Proof We prove that Algorithm 1 is $O((k \ln d)^k)$ -competitive for this class of cost functions. By Theorem 1.2, it is sufficient to verify that F is $((k \ln k)^{k-1}, \frac{k-1}{k \ln d})$ -min-locally smooth. Note that ∇F is monotone. We indeed prove a stronger inequality than (2.5), that is for any set $R \subset \mathcal{E}$,

$$\sum_{e \in S} \left[f(\mathbf{1}_{R \cup \{e\}}) - f(\mathbf{1}_R) \right] \leq O((k \ln k)^{k-1}) \cdot f(\mathbf{1}_S) + \frac{k-1}{k \ln k} \cdot f(\mathbf{1}_R)$$

or equivalently, for any set $R \subset \mathcal{E}$,

$$\sum_{e \in S} \left[g\left(a_e + \sum_{e' \in R} a_{e'}\right) - g\left(\sum_{e' \in R} a_{e'}\right) \right] \leq O((k \ln k)^{k-1}) \cdot g\left(\sum_{e \in S} a_e\right) + \frac{k-1}{k \ln k} \cdot g\left(\sum_{e' \in R} a_{e'}\right)$$

This inequality holds by Lemma 2.2 (in the appendix). Hence, the proposition follows. \square

2.2.2.2 Beyond convex functions.

Consider the following natural cost functions which represent more practical costs when serving clients as mentioned in the introduction (the cost initially increases fast then becomes more stable before growing quickly again). Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a non-convex function defined as $g(y) = y^k$ if $y \leq M_1$ or $y \geq M_2$ and $g(y) = g(M_1)$ if $M_1 \leq y \leq M_2$ where $M_1 < M_2$ are some constants. The cost function $f : \{0, 1\}^n \rightarrow \mathbb{R}^+$ defined as $f(\mathbf{1}_S) = g(\sum_{e \in S} a_e)$ where $a_e \geq 0$ for every e . In fact, the corresponding multilinear extension F is $((k \ln k)^{k-1}, \frac{k-1}{k \ln d})$ -min-locally smooth. Again, it sufficient to verify Inequality (2.5) and the proof is similar to the one in Proposition 2.6 (or more specifically, Lemma 2.2 in the appendix) and note that the derivative of g for $M_1 < y < M_2$ equals 0.

Proposition 2.7 *The algorithm is $O((k \ln d)^k)$ -competitive for minimizing the non-convex objective function defined above under covering constraints.*

2.2.2.3 Submodular functions.

Consider the class of submodular functions f satisfying

$$f(\mathbf{1}_{S \cup \{e\}}) - f(\mathbf{1}_S) \geq f(\mathbf{1}_{T \cup \{e\}}) - f(\mathbf{1}_T) \quad \forall e, S \subset T,$$

and $f(\mathbf{1}_\emptyset) = 0$. Submodular optimization has been extensively studying in optimization and machine learning. In the context of online algorithms, Buchbinder et al. [38] have considered submodular optimization with preemption, where one can reject previously accepted elements, and have given constant competitive algorithms for unconstrained and knapsack-constraint problems. To the best of our knowledge, the problem of online submodular minimization under covering constraints have not been considered.

An important concept in studying submodular functions is the *curvature*. Given a submodular function f , the *total curvature* κ_f [47] of f is defined as

$$\kappa_f = 1 - \min_e \frac{f(\mathbf{1}_\mathcal{E}) - f(\mathbf{1}_{\mathcal{E} \setminus \{e\}})}{f(\mathbf{1}_{\{e\}})}.$$

Intuitively, the total curvature measures how far away f is from being *modular*. The concept of curvature has been used to determine both upper and lower bounds on the approximation ratios for many submodular and learning problems [47, 65, 12, 124, 80, 117].

In the following, we present a competitive algorithm for minimizing a monotone submodular function under covering constraints where the competitive ratio is characterized by the curvature of the function (and also the sparsity d of the covering matrix). We first look at an useful property of the total curvature.

Lemma 2.5 *For any set S , it always holds that*

$$f(\mathbf{1}_S) \geq (1 - \kappa_f) \sum_{e \in S} f(\mathbf{1}_{\{e\}}).$$

Proof Let $S = \{e_1, \dots, e_m\}$ be an arbitrary subset of \mathcal{E} . Let $S_i = \{e_1, \dots, e_i\}$ for $1 \leq i \leq m$ and $S_0 = \emptyset$. We have

$$\begin{aligned} f(\mathbf{1}_S) &\geq f(\mathbf{1}_\mathcal{E}) - f(\mathbf{1}_{\mathcal{E} \setminus S}) = \sum_{i=0}^{m-1} f(\mathbf{1}_{\mathcal{E} \setminus S_i}) - f(\mathbf{1}_{\mathcal{E} \setminus S_{i+1}}) \geq \sum_{i=1}^m f(\mathbf{1}_\mathcal{E}) - f(\mathbf{1}_{\mathcal{E} \setminus \{e_i\}}) \\ &\geq (1 - \kappa_f) \sum_{i=1}^m f(\mathbf{1}_{e_i}) \end{aligned}$$

where the first two inequalities are due to submodularity of f and the last inequality follows by the definition of the curvature. \square

Proposition 2.8 *The algorithm is $O(\frac{\log d}{1 - \kappa_f})$ -competitive for minimizing monotone submodular function under covering constraints.*

Proof It is sufficient to verify that F is $(\frac{1}{1 - \kappa_f}, 0)$ -min-locally smooth. Indeed, the $(\frac{1}{1 - \kappa_f}, 0)$ -min-local smoothness holds due to the submodularity and Lemma 2.5: for

any subsets R^e , we have

$$\sum_{e \in S} [f(\mathbf{1}_{R^e \cup \{e\}}) - f(\mathbf{1}_{R^e})] \leq \sum_{e \in S} [f(\mathbf{1}_{\{e\}})] \leq \frac{1}{1 - \kappa_f} \cdot f(\mathbf{1}_S).$$

Therefore, the proposition follows. \square

2.3 Primal-Dual Framework for Packing Problems

Consider the following integer optimization problem. Let \mathcal{E} be a set of n resources and let $f : \{0, 1\}^n \rightarrow \mathbb{R}^+$ be an *arbitrary* cost function. Let $x_e \in \{0, 1\}$ be a variable indicating whether resource e is selected. The packings constraints $\sum_e b_{i,e} x_e \leq 1$ for every i are given in advance and resources e are revealed online one-by-one. At any time, one needs to maintain a feasible integer solution \mathbf{x} . The goal is to design an algorithm that maximizes $f(\mathbf{x})$ subject to the online packing constraints and $x_e \in \{0, 1\}$ for every e .

2.3.1 Algorithm for Fractional Packing

Recall that a differentiable function $F : [0, 1]^n \rightarrow \mathbb{R}^+$ is (λ, μ) -max-locally-smooth if for any set $S \subset \mathcal{E}$, and for any vectors $\mathbf{x}^e \in [0, 1]^n$, the following inequality holds:

$$\sum_{e \in S} \nabla_e F(\mathbf{x}^e) \geq \lambda F(\mathbf{1}_S) - \mu F(\mathbf{x}).$$

where $\mathbf{x} := \bigvee_{e \in S} \mathbf{x}^e$, meaning that $x_{e'} = \max_e \{x_{e'}^e\}$ for any coordinate e' .

Formulation. We say that $S \subset \mathcal{E}$ is a *configuration* if $\mathbf{1}_S$ corresponds to a feasible solution. Let x_e be a variable indicating whether the resource e is used. For configuration S , let z_S be a variable such that $z_S = 1$ if and only if $x_e = 1$ for every resource $e \in S$, and $x_e = 0$ for $e \notin S$. In other words, $z_S = 1$ iff $\mathbf{1}_S$ is the selected solution of the problem. We consider the following formulation and the dual of its relaxation.

Primal:

$$\begin{aligned} & \max \sum_S f(\mathbf{1}_S) z_S \\ & \sum_e b_{i,e} \cdot x_e \leq 1 \quad \forall i \\ & \sum_{S: e \in S} z_S = x_e \quad \forall e \\ & \sum_S z_S = 1 \\ & x_e, z_S \in \{0, 1\} \quad \forall e, S \end{aligned}$$

Dual:

$$\begin{aligned} & \min \sum_i \alpha_i + \gamma \\ & \sum_i b_{i,e} \cdot \alpha_i \geq \beta_e \quad \forall e \\ & \gamma + \sum_{e \in S} \beta_e \geq f(\mathbf{1}_S) \quad \forall S \\ & \alpha_i \geq 0 \quad \forall i \end{aligned}$$

In the primal, the first constraints represent the given polytope. Note that the box constraint $x_e \leq 1$ is included among these constraints. The second constraint ensures that if a resource e is chosen then the selected solution must contain e . The third constraint says that one solution (configuration) must be selected.

Algorithm. Assume that function $F(\cdot)$ is (λ, μ) -max-locally smooth. Let d be the maximal number of positive entries in a row, i.e., $d = \max_i |\{b_{ie} : b_{ie} > 0\}|$. Define $\rho = \max_i \max_{e, e' : b_{ie'} > 0} b_{ie} / b_{ie'}$. Denote $\nabla_e F(\mathbf{x}) = \partial F(\mathbf{x}) / \partial x_e$. In the algorithm, at the arrival of a new resource e , while $\nabla_e F(\mathbf{x}) > 0$ (i.e., one can still improve the cost by increasing x_e) and $\sum_i b_{i,e} \alpha_i \leq \frac{1}{\lambda} \nabla_e F(\mathbf{x})$, the primal variable x_e and dual variables α_i 's are increased by appropriate rates. We will argue in the analysis that the primal/dual solutions returned by the algorithm are feasible. Recall that by definition of the multilinear extension, $\nabla_e F(\mathbf{x}) = \mathbb{E}[f(\mathbf{1}_{R \cup \{e\}}) - f(\mathbf{1}_R)]$ where R is a random subset of resources $N \setminus \{e\}$ such that e' is included with probability $x_{e'}$. Therefore, during the iteration of the while loop with respect to resource e , only x_e is modified and $x_{e'}$ remains fixed for $e' \neq e$, so $\nabla_e F(\mathbf{x})$ is constant during the iteration.

Algorithm 2 Algorithm for Packing Constraints.

- 1: All primal and dual variables are initially set to 0.
- 2: At every step, always maintain $z_S = \prod_{e \in S} x_e \prod_{e \notin S} (1 - x_e)$.
- 3: Upon the arrival of new resource e .
- 4: **while** $\sum_i b_{i,e} \alpha_i \leq \frac{1}{\lambda} \nabla_e F(\mathbf{x})$ and $\nabla_e F(\mathbf{x}) > 0$ **do**
- 5: Increase τ at rate 1 and increase x_e at rate $\frac{1}{\nabla_e F(\mathbf{x}) \cdot \ln(1+d\rho)}$.
- 6: **for** i such that $b_{i,e} > 0$ **do**
- 7: Increase α_i according to the following function

$$\frac{\partial \alpha_i}{\partial \tau} \leftarrow \frac{b_{i,e} \cdot \alpha_i}{\nabla_e F(\mathbf{x})} + \frac{1}{d\lambda}$$

- 8: **end for**
 - 9: **end while**
-

Dual variables. Variables α_i 's are constructed in the algorithm. Let \mathbf{x} be the current solution of the algorithm and let \mathbf{x}^e be the solution after the while loop with respect to resource e . Define $\gamma = \frac{\mu}{\lambda} F(\mathbf{x})$ where \mathbf{x} and $\beta_e = \frac{1}{\lambda} \cdot \nabla_e F(\mathbf{x}^e)$. Note that by the observation above, during the while loop with respect to resource e , $\beta_e = \frac{1}{\lambda} \cdot \nabla_e F(\mathbf{x})$.

The following lemma gives a lower bound on α -variables.

Lemma 2.6 *At any moment during the execution of the algorithm, it always holds that for every i*

$$\alpha_i \geq \frac{\nabla_e F(\mathbf{x})}{\max_{e'} b_{i,e'} \cdot d\lambda} \left[\exp\left(\ln(1+d\rho) \cdot \sum_{e'} b_{i,e'} \cdot x_{e'}\right) - 1 \right].$$

Proof We prove the lemma by induction. At the beginning of the instance, while no resource has been released, both sides of the lemma are 0. Assume that the lemma holds until the arrival of a resource e . Consider a moment τ during the loop corresponding to resource e . Note that as F is a linear extension, $\partial \nabla_e F(\mathbf{x}) / \partial \tau = 0$. The

derivative of the right hand side of the lemma inequality according to τ is

$$\begin{aligned} & \frac{\nabla_e F(\mathbf{x})}{\max_{e'} b_{i,e'} \cdot d\lambda} \cdot \ln(1 + d\rho) \cdot b_{i,e} \cdot \frac{\partial x_e}{\partial \tau} \cdot \exp\left(\ln(1 + d\rho) \cdot \sum_{e'} b_{i,e'} \cdot x_{e'}\right) \\ & \leq \frac{\nabla_e F(\mathbf{x})}{\max_{e'} b_{i,e'} \cdot d\lambda} \cdot \ln(1 + d\rho) \cdot b_{i,e} \cdot \frac{1}{\nabla_e F(\mathbf{x}) \cdot \ln(1 + d\rho)} \cdot \left(\frac{\max_{e'} b_{i,e'} \cdot d\lambda \cdot \alpha_i}{\nabla_e F(\mathbf{x})} + 1\right) \\ & \leq \frac{b_{i,e} \cdot \alpha_i}{\nabla_e F(\mathbf{x})} + \frac{1}{d\lambda} = \frac{\partial \alpha_i}{\partial \tau} \end{aligned}$$

where in the first inequality, we use the induction hypothesis. So the rate in the left-hand side is always larger than that in the right-hand side. Hence, the lemma follows. \square

Lemma 2.7 *The dual variables defined as above are feasible.*

Proof We first prove the primal feasibility. During the execution of the algorithm, if $\sum_i b_{i,e'} x_{e'} > 1$ for some constraint i then by Lemma 2.3,

$$\alpha_i > \frac{\nabla_e F(\mathbf{x})}{\max_{e'} b_{i,e'} \cdot d\lambda} \left[\exp\left(\ln(1 + d\rho)\right) - 1 \right] = \frac{\rho \cdot \nabla_e F(\mathbf{x})}{\lambda \max_{e'} b_{i,e'}} \geq \frac{\nabla_e F(\mathbf{x})}{\lambda b_{i,e}}$$

Therefore, $\sum_i b_{i,e} \alpha_i > \frac{1}{\lambda} \nabla_e F(\mathbf{x})$ and hence the algorithm would have stopped increasing x_e at some earlier point. Consequently, the constraint $\sum_i b_{i,e'} x_{e'} \leq 1$ is always maintained.

The first dual constraint is satisfied by the algorithm. The second dual constraint reads

$$\frac{1}{\lambda} \sum_{e \in S} \nabla_e F(\mathbf{x}^e) + \frac{\mu}{\lambda} F(\mathbf{x}) \geq F(\mathbf{1}_S)$$

which is, by arranging terms, exactly the (λ, μ) -max-local smoothness of F . Hence, the lemma follows. \square

We are now ready to prove the main theorem.

Theorem 1.3 *Let F be the multilinear extension of the objective cost f . Denote the row sparsity $d := \max_i |\{b_{ie} : b_{ie} > 0\}|$ and $\rho := \max_i \max_{e,e': b_{ie'} > 0} b_{ie}/b_{ie'}$. Assume that F is (λ, μ) -max-locally-smooth for some parameters $\lambda > 0$ and $\mu < 1$. Then there exists a $O\left(\frac{2\ln(1+d\rho)+\mu}{\lambda}\right)$ -competitive algorithm for the fractional packing problem.*

Proof We will bound the increases of the cost and the dual objective at any time τ in the execution of Algorithm 2. The derivative of the primal with respect to τ is:

$$\nabla_e F(\mathbf{x}) \cdot \frac{\partial x_e}{\partial \tau} = \nabla_e F(\mathbf{x}) \cdot \frac{1}{\nabla_e F(\mathbf{x}) \cdot \ln(1 + d\rho)} = \frac{1}{\ln(1 + d\rho)}$$

We are now bounding the increase of the dual at time τ . The derivative of the dual with respect to τ is:

$$\begin{aligned} \frac{\partial D}{\partial \tau} &= \sum_i \frac{\partial \alpha_i}{\partial \tau} + \frac{\partial \gamma}{\partial \tau} = \sum_{i:b_{i,e}>0} \left(\frac{b_{i,e} \cdot \alpha_i}{\nabla_e F(\mathbf{x})} + \frac{1}{d\lambda} \right) + \frac{\mu}{\lambda} \frac{\partial F(\mathbf{x})}{\partial \tau} \\ &= \sum_{i:b_{i,e}>0} \frac{b_{i,e} \cdot \alpha_i}{\nabla_e F(\mathbf{x})} + \sum_{i:b_{i,e}>0} \frac{1}{d\lambda} + \frac{\mu}{\lambda} \cdot \frac{1}{\ln(1+d\rho)} \\ &\leq \frac{2}{\lambda} + \frac{\mu}{\lambda \cdot \ln(1+d\rho)} = \frac{2 \ln(1+d\rho) + \mu}{\lambda \cdot \ln(1+d\rho)} \end{aligned}$$

where the inequality holds since during the algorithm $\sum_i b_{i,e} \cdot \alpha_i \leq \beta_e = \frac{1}{\lambda} \nabla_e F(\mathbf{x})$. Hence, the competitive ratio is $O\left(\frac{2 \ln(1+d\rho) + \mu}{\lambda}\right)$. \square

Note that the competitive ratio is the same up to a constant factor as the performance guarantee for maximizing a linear function under packing constraints. Specifically, if function f is linear then the smooth parameters are $\lambda = \mu = 1$.

2.3.2 Applications to online submodular maximization

Consider a online submodular maximization subject to packing constraints. We incorporate additional constraints $x_e \leq 2/3$ (instead of box constraint $x_e \leq 1$) for every e . The advantage of these stronger constraints, as shown below, is that we can bound the smooth parameters while losing only a constant factor in the competitive ratio. We are now determining smooth parameters of the multilinear extension F .

Lemma 2.8 *Let f be an arbitrary submodular function. Then, the multilinear extension F is $(1,1)$ -smooth if f is monotone and is $(1/3, 1)$ -smooth if f is non-monotone.*

Proof For arbitrary submodular function f , it holds that [56, Lemma III.5] for any vector \mathbf{y} and any subset S , $F(\mathbf{1}_S \vee \mathbf{y}) \geq (1 - \max_e y_e)F(\mathbf{1}_S)$. Moreover, if f is monotone, $F(\mathbf{1}_S \vee \mathbf{y}) \geq F(\mathbf{1}_S)$.

Consider arbitrary vectors \mathbf{x}^e and let $\mathbf{x} = \bigvee_e \mathbf{x}^e$. As F is the linear extension of a submodular function, $\nabla_e F(\mathbf{x}^e) \geq \nabla_e F(\mathbf{x}) = \mathbb{E}[f(\mathbf{1}_{R \cup \{e\}}) - f(\mathbf{1}_R)]$ where R is a random subset of resources $N \setminus \{e\}$ such that e' is included with probability $x_{e'}$. Therefore, for any subset S ,

$$\begin{aligned} F(\mathbf{x}) + \sum_{e \in S} \nabla_e F(\mathbf{x}^e) &\geq F(\mathbf{x}) + \sum_{e \in S} \mathbb{E}[f(\mathbf{1}_{R \cup \{e\}}) - f(\mathbf{1}_R)] \\ &= \mathbb{E} \left[f(\mathbf{1}_R) + \sum_{e \in S} [f(\mathbf{1}_{R \cup \{e\}}) - f(\mathbf{1}_R)] \right] \geq \mathbb{E}[f(\mathbf{1}_{R \cup S})] = F(\mathbf{1}_S \vee \mathbf{x}) \\ &\geq \begin{cases} F(\mathbf{1}_S) & \text{if } f \text{ monotone,} \\ (1 - \max_e x_e)F(\mathbf{1}_S) & \text{otherwise} \end{cases} \geq \begin{cases} F(\mathbf{1}_S) & \text{if } f \text{ monotone,} \\ 1/3 \cdot F(\mathbf{1}_S) & \text{otherwise} \end{cases} \end{aligned}$$

where the second inequality is due to the submodularity of f , and the last inequality holds since $x_e \leq 2/3$ for every e . The lemma follows. \square

The previous lemma and Theorem 1.3 lead to the following result.

Proposition 2.9 *Algorithm 2 yields a $O(\ln(1+d\rho))$ -competitive fractional solution for maximizing (arbitrary) submodular functions under packing constraints.*

One can derive online randomized algorithms for specific problems by rounding the fractional solutions. For example, using the online contention resolution rounding schemes [57], one can obtain randomized algorithms for several specific constraint polytopes, for example, knapsack polytopes, matching polytopes and matroid polytopes.

Chapter 3

Game Efficiency through Linear Programming Duality

In this chapter, we present the primal-dual approach to algorithmic game theory. We show the natural development and the applicability of our framework from full-information games (Sections 3.1 and 3.2) to incomplete-information games (Sections 3.3, 3.4 and 3.5) by revisiting known results and establishing new ones. The approach provides a general recipe to analyze the efficiency of games and also to derive concepts leading to improvements. In this chapter, we consider games in discrete settings of games (for example, discrete values of valuations and payments, i.e., there are only a finite (large) number of possible valuations and payments). The main purpose of restricting to discrete settings is that we can use tools from linear programming. The continuous settings can be done by considering successively finer discrete spaces.

3.1 Smooth Games under the Lens of Duality

We consider smooth games [108] in the point of view of configuration LPs and duality. Recall that in a game, each player i selects a strategy s_i from a set \mathcal{S}_i for $1 \leq i \leq n$ and that forms a *strategy profile* $\mathbf{s} = (s_1, \dots, s_n)$. The cost $C_i(\mathbf{s})$ of player i is a function of the strategy profile \mathbf{s} — the chosen strategies of all players. A game with a joint cost objective function $C(\mathbf{s}) = \sum_{i=1}^n C_i(\mathbf{s})$ is (λ, μ) -smooth if for every two outcomes \mathbf{s} and \mathbf{s}^* ,

$$\sum_{i=1}^n C_i(\mathbf{s}_i^*, \mathbf{s}_{-i}) \leq \lambda \cdot C(\mathbf{s}^*) + \mu \cdot C(\mathbf{s})$$

The *robust price of anarchy* of a game G is

$$\rho(G) := \inf \left\{ \frac{\lambda}{1 - \mu} : \text{the game is } (\lambda, \mu)\text{-smooth where } \mu < 1 \right\}$$

In his seminal paper, Roughgarden [108] has introduced the notion of smooth games and characterize the efficiency of equilibria as follows.

Theorem 3.1 ([108]) *For every game G with robust PoA $\rho(G)$, every coarse correlated equilibrium σ of G and every strategy profile \mathbf{s}^* , it holds that*

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

Until the end of the section, we revisit this theorem by our primal-dual approach.

Formulation. Given a game, we formulate the corresponding optimization problem by a configuration LP. Let x_{ij} be a variable indicating whether player i chooses strategy $s_{ij} \in \mathcal{S}_i$. Informally, a *configuration* A in the formulation is a strategy profile of the game. Formally, a configuration A consists of pairs (i, j) such that $(i, j) \in A$ means that in configuration A , $x_{ij} = 1$. (In other words, in this configuration, player i selects strategy $s_{ij} \in \mathcal{S}_i$.) For every configuration A , let z_A be a variable such that $z_A = 1$ if and only if $x_{ij} = 1$ for all $(i, j) \in A$. Intuitively, $z_A = 1$ if configuration A is the outcome of the game. For each configuration A , let $c(A)$ be the cost of the outcome (strategy profile) corresponding to configuration A . Consider the following formulation and the dual of its relaxation.

Primal:

$$\begin{aligned}
& \min \sum_A c(A) z_A \\
& \sum_{j: s_{ij} \in \mathcal{S}_i} x_{ij} \geq 1 \quad \forall i \\
& \sum_A z_A = 1 \\
& \sum_{A: (i,j) \in A} z_A = x_{ij} \quad \forall i, j \\
& x_{ij}, z_A \in \{0, 1\} \quad \forall i, j, A
\end{aligned}$$

Dual:

$$\begin{aligned}
& \max \sum_i \alpha_i + \beta \\
& \alpha_i \leq \gamma_{ij} \quad \forall i, j \\
& \beta + \sum_{(i,j) \in A} \gamma_{ij} \leq c(A) \quad \forall A \\
& \alpha_i \geq 0 \quad \forall i
\end{aligned}$$

In the formulation, the first constraint ensures that a player i chooses a strategy $s_{ij} \in \mathcal{S}_i$. The second constraint means that there must be an outcome of the game. The third constraint guarantees that if a player i selects some strategy s_{ij} then the outcome configuration A must contain (i, j) .

Construction of dual variables. Assuming that the game is (λ, μ) -smooth. Fix the parameters λ and μ . Given a (arbitrary) coarse correlated equilibrium σ , define dual variables as follows:

$$\alpha_i := \frac{1}{\lambda} \mathbb{E}_{\mathbf{s} \sim \sigma} [C_i(\mathbf{s})], \quad \beta := -\frac{\mu}{\lambda} \mathbb{E}_{\mathbf{s} \sim \sigma} [C(\mathbf{s})], \quad \gamma_{ij} := \frac{1}{\lambda} \mathbb{E}_{\mathbf{s} \sim \sigma} [C_i(s_{ij}, \mathbf{s}_{-i})].$$

Informally, up to some constant factors depending on λ and μ , α_i is the cost of player i in equilibrium σ , $-\beta$ stands for the cost of the game in equilibrium σ and γ_{ij} represents the cost of player i if player i uses strategy s_{ij} while other players $i' \neq i$ follows strategies in σ . Notice that β has negative value.

Feasibility. We show that the constructed dual variables form a feasible solution. The first constraint follows exactly the definition of (coarse correlated) equilibrium. The second constraint is exactly the smoothness definition. Specifically, let \mathbf{s}^* be the strategy profile corresponding to configuration A . Note that $\mathbb{E}_{\mathbf{s} \sim \sigma} [C_i(\mathbf{s}^*)] = C_i(\mathbf{s}^*)$. The dual constraint reads

$$-\frac{\mu}{\lambda} \mathbb{E}_{\mathbf{s} \sim \sigma} [C(\mathbf{s})] + \sum_i \frac{1}{\lambda} \mathbb{E}_{\mathbf{s} \sim \sigma} [C_i(s_i^*, \mathbf{s}_{-i})] \leq \mathbb{E}_{\mathbf{s} \sim \sigma} [C_i(\mathbf{s}^*)]$$

which is the definition of (λ, μ) -smoothness by rearranging the terms and using the linearity of expectation.

Price of Anarchy. By weak duality, the optimal cost among all outcomes of the problem (strategy profiles of the game) is at least the dual objective of the constructed dual variables. Hence, in order to bound the PoA, we will bound the ratio between the cost of an (arbitrary) equilibrium σ and the dual objective of the corresponding dual variables. The cost of equilibrium σ is $\mathbb{E}_{s \sim \sigma}[C(s)]$ while the dual objective of the constructed dual variables is

$$\sum_{i=1}^n \frac{1}{\lambda} \mathbb{E}_{s \sim \sigma}[C_i(s)] - \frac{\mu}{\lambda} \mathbb{E}_{s \sim \sigma}[C(s)] = \frac{1-\mu}{\lambda} \mathbb{E}_{s \sim \sigma}[C(s)].$$

Therefore, for a (λ, μ) -smooth game, the PoA is at most $\lambda/(1-\mu)$.

Remark. As shown in [108], Theorem 3.1 applies also to outcome sequences generated by repeated play such as vanishing average regret. By the same duality approach, we can also recover this result (by setting dual variables related to the average cost during the play).

3.2 Congestion Games

3.2.1 Atomic Congestion Games

Model. Atomic congestion games were defined by Rosenthal [106]. In this section, we consider atomic weighted congestion games, a generalized version of the standard congestion game. In the game, we are given a ground set E of resources, a set of n players with strategy sets $\mathcal{S}_1, \dots, \mathcal{S}_n \subseteq 2^E$ and weights w_1, \dots, w_n and a cost function $\ell_e : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ for each resource $e \in E$. Note that the weighted setting generalizes the standard congestion games in which $w_i = 1$ for all players i . Given a strategy profile $s = (s_1, \dots, s_n)$ where $s_i \in \mathcal{S}_i$ for each player i , we say that $w_e(s) = \sum_{i:e \in s_i} w_i$ is the *load* induced on e by s . The cost of a player i is defined as $C_i(s) = \sum_{e:e \in s_i} w_i \cdot \ell_e(w_e)$ where w_e is the load on resource e induced by profile s . The total cost of the game in profile s is $C(s) = \sum_{i=1}^n C_i(s) = \sum_{e:e \in s_i} w_e(s) \cdot \ell_e(w_e(s))$.

The PoA of atomic congestion games has been an extensively studied topic in algorithmic game theory. Most notably, Roughgarden [108] proved that the smoothness argument gave tight bounds for (unweighted) atomic congestion games. For the weighted setting, Bhawalkar et al. [24] showed that the smoothness framework also gave tight bounds for large classes of congestion games.

In this section, we reprove the upper bound [108, 24] on the PoA in atomic congestion games. The result is proved by the same duality approach described in Section 3.1. Nevertheless, we present a proof for this result for the following reasons. First, we give a slightly different formulation of the configuration LP. To establish smoothness, all current proofs are based on smooth-inequalities related to resources. The new formulation is given to capture the smooth-inequality notion on resources. Second, the new proof will be used later to show that in terms of PoA, the atomic congestion games have a strong connection with non-atomic and splittable congestion games under the viewpoint of duality.

We say that a cost function $\ell_e : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ for a resource e is (λ, μ) -*resource-smooth* if for all sequences of non-negative real numbers $(a_i)_{i=1}^n$ and $(b_i)_{i=1}^n$, it holds that

$$\sum_{i=1}^n \ell_e \left(\sum_{j=1}^i a_j + b_i \right) \leq \lambda \cdot \ell_e \left(\sum_{i=1}^n b_i \right) + \mu \cdot \ell_e \left(\sum_{i=1}^n a_i \right)$$

Theorem 3.2 ([108, 24]) Let \mathcal{L} be a non-empty set of cost functions. The PoA of every coarse correlated equilibrium of every (weighted) atomic congestion game with cost functions $\ell_e \in \mathcal{L}$ is at most

$$\inf \left\{ \frac{\lambda}{1 - \mu} : \ell_e \text{ is } (\lambda, \mu)\text{-resource-smooth where } \mu < 1 \forall e \in E \right\}$$

Proof

Formulation. Let x_{ij} be variable indicating whether player i chooses strategy $s_{ij} \in \mathcal{S}_i$. For every resource e and every subset of players T , let z_{eT} be a variable such that $z_{eT} = 1$ if and only if every player $i \in T$ uses resource e , i.e., $e \in s_i$, and player $i \notin T$ does not use resource e . Denote $w(T) = \sum_{i \in T} w_i$. Consider the following integer program and its dual. In the primal, the first constraint says that a player i has to select a strategy $s_{ij} \in \mathcal{S}_i$. The second constraint means that a subset of players T will use resource e . The third constraint guarantees that if a player i chooses some strategy $s_{ij} \in \mathcal{S}_i$ containing resource e then there must be a subset of players T such that $i \in T$ and $z_{eT} = 1$.

Primal:

$$\begin{aligned} \min \quad & \sum_e w(T) \ell_e(w(T)) z_{eT} \\ & \sum_j x_{ij} \geq 1 \quad \forall i \\ & \sum_T z_{eT} = 1 \quad \forall e \\ & \sum_{T:i \in T} z_{eT} = \sum_{j:e \in s_{ij}} x_{ij} \quad \forall i, e \\ & x_{ij}, z_{eT} \in \{0, 1\} \quad \forall i, j, e, T \end{aligned}$$

Dual:

$$\begin{aligned} \max \quad & \sum_i \alpha_i + \sum_e \beta_e \\ & \alpha_i \leq \sum_{e:e \in s_{ij}} \gamma_{i,e} \quad \forall i, j \\ & \beta_e + \sum_{i \in T} \gamma_{i,e} \leq w(T) \ell_e(w(T)) \quad \forall e, T \\ & \alpha_i \geq 0 \quad \forall i \end{aligned}$$

Dual Variables. Fix parameters λ and μ . Given a coarse correlated equilibrium σ , define corresponding dual variables as follows.

$$\begin{aligned} \alpha_i &:= \frac{1}{\lambda} \mathbb{E}_{s \sim \sigma} [C_i(s)], \\ \beta_e &:= -\frac{\mu}{\lambda} \mathbb{E}_{s \sim \sigma} \left[\sum_{i:e \in s_i} w_i \ell_e(w_e(s)) \right], \\ \gamma_{i,e} &:= \frac{1}{\lambda} \mathbb{E}_{s \sim \sigma} [w_i \cdot \ell_e(w_e(s_{-i}) + w_i)] \end{aligned}$$

where $w_e(s_{-i}) = \sum_{i' \neq i, e \in s_{i'}} w_{i'}$. Informally, up to some constant factors, α_i is the cost of player i in equilibrium σ , $-\beta_e$ stands for the total cost of players on resource e in this equilibrium and $\gamma_{i,e}$ represents the cost of player i on resource e if player i uses strategy containing e while other players i' follows strategy $s_{i'}$ for all $i' \neq i$.

Feasibility. By this definition of dual variables, the first dual constraint follows from the definition of coarse correlated equilibrium. The second dual constraint is satisfied due to the smoothness definition. Specifically, the constraint for a resource

e and a subset of players T reads

$$-\mathbb{E}_{\mathbf{s} \sim \sigma} \left[\frac{\mu}{\lambda} w_e(\mathbf{s}) \ell_e(w_e(\mathbf{s})) \right] + \mathbb{E}_{\mathbf{s} \sim \sigma} \left[\frac{1}{\lambda} \sum_{i \in T} w_i \ell_e(w_e(\mathbf{s}_{-i}) + w_i) \right] \leq w(T) \ell_e(w(T)).$$

The inequality holds since without expectation and by linearity of expectation (and also $\mathbb{E}_{\mathbf{s} \sim \sigma} [w(T) \cdot \ell_e(w(T))] = w(T) \ell_e(w(T))$), it is exactly the smoothness definition.

Bounding primal and dual. The PoA is bounded by the ratio between the primal objective and the dual one. Note that

$$\sum_i \alpha_i = \sum_i \frac{1}{\lambda} \mathbb{E}_{\mathbf{s} \sim \sigma} [C_i(\mathbf{s})] = \mathbb{E}_{\mathbf{s} \sim \sigma} \left[\frac{1}{\lambda} \sum_e w_e(\mathbf{s}) \ell_e(w_e(\mathbf{s})) \right].$$

Therefore,

$$\sum_i \alpha_i + \sum_e \beta_e = \frac{1 - \mu}{\lambda} \sum_e w_e(\mathbf{s}) \ell_e(w_e(\mathbf{s}))$$

Hence, $\text{PoA} \leq \lambda / (1 - \mu)$. □

3.2.2 Nonatomic Congestion Games

Model. Non-atomic congestion games were defined by Roughgarden and Tardos [112], motivated by the non-atomic routing games of Wardrop [125] and Beckmann et al. [20] and the congestion games of Rosenthal [106]. We consider a discrete version of non-atomic congestion games. The main purpose of restricting to discrete settings is that we can use tools from linear programming. The continuous settings can be done by considering successively finer discrete spaces.

Fix a constant ϵ (arbitrarily small). A non-atomic congestion game consists of a ground set E of resources and n different types of players. The set of strategies of players of type i is \mathcal{S}_i and each strategy consists of a subset of resources. Players of type i are associated to an integer number m_i that corresponds to a total amount $w_i := m_i \cdot \epsilon$. Players of type i select strategies $s_{ij} \in \mathcal{S}_i$ and distribute amounts $f_{s_{ij}}$ — a non-negative multiple of ϵ — to strategy s_{ij} , which lead to a strategy distribution $\mathbf{f} = (f_{s_{ij}})$ with $\sum_{s_{ij} \in \mathcal{S}_i} f_{s_{ij}} = w_i = m_i \epsilon$ for player type i . We abuse notation and let f_e be the total amount of congestion induced on resource e by the strategy distribution \mathbf{f} . That is, $f_e := \sum_{i=1}^n \sum_{e \in s_{ij}} f_{s_{ij}}$. Each resource has a non-decreasing cost function $\ell_e : \mathbb{R}^+ \rightarrow \mathbb{R}^+$. With respect to a strategy distribution \mathbf{f} , players of type i selecting strategy $s_{ij} \in \mathcal{S}_i$ incurs a cost $C_{s_{ij}}(\mathbf{f}) = \sum_{e \in s_{ij}} \ell_e(f_e)$. A strategy distribution \mathbf{f} is an *pure equilibrium* if for each player type i and strategy $s_{ij}, s_{ij'} \in \mathcal{S}_i$ with $f_{s_{ij}} > 0$,

$$C_{s_{ij}}(\mathbf{f}) \leq C_{s_{ij'}}(\mathbf{f}).$$

The more general equilibrium concept such as mixed, correlated and coarse correlated equilibria, are defined similarly as in Section 1.5.1. The social cost of a strategy distribution \mathbf{f} is

$$C(\mathbf{f}) = \sum_{i=1}^n \sum_{s_{ij} \in \mathcal{S}_i} f_{s_{ij}} \cdot C_{s_{ij}}(\mathbf{f}) = \sum_e f_e \cdot \ell_e(f_e)$$

For non-atomic congestion games, tight bounds on the PoA for almost all classes of cost functions have been given in [112]. The core of all analyses for PoA bounds is indeed the characterization of the unique equilibrium via a variational inequality due to Beckmann et al. [20]. This argument is explained in [48, 46]. Moreover, the connection between smoothness arguments and PoA bounds for non-atomic congestion games was revealed in [48].

3.2.2.1 Efficiency of Non-Atomic Congestion Games

In this section, we reprove the tight bound for non-atomic congestion games by the duality approach. It has been shown that in non-atomic congestion games all equilibria are essentially unique; specifically, all coarse correlated equilibria of a non-atomic congestion game have the same cost [27]. Hence, the robust PoA is indeed the PoA of pure Nash equilibrium. However, as we do not use the equilibrium characterization from [20], we will prove the PoA bound for coarse correlated equilibria. Consequently, the tight PoA bound can be proved for non-regret sequences and short best-reponse sequences. Moreover, we avoid the standard assumptions on the cost functions: $x\ell_e(x)$ is convex and $\ell_e(x)$ is differentiable.

Let \mathcal{L} be a non-empty set of cost functions. The *Pigou bound* $\zeta(\mathcal{L})$ for \mathcal{L} is defined as

$$\zeta(\mathcal{L}) := \sup_{\ell \in \mathcal{L}} \sup_{u, v} \frac{u \cdot \ell(u)}{v \cdot \ell(v) + (u - v) \cdot \ell(u)}.$$

Theorem 3.3 ([112]) *Let \mathcal{L} be a set of cost functions. Then, for every splittable congestion game G with cost functions in \mathcal{L} , the price of anarchy of G is at most $\zeta(\mathcal{L})$.*

Proof

Formulation. Denote a finite set of integer multiples of ϵ as $\{a_0, a_1, \dots, a_m\}$ where $a_k = k \cdot \epsilon$ and $m = \max_{i=1}^n m_i$. We say that T_e is a *configuration* of a resource e if $T_e = \{(i, k) : 1 \leq i \leq n, 0 \leq k \leq m\}$ in which a couple (i, k) specifies the player type i and the amount a_k that the player type i distributes to some strategy $s_{ij} \in \mathcal{S}_i$ where $e \in s_{ij}$. Note that in a configuration T_e of a resource e , there might be multiple couples $(i, k) \in T_e$ and $(i, k') \in T_e$ corresponding to players of the same type. It simply means that players of type i distribute the amounts a_k and $a_{k'}$ to some strategies s_{ij} and $s_{ij'}$ respectively that contain resource e , i.e., $e \in s_{ij}$ and $e \in s_{ij'}$. Intuitively, a configuration of a resource is a strategy distribution of a game restricted on the resource.

Let x_{ijk} be a variable indicating whether player type i distributes an amount a_k to strategy $s_{ij} \in \mathcal{S}_i$. For every resource e and a configuration T_e , let z_{e, T_e} be a variable such that $z_{e, T_e} = 1$ if and only if players of type i distribute a_k to some strategy containing resource e for $(i, k) \in T_e$. In other words, $z_{e, T_e} = 1$ if and only if for $(i, k) \in T_e$, $x_{ijk} = 1$ for some $s_{ij} \in \mathcal{S}_i$ such that $e \in s_{ij}$. For a configuration T_e of a resource e , let $w(T_e)$ be the total amount distributed by players on resource e in this configuration. Consider the following configuration integer program and the dual of its relaxation.

Primal:

$$\begin{aligned}
\min \sum_{e, T_e} w(T_e) \ell_e(w(T_e)) z_{e, T_e} \\
\sum_{j, k} a_k x_{ijk} = w_i & \quad \forall i \\
\sum_{T_e} z_{e, T_e} = 1 & \quad \forall e \\
\sum_{T_e: (i, k) \in T} z_{e, T_e} = \sum_{j: e \in s_{ij}} x_{ijk} & \quad \forall (i, k), e \\
x_{ijk}, z_{e, T_e} \in \{0, 1\} & \quad \forall i, j, e, T_e
\end{aligned}$$

Dual:

$$\begin{aligned}
\max \sum_i w_i \alpha_i + \sum_e \beta_e \\
a_k \alpha_i \leq \sum_{e: e \in s_{ij}} \gamma_{i, k, e} & \quad \forall i, k, j \\
\beta_e + \sum_{(i, k) \in T_e} \gamma_{i, k, e} \leq w(T_e) \ell_e(w(T_e)) \\
& \quad \forall e, T_e
\end{aligned}$$

In the primal, the first constraint ensures that players of type i distribute the total amount w_i among its strategies. The second constraint means that a resource e is always associated to a configuration (possibly empty). The third constraint guarantees that if player type i distributes an amount a_k to some strategy s_{ij} containing resource e then there must be a configuration T_e such that $(i, k) \in T_e$ and $z_{e, T_e} = 1$.

Dual Variables. Given a coarse correlated equilibrium σ , define the corresponding dual variables as follows.

$$\begin{aligned}
\alpha_i &:= \mathbb{E}_{f \sim \sigma} \left[\sum_{e \in s_{ij}} \ell_e(f_e) \right] \text{ for some } s_{ij} \in \mathcal{S}_i : f_{s_{ij}} > 0, \\
\gamma_{i, k, e} &:= \mathbb{E}_{f \sim \sigma} [a_k \cdot \ell_e(f_e)], \\
\beta_e &:= \inf_{T_e} \left\{ w(T_e) \ell_e(w(T_e)) - \mathbb{E}_{f \sim \sigma} \left[\sum_{(i, k) \in T_e} a_k \cdot \ell_e(f_e) \right] \right\}
\end{aligned}$$

The dual variables have similar interpretations as in previous analysis. Variable α_i is the total cost of resources in a strategy used by player type i in equilibrium σ and $\gamma_{i, k, e}$ represents an estimation of the cost of player i on resource e if player type i distributes an amount a_k in some strategy containing e while other players i' follows their strategies in σ .

Feasibility. By this definition of dual variables, the first dual constraint holds since it is the definition of coarse correlated equilibrium. The second dual constraint for a resource e and a configuration T_e reads

$$\beta_e + \sum_{(i, k) \in T_e} \mathbb{E}_{f \sim \sigma} [a_k \cdot \ell_e(f_e)] \leq w(T_e) \ell_e(w(T_e)).$$

This inequality follows directly from the definition of β -variables and linearity of expectation.

Bounding primal and dual. For each resource e , let v_e be the amount in T_e corresponding the infimum in the definition of β_e . (As we consider discrete and finite settings, the infimum is indeed the minimum.) The dual objective is

$$\sum_i w_i \alpha_i + \sum_e \beta_e = \mathbb{E}_{f \sim \sigma} \left[\sum_e \left(f_e \ell_e(f_e) + v_e \ell_e(v_e) - v_e \ell_e(f_e) \right) \right]$$

where in the equalities, we use the definition of dual variables. Note that the term $(f_e \ell_e(f_e) + v_e \ell_e(v_e) - v_e \ell_e(f_e)) \geq 0$ for every resource e . Specifically, since ℓ_e is non-decreasing, if $f_e \geq v_e$ then $f_e \ell_e(f_e) \geq v_e \ell_e(f_e)$; else $v_e \ell_e(v_e) \geq v_e \ell_e(f_e)$.

Besides, the primal objective is $\mathbb{E}_{f \sim \sigma} [\sum_e f_e \ell_e(f_e)]$. Hence, the ratio between primal and dual is at most

$$\max_e \frac{f_e \ell_e(f_e)}{v_e \ell_e(v_e) + (f_e - v_e) \ell_e(f_e)}$$

which is bounded by the Pigou bound $\zeta(\mathcal{L})$ where \mathcal{L} is the class of cost functions on resources in the game. \square

Remark. The proofs of Theorem 3.2 and Theorem 3.3 are essentially the same. By the duality approach as a unifying tool, the main difference in term of equilibrium efficiency between atomic and non-atomic congestion games is due to the definition of player cost. In the context of large games [55], while the weight of a player is negligible then the player cost in a atomic congestion game coincides with the one in the corresponding non-atomic congestion game. In this context, the PoA in atomic congestion game tends to that in non-atomic setting.

3.2.2.2 Resource Augmentation in Non-Atomic Congestion Games

Roughgarden and Tardos [111] proved that in every non-atomic selfish routing game, the cost of an equilibrium is upper bounded by that of an optimal solution routing twice as much traffic. In this section, we recover this result by the mean of linear programming duality. Resource augmentation have been widely studied in many contexts in algorithms. Recently, Lucarelli et al. [92] have presented an unified approach to study resource augmentation in online (scheduling) problems based on primal-dual techniques. We will follow this framework to prove the resource augmentation result in non-atomic congestion games.

Let $(G, (1+r)w, \ell)$ for some constant r be a non-atomic congestion game in which the total amount for players of type i is $(1+r)w_i$ and the cost function on each resource e is ℓ_e . Our purpose is to bound the cost of an arbitrary equilibrium in (G, w, ℓ) by that of an optimal solution in $(G, (1+r)w, \ell)$ for some $r > 0$. Consider the following formulation (similar to the previous section) (\mathcal{P}_r) for $(G, (1+r)w, \ell)$. By weak duality, the optimal cost in $(G, (1+r)w, \ell)$ is at least the objective of a dual feasible solution in (\mathcal{D}_r) .

Primal:

$$\begin{aligned} \min \sum_{e, T_e} w(T_e) \ell_e(w(T_e)) z_{e, T_e} \quad (\mathcal{P}_r) \\ \sum_{j, k} a_k x_{ijk} &= (1+r) \cdot w_i \quad \forall i \\ \sum_{T_e} z_{e, T_e} &= 1 \quad \forall e \\ \sum_{T_e: (i, k) \in T_e} z_{e, T_e} &= \sum_{j: e \in s_{ij}} x_{ijk} \quad \forall (i, k), e \\ x_{ijk}, z_{e, T_e} &\in \{0, 1\} \quad \forall i, j, e, T_e \end{aligned}$$

Dual:

$$\begin{aligned} \max \sum_i (1+r) w_i \alpha_i + \sum_e \beta_e \quad (\mathcal{D}_r) \\ a_k \alpha_i &\leq \sum_{e: e \in s_{ij}} \gamma_{i, k, e} \quad \forall i, k, j \\ \beta_e + \sum_{(i, k) \in T_e} \gamma_{i, k, e} &\leq w(T_e) \ell_e(w(T_e)) \quad \forall e, T_e \end{aligned}$$

Hence, our scheme consists of bounding the cost of an arbitrary equilibrium in (G, w, ℓ) and the objective (\mathcal{D}_r) of an appropriate dual feasible solution.

Theorem 3.4 *In every non-atomic congestion game, for any constant $r > 0$, the cost of an equilibrium in (G, w, ℓ) is at most $1/r$ times the cost of an optimal solution in $(G, (1+r)w, \ell)$.*

Proof Let σ be a coarse correlated equilibrium of the game where the amount for players of type i is w_i . Construct the dual feasible solution for (\mathcal{D}_r) as in the proof of Theorem 3.3. As the dual constraints of (\mathcal{D}_r) and (\mathcal{D}_0) are the same, the construction in the proof of Theorem 3.3 gives a dual feasible solution for (\mathcal{D}_r) . It remains to bound the objective of (\mathcal{D}_r) of this dual solution to the cost of equilibrium σ , which is $\mathbb{E}_{f \sim \sigma} [\sum_e f_e \ell_e(f_e)]$. The former is

$$\begin{aligned} \sum_i (1+r)w_i \alpha_i + \sum_e \beta_e &= \mathbb{E}_{f \sim \sigma} \left[\sum_e \left((1+r) \cdot f_e \ell_e(f_e) + v_e \ell_e(v_e) - v_e \ell_e(f_e) \right) \right] \\ &\geq \mathbb{E}_{f \sim \sigma} \left[\sum_e r \cdot f_e \ell_e(f_e) \right] \end{aligned}$$

where the inequality holds since $f_e \ell_e(f_e) + v_e \ell_e(v_e) \geq v_e \ell_e(f_e)$. Precisely, if $f_e \geq v_e$ then $f_e \ell_e(f_e) \geq v_e \ell_e(f_e)$ and if $f_e < v_e$ then $v_e \ell_e(v_e) > v_e \ell_e(f_e)$ (since ℓ_e is non-decreasing). Hence, we deduce that the objective of (\mathcal{D}_r) is at least r times the cost of equilibrium σ . \square

3.2.3 Splittable Congestion Games

Model. In this section we consider the splittable congestion games also in the discrete setting. Fix an arbitrarily small constant $\epsilon > 0$. In a splittable congestion game, there is a set E of resources where each resource is associated to a non-decreasing differentiable cost function $\ell_e : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ such that $x \ell_e(x)$ is convex. There are n players, a player i has a set of strategies $\mathcal{S}_i \subseteq 2^E$ and has weight w_i , a multiple of ϵ . A strategy of player i is a distribution u^i of its weight w_i among strategies s_{ij} in \mathcal{S}_i such that $\sum_{s_{ij} \in \mathcal{S}_i} u_{s_{ij}}^i = w_i$ and $u_{s_{ij}}^i \geq 0$ is a multiple of ϵ . A strategy profile is a vector $\mathbf{u} = (u^1, \dots, u^n)$ of all players' strategies. We abuse notation and define $u_e^i = \sum_{e \in s_{ij}} u_{s_{ij}}^i$ as the load player i distributes on resource e and $u_e = \sum_{i=1}^n u_e^i$ the total load on e . Given a strategy profile \mathbf{u} , the cost of player i is defined as $C_i(\mathbf{u}) := \sum_e u_e^i \cdot \ell_e(u_e)$. A strategy profile \mathbf{u} is a pure Nash equilibrium if and only if for every player i and all $s_{ij}, s_{ij'} \in \mathcal{S}_i$ with $u_{s_{ij}}^i > 0$:

$$\sum_{e \in s_{ij}} (\ell_e(u_e) + u_e^i \cdot \ell'_e(u_e)) \leq \sum_{e \in s_{ij'}} (\ell_e(u_e) + u_e^i \cdot \ell'_e(u_e))$$

The proof of this equilibrium characterization can be found in [69]. Again, the more general concepts of mixed, correlated and coarse correlated equilibria are defined similarly as in Section 1.5.1. In the game, the social cost is defined as $C(\mathbf{u}) := \sum_{i=1}^n C_i(\mathbf{u}) = \sum_e u_e \ell_e(u_e)$.

The PoA bounds has been recently established for a large class of cost functions by Roughgarden and Schoppmann [110]. The authors proposed a *local smoothness* framework and showed that the local smoothness arguments give optimal PoA bounds for a large class of cost functions in splittable congestion games. Prior to

Roughgarden and Schoppmann [110], the works of Cominetti et al. [46] and Harks [69] have also the flavour of local smoothness, though their bounds are not tight. The local smooth arguments extends to the correlated equilibria of a game but not to the coarse correlated equilibria. Motivated by the duality approach, we define a new notion of smoothness and prove a bound on the PoA of coarse correlated equilibria. It turns out that this PoA bound for coarse correlated equilibria is indeed *tight* for all classes of scale-invariant cost functions by the lower bound given by Roughgarden and Schoppmann [110, Section 5]. A class of cost functions \mathcal{L} is *scale-invariant* if $\ell \in \mathcal{L}$ implies that $a \cdot \ell(b \cdot x) \in \mathcal{L}$ for every $a, b > 0$.

Formulation. Given a splittable congestion game, we formulate the problem by the same configuration program used for non-atomic congestion game. Denote a finite set of multiples of ϵ as $\{a_0, a_1, \dots, a_m\}$ where $a_k = k \cdot \epsilon$ and $m = \max_{i=1}^n \lceil w_i / \epsilon \rceil$. We say that T_e is a *configuration* of a resource e if $T_e = \{(i, k) : 1 \leq i \leq n, 0 \leq k \leq m\}$ in which a couple (i, k) specifies the player (i) and the amount a_k of the weight w_i that player i distributes to some strategy $s_{ij} \in \mathcal{S}_i$ with $e \in s_{ij}$. Intuitively, a configuration of a resource is a strategy profile of a game restricted to the resource. Let x_{ijk} be variable indicating whether player i distributes an amount a_k of its weight to strategy $s_{ij} \in \mathcal{S}_i$. For every resource e and a configuration T_e on resource e , let z_{e, T_e} be a variable such that $z_{e, T_e} = 1$ if and only if for $(i, k) \in T_e$, $x_{ijk} = 1$ for some $s_{ij} \in \mathcal{S}_i$ such that $e \in s_{ij}$. For a configuration T_e of resource e , denote $w(T_e)$ the total amount distributed by players in T_e to e .

$$\begin{array}{ll}
\min \sum_{e, T_e} w(T_e) \ell_e(w(T_e)) z_{e, T_e} & \max \sum_i w_i \alpha_i + \sum_e \beta_e \\
\sum_{j, k} a_k x_{ijk} = w_i & \forall i \quad a_k \alpha_i \leq \sum_{e: e \in s_{ij}} \gamma_{i, k, e} \quad \forall i, k, j \\
\sum_{T_e} z_{e, T_e} = 1 & \forall e \quad \beta_e + \sum_{(i, k) \in T_e} \gamma_{i, k, e} \leq w(T_e) \ell_e(w(T_e)) \\
\sum_{T_e: (i, k) \in T_e} z_{e, T_e} = \sum_{j: e \in s_{ij}} x_{ijk} & \forall (i, k), e \quad \forall e, T_e \\
x_{ij}, z_{e, T_e} \in \{0, 1\} & \forall i, j, e, T_e
\end{array}$$

Again, in the primal, the first constraint says that a player i distributes the total weight w_i among its strategies. The second constraint means that a resource e is always associated to a configuration (possibly empty). The third constraint guarantees that if a player i distributes an amount a_k to some strategy s_{ij} containing resource e then there must be a configuration T_e such that $(i, k) \in T_e$ and $z_{e, T_e} = 1$.

All previous duality proofs have the same structure: in the dual LP, the first constraint gives the characterization of an equilibrium and the second one settles the PoA bounds. Following this line, we give the following definition.

Definition 3.1 A cost function $\ell : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is (λ, μ) -dual-smooth if for every vectors $\mathbf{u} = (u_1, \dots, u_n)$ and $\mathbf{v} = (v_1, \dots, v_n)$, it holds that

$$v \ell(\mathbf{u}) + \sum_{i=1}^n u_i (v_i - u_i) \cdot \ell'(\mathbf{u}) \leq \lambda \cdot v \ell(\mathbf{v}) + \mu \cdot u \ell(\mathbf{u})$$

where $u = \sum_i u_i$ and $v = \sum_i v_i$. A splittable congestion game is (λ, μ) -dual-smooth if every resource e in the game, function ℓ_e is (λ, μ) -dual-smooth.

Theorem 3.5 For every (λ, μ) -dual-smooth splittable congestion game G , the price of anarchy of coarse correlated equilibria of G is at most $\lambda/(1 - \mu)$. This bound is tight for the class of scalable cost functions.

Proof The proof follows the duality scheme.

Dual Variables. Fix parameter λ and μ . Given a coarse correlated equilibrium σ , define corresponding dual variables as follows.

$$\begin{aligned}\alpha_i &= \frac{1}{\lambda} \mathbb{E}_{\mathbf{u} \sim \sigma} \left[\sum_{e \in \mathcal{S}_{ij}} \ell_e(u_e) + u_e^i \ell'_e(u_e) \right] \text{ for some } s_{ij} \in \mathcal{S}_i : u_{s_{ij}}^i > 0, \\ \beta_e &= -\frac{1}{\lambda} \mathbb{E}_{\mathbf{u} \sim \sigma} \left[\mu \cdot u_e \ell_e(u_e) + \sum_i (u_e^i)^2 \cdot \ell'_e(u_e) \right], \\ \gamma_{i,k,e} &= \frac{1}{\lambda} \mathbb{E}_{\mathbf{u} \sim \sigma} [a_k (\ell_e(u_e) + u_e^i \ell'_e(u_e))].\end{aligned}$$

The dual variables have similar interpretations as in previous analysis. Up to some constant factors, variable α_i is the marginal cost of a strategy used by player i in the equilibrium; and $\gamma_{i,k,e}$ represents an estimation of the cost of player i on resource e if player i distributes an amount a_k of its weight to some strategy containing e while other players follow their strategies in the equilibrium.

Feasibility. By this definition of dual variables, the first dual constraint holds since it is the definition of coarse correlated equilibrium. Rearranging the terms, the second dual constraint for a resource e and a configuration T_e reads

$$\begin{aligned}\frac{1}{\lambda} \sum_{(i,k) \in T_e} \mathbb{E}_{\mathbf{u} \sim \sigma} [a_k \cdot \ell_e(u_e) + u_e^i (a_k - u_e^i) \ell'_e(u_e)] \\ \leq w(T_e) \ell_e(w(T_e)) + \frac{\mu}{\lambda} \mathbb{E}_{\mathbf{u} \sim \sigma} [u_e \ell_e(u_e)].\end{aligned}$$

This inequality follows directly from the definition of (λ, μ) -dual-smoothness and linearity of expectation (and note that $w(T_e) \ell_e(w(T_e)) = \mathbb{E}_{\mathbf{u} \sim \sigma} [w(T_e) \ell_e(w(T_e))]$ and $w(T_e) = \sum_{(i,k) \in T_e} a_k$).

Bounding primal and dual. By the definition of dual variables, the dual objective is

$$\begin{aligned}\sum_i w_i \alpha_i + \sum_e \beta_e &= \sum_e \left(\sum_i u_e^i \alpha_i + \beta_e \right) \\ &= \frac{1}{\lambda} \mathbb{E}_{\mathbf{u} \sim \sigma} \left[\sum_e u_e \ell_e(u_e) + \sum_i (u_e^i)^2 \cdot \ell'_e(u_e) \right] - \frac{1}{\lambda} \mathbb{E}_{\mathbf{u} \sim \sigma} \left[\mu \cdot u_e \ell_e(u_e) + \sum_i (u_e^i)^2 \cdot \ell'_e(u_e) \right] \\ &= \frac{1 - \mu}{\lambda} \mathbb{E}_{\mathbf{u} \sim \sigma} \left[\sum_e u_e \ell_e(u_e) \right]\end{aligned}$$

while the cost of the equilibrium σ is $\mathbb{E}_{\mathbf{u} \sim \sigma} [\sum_e u_e \ell_e(u_e)]$. The theorem follows. \square

3.3 Efficiency of Smooth Auctions in Welfare Maximization

In this section, we show that the primal-dual approach also captures the smoothness framework in studying the inefficiency of Bayes-Nash equilibria in incomplete-information settings. Recall the definition of smooth games.

Definition 1.8 ([109]) For parameters $\lambda, \mu \geq 0$, an auction is (λ, μ) -smooth if for every valuation profile $\mathbf{v} = (v_1, \dots, v_n)$, there exist action distributions $D_1^*(\mathbf{v}), \dots, D_n^*(\mathbf{v})$ over $\mathcal{A}_1, \dots, \mathcal{A}_n$ such that, for every action profile \mathbf{a} ,

$$\sum_i \mathbb{E}_{a_i^* \sim D_i^*(\mathbf{v})} [u_i(a_i^*, \mathbf{a}_{-i}; v_i)] \geq \lambda \cdot \text{SW}(\mathbf{a}^*; \mathbf{v}) - \mu \cdot \text{SW}(\mathbf{a}; \mathbf{v}). \quad (1.3)$$

Definition 1.9 ([119]) For parameters $\lambda, \mu \geq 0$, an auction is (λ, μ) -smooth if for every valuation profile $\mathbf{v} = (v_1, \dots, v_n)$, there exist action distributions $D_1^*(\mathbf{v}), \dots, D_n^*(\mathbf{v})$ over $\mathcal{A}_1, \dots, \mathcal{A}_n$ such that, for every action profile \mathbf{a} ,

$$\sum_i \mathbb{E}_{a_i^* \sim D_i^*(\mathbf{v})} [u_i(a_i^*, \mathbf{a}_{-i}; v_i)] \geq \lambda \cdot \text{OPT}(\mathbf{v}) - \mu \cdot \text{REV}(\mathbf{a}; \mathbf{v}). \quad (1.4)$$

In this section, we consider Definition 1.8 of smooth auctions in [109] and revisit the price of anarchy bound of smooth auctions. In the end of the section, we show that a similar proof carries through the smooth auctions defined by Syrgkanis and Tardos [119].

Theorem 3.6 ([109]) If an auction is (λ, μ) -smooth and the distributions of player valuations are independent then every Bayes-Nash equilibrium has expected welfare at least $\frac{\lambda}{1+\mu}$ times the optimal expected welfare.

Proof Given an auction, we formulate the corresponding optimization problem by a configuration LP. A configuration A consists of pairs (i, a_i) specifying that in configuration A , player i chooses action a_i . Intuitively, a configuration is an action profile of players. For every player i , every valuation $v_i \in \mathcal{V}_i$ and every action $a_i \in \mathcal{A}_i$, let $x_{i,a_i}(v_i)$ be the variable representing the probability that player i chooses action a_i . Besides, for every valuation profile \mathbf{v} , let $z_A(\mathbf{v})$ be the variable indicating the probability that the chosen configuration (action profile) is A . For each configuration A and valuation profile \mathbf{v} , the auctioneer outcomes an allocation and a payment and that results in a social welfare denoted as $c_A(\mathbf{v})$. In the other words, if \mathbf{a} is the action profile corresponding to the configuration A then $c_A(\mathbf{v})$ is in fact $\text{SW}(\mathbf{a}; \mathbf{v})$. Consider the following formulation.

$$\begin{aligned} & \max \sum_{\mathbf{v}} c_A(\mathbf{v}) z_A(\mathbf{v}) \\ & \sum_{a_i \in \mathcal{A}_i} x_{i,a_i}(v_i) \leq f_i(v_i) && \forall i, v_i \\ & \sum_A z_A(\mathbf{v}) \leq f(\mathbf{v}) && \forall \mathbf{v} \\ & \sum_{A: (i,a_i) \in A} z_A(v_i, \mathbf{v}_{-i}) \leq f_{-i}(\mathbf{v}_{-i}) \cdot x_{i,a_i}(v_i) && \forall i, a_i, v_i, \mathbf{v}_{-i} \\ & x_{i,a_i}(v_i), z_A(\mathbf{v}) \geq 0 && \forall i, a_i, A, v_i, \mathbf{v} \end{aligned}$$

The first and second constraints guarantee that variables x and z represent indeed the probability distribution of each player and the joint distribution, respectively.

The third constraint makes the connection between variables x and z . It ensures that if a player i with valuation v_i selects some action a_i then in the valuation profile (v_i, \mathbf{v}_{-i}) , the probability that the configuration A contains (i, a_i) must be $f_{-i}(\mathbf{v}_{-i}) \cdot x_{i,a_i}(v_i)$. The primal objective is the expected welfare of the auction.

The dual is the following.

$$\begin{aligned} \min \quad & \sum_{i,v_i} f_i(v_i) \cdot \alpha_i(v_i) + \sum_{\mathbf{v}} f(\mathbf{v}) \cdot \beta(\mathbf{v}) \\ \alpha_i(v_i) \geq \quad & \sum_{\mathbf{v}_{-i}} f_{-i}(\mathbf{v}_{-i}) \cdot \gamma_{i,a_i}(v_i, \mathbf{v}_{-i}) & \forall i, a_i, v_i \\ \beta(\mathbf{v}) + \sum_{(i,a_i) \in A} \gamma_{i,a_i}(\mathbf{v}) \geq \quad & c_A(\mathbf{v}) & \forall A, \mathbf{v} \\ \alpha_i(v_i), \beta(\mathbf{v}), \gamma_{i,a_i}(\mathbf{v}) \geq \quad & 0 & \forall i, v_i, \mathbf{v} \end{aligned}$$

Construction of dual variables. Assuming that the auction is (λ, μ) -smooth. Fix the parameters λ and μ . Given an arbitrary Bayes-Nash equilibrium σ , define the dual variables as follows.

$$\begin{aligned} \alpha_i(v_i) &:= \frac{1}{\lambda} \mathbb{E}_{v_{-i}} [\mathbb{E}_{\mathbf{b} \sim \sigma(v_i, \mathbf{v}_{-i})} [u_i(\mathbf{b}; v_i)]], \\ \beta(\mathbf{v}) &:= \frac{\mu}{\lambda} \mathbb{E}_{\mathbf{b} \sim \sigma(\mathbf{v})} [\text{SW}(\mathbf{b}; \mathbf{v})], \\ \gamma_{i,a_i}(\mathbf{v}) &:= \frac{1}{\lambda} \mathbb{E}_{\mathbf{b}_{-i} \sim \sigma_{-i}(\mathbf{v}_{-i})} [u_i(a_i, \mathbf{b}_{-i}; v_i)]. \end{aligned}$$

Informally, up to some constant factors depending on λ and μ , $\alpha_i(v_i)$ is the expected utility of player i in equilibrium σ ; $\beta(\mathbf{v})$ stands for the social welfare of the auction where the valuation profile is \mathbf{v} and players follow the equilibrium actions $\sigma(\mathbf{v})$; and $\gamma_{i,a_i}(\mathbf{v})$ represents the utility of player i in valuation profile \mathbf{v} if player i chooses action a_i while other players $i' \neq i$ follows their equilibrium strategies $\sigma_{-i}(\mathbf{v}_{-i})$.

Feasibility. We show that the constructed dual variables form a feasible solution. By the definition of dual variables, the first dual constraint reads

$$\begin{aligned} \frac{1}{\lambda} \mathbb{E}_{v_{-i}} [\mathbb{E}_{\mathbf{b} \sim \sigma(\mathbf{v})} [u_i(\mathbf{b}; v_i)]] &\geq \frac{1}{\lambda} \sum_{\mathbf{v}_{-i}} f_{-i}(\mathbf{v}_{-i}) \cdot \mathbb{E}_{\mathbf{b}_{-i} \sim \sigma_{-i}(\mathbf{v}_{-i})} [u_i(a_i, \mathbf{b}_{-i}; v_i)] \\ &= \frac{1}{\lambda} \mathbb{E}_{v_{-i}} [\mathbb{E}_{\mathbf{b}_{-i} \sim \sigma_{-i}(\mathbf{v}_{-i})} [u_i(a_i, \mathbf{b}_{-i}; v_i)]] \end{aligned}$$

This is exactly the definition of a Bayes-Nash equilibrium.

For every valuation profile $\mathbf{v} = (v_1, \dots, v_n)$ and for any configuration A (corresponding action profile $\mathbf{a} = (a_1, \dots, a_n)$), the second constraint reads:

$$\frac{\mu}{\lambda} \mathbb{E}_{\mathbf{b} \sim \sigma(\mathbf{v})} [\text{SW}(\mathbf{b}; \mathbf{v})] + \sum_{(i,a_i) \in A} \frac{1}{\lambda} \mathbb{E}_{\mathbf{b}_{-i} \sim \sigma_{-i}(\mathbf{v}_{-i})} [u_i(a_i, \mathbf{b}_{-i}; v_i)] \geq \text{SW}(\mathbf{a}; \mathbf{v}). \quad (3.1)$$

Note that we can write $\text{SW}(\mathbf{a}; \mathbf{v}) = \mathbb{E}_{\mathbf{b} \sim \sigma(\mathbf{v})} [\text{SW}(\mathbf{a}; \mathbf{v})]$. For any fixed realization \mathbf{b} of $\sigma(\mathbf{v})$, by (λ, μ) -smoothness

$$\frac{\mu}{\lambda} \text{SW}(\mathbf{b}; \mathbf{v}) + \sum_i \frac{1}{\lambda} u_i(a_i, \mathbf{b}_{-i}; v_i) \geq \text{SW}(\mathbf{a}; \mathbf{v}).$$

Hence, by taking expectation over $\sigma(\mathbf{v})$, Inequality (3.1) follows.

Price of Anarchy. The welfare of equilibrium σ is $\mathbb{E}_{\mathbf{v}} \mathbb{E}_{\mathbf{b} \sim \sigma(\mathbf{v})} [\text{SW}(\mathbf{b}; \mathbf{v})]$ while the dual objective of the constructed dual variables is

$$\begin{aligned} & \sum_{i, v_i} f_i(v_i) \cdot \frac{1}{\lambda} \mathbb{E}_{v_i} [\mathbb{E}_{\mathbf{b} \sim \sigma(\mathbf{v})} [u_i(\mathbf{b}; v_i)]] + \sum_{\mathbf{v}} f(\mathbf{v}) \cdot \frac{\mu}{\lambda} \mathbb{E}_{\mathbf{b} \sim \sigma(\mathbf{v})} [\text{SW}(\mathbf{b}; \mathbf{v})] \\ & \leq \frac{1 + \mu}{\lambda} \cdot \mathbb{E}_{\mathbf{v}} \mathbb{E}_{\mathbf{b} \sim \sigma(\mathbf{v})} [\text{SW}(\mathbf{b}; \mathbf{v})]. \end{aligned}$$

Therefore, the PoA of a (λ, μ) -smooth auction is at most $\lambda / (1 + \mu)$. \square

Remark. Consider the notion of (λ, μ) -smooth auctions defined by Syrgkanis and Tardos [119]. In order to bound the price of anarchy, Inequality (1.4) can be replaced by a weaker one, which is:

$$\sum_i \mathbb{E}_{a_i^* \sim D_i^*(\mathbf{v})} [u_i(a_i^*, \mathbf{a}_{-i}; v_i)] \geq \lambda \cdot \text{SW}(\mathbf{a}^*; \mathbf{v}) - \mu \cdot \text{REV}(\mathbf{a}; \mathbf{v}). \quad (3.2)$$

Using the same proof structure of Theorem 3.6, we can prove that the price of anarchy is at most λ / μ [119]. Specifically, define dual variables α and γ as previous and

$$\beta(\mathbf{v}) = \frac{\mu}{\lambda} \mathbb{E}_{\mathbf{b} \sim \sigma(\mathbf{v})} [\text{REV}(\mathbf{b}; \mathbf{v})].$$

The feasibility follows the definitions of Bayes-Nash equilibria and smooth auctions, in particular Inequality (3.2). To bound the price of anarchy, as $\mu \geq 1$, we have

$$\begin{aligned} & \sum_{i, v_i} f_i(v_i) \cdot \frac{1}{\lambda} \mathbb{E}_{v_i} [\mathbb{E}_{\mathbf{b} \sim \sigma(\mathbf{v})} [u_i(\mathbf{b}; v_i)]] + \sum_{\mathbf{v}} f(\mathbf{v}) \cdot \frac{\mu}{\lambda} \mathbb{E}_{\mathbf{b} \sim \sigma(\mathbf{v})} [\text{REV}(\mathbf{b}; \mathbf{v})] \\ & \leq \frac{\mu}{\lambda} \cdot \mathbb{E}_{\mathbf{v}} \mathbb{E}_{\mathbf{b} \sim \sigma(\mathbf{v})} [\text{SW}(\mathbf{b}; \mathbf{v})]. \end{aligned}$$

Therefore, the price of anarchy is at most λ / μ .

3.4 Simultaneous Item-Bidding Auctions

Model. In this section, we consider the following Bayesian combinatorial auctions. In the setting, there are m items to be sold to n players. Each player i has a private monotone valuation $v_i : 2^{[m]} \rightarrow \mathbb{R}^+$ over different subsets of items $S \subset 2^{[m]}$. For simplicity, we denote $v_i(S)$ as v_{iS} . The valuation profile $\mathbf{v} = (v_1, \dots, v_n)$ is drawn from a *product* distribution \mathbf{F} . In other words, the probability distributions F_i of valuations v_i are independent. Designing efficient combinatorial auctions are in general complex and a major direction in literature is to seek simple and efficient auctions in term of PoA. Among others, simultaneous item-bidding auctions are of particular interest.

We consider two forms of simultaneous item-bidding auctions: *simultaneous first-price auctions (S1A)* and *simultaneous second-price auctions (S2A)*. In the auctions, each player submits simultaneously a vector of bids, one for each item. A typical assumption is the *non-overbidding* property in which each player submits a vector b_i of bids such that for any set of items S , $\sum_{j \in S} b_{ij} \leq v_{iS}$. Given the bid profile, each item is allocated to the player with highest bid. In a simultaneous first-price auction, the

payment of the winner of each item is its bid on the item; while in a simultaneous second-price auction, the winner of each item pays the second highest bid on the item.

3.4.1 Connection between Primal-Dual and Non-Smooth Techniques

In this section, we consider the setting in which all player valuations are *sub-additive*. That is, $v_i(S \cup T) \leq v_i(S) + v_i(T)$ for every player i and all subsets $S, T \subset 2^{[m]}$. The PoA of simultaneous item-bidding auctions has been widely studied in this setting. Using smoothness framework in auctions, logarithmic bounds on PoA for S1A and S2A are given by Hassidim et al. [71] and Bhawalkar and Roughgarden [22], respectively. Recently, Feldman et al. [54] presented a significant improvement by establishing the PoA bounds 2 and 4 for S1A and S2A, respectively. Their proof arguments go beyond the smoothness framework. In the following, we revisit the results of Feldman et al. [54] and show that the duality approach captures the non-smooth technique from [54].

Formulation. Given a valuation profile \mathbf{v} , let $\bar{x}_{ij}(\mathbf{v})$ be the variable indicating whether player i receives item j in valuation profile \mathbf{v} . Let $\bar{z}_{iS}(\mathbf{v})$ be the variable indicating whether player i receives a set of items S . Then for any profile \mathbf{v} and for any item j , $\sum_i \bar{x}_{ij}(\mathbf{v}) \leq 1$, meaning that an item j is allocated to at most one player. Moreover, $\sum_{S:j \in S} \bar{z}_{iS}(\mathbf{v}) = \bar{x}_{ij}(\mathbf{v})$, meaning that if player i receives item j then some subset of items S allocated to i must contain j . Besides, $\sum_S \bar{z}_{iS}(\mathbf{v}) = 1$ since some subset of items (possibly empty) is allocated to i .

Let $x_{ij}(v_i)$ and $z_{iS}(v_i)$ be *interim* variables corresponding to $\bar{x}_{ij}(\mathbf{v})$ and $\bar{z}_{iS}(\mathbf{v})$ and are defined as follows:

$$x_{ij}(v_i) := \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}} [\bar{x}_{ij}(v_i, \mathbf{v}_{-i})], \quad z_{iS}(v_i) := \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}} [\bar{z}_{iS}(v_i, \mathbf{v}_{-i})],$$

where \mathbf{F}_{-i} is the product distribution of all players other than i . Consider the following relaxation with interim variables. The constraints in the primal follow the relationship between the interim variables $x_{ij}(v_i), z_{iS}(v_i)$ and variables $\bar{x}_{ij}(\mathbf{v}), \bar{z}_{iS}(\mathbf{v})$.

$$\begin{aligned} & \max \sum_{i,S} \sum_{v_i} f_i(v_i) [v_{iS} \cdot z_{iS}(v_i)] \\ & \sum_i \sum_{v_i \in V_i} f_i(v_i) x_{ij}(v_i) \leq 1 && \forall j \\ & \sum_S z_{iS}(v_i) = 1 && \forall i, v_i \\ & \sum_{S:j \in S} z_{iS}(v_i) = x_{ij}(v_i) && \forall i, j, v_i \\ & x_{ij}(v_i), z_{iS}(v_i) \geq 0 && \forall i, j, S, v_i \end{aligned}$$

The dual is the following.

$$\begin{aligned}
\min \quad & \sum_{i,v_i} \alpha_i(v_i) + \sum_j \beta_j \\
& f_i(v_i) \cdot \beta_j \geq \gamma_{i,j}(v_i) && \forall i, j, v_i \\
\alpha_i(v_i) + \sum_{j \in S} \gamma_{i,j}(v_i) & \geq f_i(v_i) \cdot v_{iS} && \forall i, S, v_i \\
& \alpha_i(v_i) \geq 0 && \forall i, v_i
\end{aligned}$$

Dual Variables. Fix a Bayes-Nash equilibrium σ . Given a valuation v , denote $\mathbf{b} = (b_1, \dots, b_n) = \sigma(v)$ as the bid equilibrium. Let \mathbf{B} be the distribution of \mathbf{b} over the randomness of v and σ . Let $\mathbf{B}(v_i)$ be the distribution of \mathbf{b} over the randomness of v and σ while the valuation v_i of player i is fixed. Since v_i and v_{-i} are independent and each σ_i is a mapping $\mathcal{V}_i \rightarrow \Delta(\mathcal{A}_i)$, strategy b_i is independent of \mathbf{b}_{-i} . Let \mathbf{B}_{-i} be the distribution of \mathbf{b}_{-i} . We define the dual variables as follows.

Let $\alpha_i(v_i)$ be proportional to the expected utility of player i with valuation v_i , over the randomness of valuations v_{-i} of other players. Specifically,

$$\alpha_i(v_i) := 2f_i(v_i) \cdot \mathbb{E}_{v_{-i} \sim F_{-i}} [\mathbb{E}_{\sigma} [u_i(\sigma(v_i, v_{-i}), v_i)]] = 2f_i(v_i) \cdot \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i)} [u_i(\mathbf{b}, v_i)]$$

Besides, let $\gamma_{i,j}(v_i)$ be proportional to the expected value of the bid on item j if player i with valuation v_i wants to win item j while other players follow the equilibrium strategies. Formally,

$$\gamma_{i,j}(v_i) := 2f_i(v_i) \cdot \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}} \left[\max_{k \neq i} b_{kj} \right].$$

Finally, define $\beta_j := 2 \max_i \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}} [\max_{k \neq i} b_{kj}]$.

The following lemma shows the feasibility of the variables. The main core of the proof relies on an argument in [54].

Lemma 3.1 *The dual vector (α, β, γ) defined above constitutes a dual feasible solution.*

Proof The first dual constraint follows immediately by the definitions of dual variables β and γ . We are now proving the second dual constraint. Fix a player i with sub-additive valuation v_i and assume that $f_i(v_i) > 0$ (otherwise, it is trivial). By [54] (or see [107, Lemma 1.3] for another clear exposition), for any set of items S , there exists an action b_i^* such that

$$\mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}} \left[u_i((b_i^*, \mathbf{b}_{-i}), v_i) \right] + \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}} \left[\sum_{j \in S} \max_{k \neq i} b_{kj} \right] \geq \frac{1}{2} v_{iS}.$$

Moreover, the first term in the left-hand side is at most the utility of player i with valuation v_i since (b_i, \mathbf{b}_{-i}) is a Bayes-Nash equilibrium. Therefore,

$$\mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i)} [u_i(\mathbf{b}, v_i)] + \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i)} \left[\sum_{j \in S} \max_{k \neq i} b_{kj} \right] \geq \frac{1}{2} v_{iS}.$$

By the definition of dual variables, this inequality is exactly the second constraint when multiplying both sides by $2f_i(v_i)$. \square

Theorem 3.7 ([54]) *If player valuations are sub-additive then every Bayes-Nash equilibrium of a S1A (or S2A) has expected welfare at least 1/2 (or 1/4, resp) of the optimal one.*

Proof For an item j , let $i^*(j) \in \arg \max_i \mathbb{E}_{v_{-i} \sim F_{-i}} [\max_{k \neq i} b_{kj}]$. Hence,

$$\begin{aligned} \beta_j &= 2\mathbb{E}_{v_{-i^*(j)} \sim F_{-i^*(j)}} \mathbb{E}_{\sigma} \left[\max_{k \neq i^*(j)} b_{kj} \right] = 2\mathbb{E}_{v_{i^*(j)} \sim F_i} \mathbb{E}_{v_{-i^*(j)} \sim F_{-i^*(j)}} \mathbb{E}_{\sigma} \left[\max_{k \neq i^*(j)} b_{kj} \right] \\ &= 2\mathbb{E}_{v \sim F} \mathbb{E}_{\sigma} \left[\max_{k \neq i^*(j)} b_{kj} \right], \end{aligned}$$

where the second equality is due to the fact that the term $\mathbb{E}_{v_{-i^*(j)} \sim F_{-i^*(j)}} \mathbb{E}_{\sigma} [\max_{k \neq i^*(j)} b_{kj}]$ is independent of $v_{i^*(j)}$. Therefore, the dual objective is

$$\sum_{i, v_i} \alpha_i(v_i) + \sum_j \beta_j = 2\mathbb{E}_{v \sim F} \mathbb{E}_{\sigma} \left[\sum_i u_i(\mathbf{b}, v_i) + \sum_j \max_{k \neq i^*(j)} b_{kj} \right]$$

Fix a random choice of profile v and σ (so the bid profile \mathbf{b} is fixed). We bound the dual objective, i.e., the right-hand side of the above equality, in S1A and S2A. Note that the utility of a player winning no item is 0.

First Price Auction. Partition the set of items into the winning items of each player. Consider a player i with the set of winning items S . The utility of this player i is $v_{iS} - \sum_{j \in S} \max_k b_{kj}$. Hence, $v_{iS} - \sum_{j \in S} b_{ij} + \sum_{j \in S} \max_{k \neq i^*(j)} b_{kj} \leq v_{iS}$ since by the allocation rule, $b_{ij} = \max_k b_{kj}$ for every $j \in S$. Hence, summing over all players, the dual objective is bounded by twice the total expected valuation of winning players, which is the primal. So the price of anarchy is at most 2.

Second Price Auction. Similarly, consider a player i with the set of winning items S . The utility of player i as well as its payment (by no-overbidding) are at most v_{iS} . Therefore, summing over all players, the dual objective is bounded by four times the total expected valuation of winning players. Hence, the price of anarchy is at most 4. \square

Remark. The non-overbidding assumption, a risk-aversion assumption, is given in order to prevent players from suffering negative utility while receiving items. We use this assumption in the proof only in settling the ratio between the primal and the dual; specifically to argue that the payment of a player does not exceed its valuation on the received items. The above analysis holds even without this assumption in the following sense. Assume that players are allowed to bid up to a constant r times their valuation (hence, players risk to have negative utility). Then, the PoA for S2A is $2(1+r)$.

3.4.2 Connection between Primal-Dual and No-Envy Learning

Very recently, Daskalakis and Syrgkanis [49] have introduced *no-envy learning* — a novel concept of learning in auctions. The notion is inspired by the concept of Walrasian equilibrium and it is motivated by the fact that no-regret learning algorithms (which converge to coarse correlated equilibria) for the simultaneous item-bidding auctions are computationally inefficient as the number of player actions are exponential. When the players have fractionally sub-additive (XOS) valuation, Daskalakis

and Syrgkanis [49] showed that no-envy outcomes are a relaxation of no-regret outcomes. Moreover, no-envy outcomes maintain the approximate welfare optimality of no-regret outcomes while ensuring the computational tractability. In this section, we explore the connection between the no-envy learning and the primal-dual approach. Indeed, the notion of no-envy learning would be naturally derived from the dual constraints very much in the same way as the smoothness argument is.

We recall the notion of no-envy learning algorithms [49]. We first define the *online learning problem*. In the online learning problem, at each step t , the player chooses a bid vector $b^t = (b_1^t, \dots, b_m^t)$ where b_j^t is the bid on item j for $1 \leq j \leq m$; and the adversary picks adaptively (depending on the history of the play but not on the current bid b^t) a threshold vector $\theta^t = (\theta_1^t, \dots, \theta_m^t)$. The player wins the set $S^*(b^t, \theta^t) = \{j : b_j^t \geq \theta_j^t\}$ and gets reward:

$$u(b^t, \theta^t) := v(S^*(b^t, \theta^t)) - \sum_{j \in S^*(b^t, \theta^t)} \theta_j^t$$

where $v : 2^{[m]} \rightarrow \mathbb{R}$ is the valuation of the player.

Definition 3.2 ([49]) *An algorithm for the online learning problem is r -approximate no-envy if, for any adaptively chosen sequence of (random) threshold vector $\theta^{1:T}$ by the adversary, the (random) bid vector $b^{1:T}$ chosen by the algorithm satisfies:*

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[u(b^t, \theta^t)] \geq \max_{S \subseteq [m]} \left(\frac{1}{r} \cdot v(S) - \sum_{j \in S} \frac{1}{T} \sum_{t=1}^T \mathbb{E}[\theta_j^t] \right) - \epsilon(T) \quad (3.3)$$

where the no-envy rate $\epsilon(T) \rightarrow 0$ when $T \rightarrow \infty$. An algorithm is no-envy if it is 1-approximate no-envy.

Now we show the connection between primal-dual and no-envy learning by revisiting the following theorem. As we will see, the notion of no-envy learning corresponds exactly to a constraint of the dual program.

Theorem 3.8 ([49]) *If n players in a S2A use a r -approximate no-envy learning algorithm with envy rate $\epsilon(T)$ then in T steps, the average welfare is at least $\frac{1}{2r} \text{OPT} - n \cdot \epsilon(T)$ where OPT is the expected optimal welfare.*

Proof Let b_i^t be the bid vector of player i where b_{ij}^t is the bid of player i on item j in step t . In a S2A the threshold $\theta_{ij}^t = \max_{k \neq i} b_{kj}^t$. Consider the same primal and dual LPs in Section 3.4.1.

Dual variables. Recall that r is the approximation factor and $\epsilon(T)$ the no-envy rate of the learning algorithm. Define dual variables (similar to the ones in Section 3.4.1)

as follows.

$$\begin{aligned}
\alpha_i(v_i) &:= r \cdot f_i(v_i) \cdot \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}} \left[\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathbf{b}^t(v_i, \mathbf{v}_{-i})} [u_i(b_i^t, \theta_i^t)] \right] + r \cdot \epsilon(T) \\
\gamma_{i,j}(v_i) &:= r \cdot f_i(v_i) \cdot \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}} \left[\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathbf{b}^t(v_i, \mathbf{v}_{-i})} [\theta_{ij}^t] \right] \\
&= r \cdot f_i(v_i) \cdot \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}} \left[\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathbf{b}_{-i}^t(\mathbf{v}_{-i})} [\theta_{ij}^t] \right] \\
\beta_j &:= r \cdot \max_i \max_{v_i} \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}} \left[\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathbf{b}^t(v_i, \mathbf{v}_{-i})} [\theta_{ij}^t] \right] \\
&= r \cdot \max_i \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}} \left[\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathbf{b}_{-i}^t(\mathbf{v}_{-i})} [\theta_{ij}^t] \right]
\end{aligned}$$

where the second equalities in the definitions of γ and β follow the fact that player valuations are independent and θ_{ij}^t does not depend on b_{ij}^t for every i, j .

Feasibility. The first dual constraint follows immediately by the definitions of dual variables β and γ . For a fixed set S and a player i with valuation v_i , the second dual constraint reads

$$\begin{aligned}
r \cdot f_i(v_i) \cdot \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}} \left[\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathbf{b}^t(v_i, \mathbf{v}_{-i})} [u_i(b_i^t, \theta_i^t)] \right] + r \cdot \epsilon(T) \\
+ r \cdot \sum_{j \in S} f_i(v_i) \cdot \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}} \left[\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathbf{b}_{-i}^t(\mathbf{v}_{-i})} [\theta_{ij}^t] \right] \geq f_i(v_i) \cdot v_{iS}.
\end{aligned}$$

This inequality follows immediately from the definition of r -approximate no-envy learning algorithms (specifically, Inequality (3.3)) by simplifying and rearranging terms. (Note that $\mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}} [f_i(v_i) \cdot v_{iS}] = f_i(v_i) \cdot v_{iS}$.)

Bounding the cost. In T steps, the average welfare is

$$\mathbb{E}_{\mathbf{v}} \left[\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathbf{b}^t(\mathbf{v})} [v_i(b_i^t, \theta_i^t)] \right] = \mathbb{E}_{\mathbf{v}} \left[\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathbf{b}^t(\mathbf{v})} [v_i(S^*(b_i^t, \theta_i^t))] \right].$$

Besides, in the dual objective,

$$\begin{aligned}
\sum_{i, v_i} \alpha_i(v_i) &\leq r \cdot \mathbb{E}_{\mathbf{v}} \left[\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathbf{b}^t(\mathbf{v})} [v_i(S^*(b_i^t, \theta_i^t))] \right] + n \cdot r \cdot \epsilon(T), \\
\sum_j \beta_j &\leq r \cdot \mathbb{E}_{\mathbf{v}} \left[\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathbf{b}^t(\mathbf{v})} [v_i(S^*(b_i^t, \theta_i^t))] \right]
\end{aligned}$$

where the last inequality is due to the non-overbidding property. Hence, the theorem follows by weak duality. \square

3.5 Sequential Auctions

3.5.1 Sequential Second Price Auctions in Sponsored Search

Model. In the sponsored search problem, there are n players and n slots for online advertisement. Each player i has a *private valuation* v_i , representing its valuation per click. We use $\mathbf{v} = (v_1, \dots, v_n)$ to denote the *valuation profile* of players. Additionally, each player i has a *quality factor* α_i that reflect the click-ability of the ad. The couple of valuation and quality factor (v_i, α_i) of player i is drawn from a publicly known distribution F_i . In the model, we assume that the distributions F_i 's are mutually independent. The slots have associated *click-through-rates* $\beta_1 \geq \beta_2 \geq \dots \geq \beta_n$. An *outcome* is an one-to-one assignment of slots to players. When player i is assigned to the j -th slot, the player gets $\alpha_i \beta_j$ clicks.

In the auction, the auctioneer sells slots sequentially one-by-one in non-increasing order of β_j via the second price mechanisms. At the consideration of slot j , the auctioneer collects all the bid b_{ij} on item j from every player i , which is interpreted as a valuation declaration. We also assume that the non-overbidding property, meaning that $b_{ij} \leq v_i$ for all i and j . The auctioneer then assigns slot j to the player (that has not received any slot so far) with highest *effective* bid, defined as $\alpha_i b_i$. The payment of the winning player is set according to *critical* value: the smallest bid that guarantees the player still gets the slot. Specifically, if a slot j is assigned to player i then the payment of i is $p_i = \alpha_{i'} \beta_{i'} / \alpha_i$ where $\alpha_{i'} \beta_{i'}$ is the second highest effective bid on slot j . The *utility* of player i is $\alpha_i \beta_j (v_i - p_i)$. The *social welfare* of the outcome is $\sum_{i,j} \beta_j \alpha_i v_i$ where the sum is taken over all player i with their allocated slots j .

This setting is captured by extensive form games (see [63, 105] for comprehensive treatments). The strategy of each player is an adaptive bidding policy: the bid of player i for slot j is a function of its valuation v_i , the common knowledge about the distributions of player valuations \mathbf{F} and the history h_j of outcomes in auctions before the consideration of slot j . Thus a player strategy can be denoted as $b_{ij}(v_i, h_j)$. We are interested in the *perfect Bayesian equilibria* which is a refinement of the concepts of Bayes-Nash equilibria and subgame perfect equilibria. A profile of bidding policies is a *perfect Bayesian equilibrium* if it is a Bayes-Nash equilibrium of the original game and given an arbitrary history (of some t first rounds), the policy profile remains also a Bayes-Nash equilibrium of this induced game.

The sponsored search problem has been extensively studied via the generalized second-price (GSP) auctions. It was first considered by Mehta et al. [95] from optimization perspective and was proposed simultaneously by Edelman et al. [53] and Varian [121] from a game theoretical point of view (see [89, 93] for surveys on the topic). Recently, Caragiannis et al. [41] have proved the PoA upper bound of 2.927 (without the independence assumption on distributions F_i 's), the currently best known PoA bound, using a technique called semi-smoothness, an extension of the smoothness framework in [108]. The study of PoA in sequential auctions has been initiated by Leme et al. [90]. The authors studied sequential first price auctions for matching markets and matroid auctions in the full-information environments and showed that the PoA (of pure Nash equilibria) is at most 2. Subsequently, Syrgkanis and Tardos [118] extended the results to incomplete-informations settings and gave constant bounds for both auctions. Leme et al. [90], Syrgkanis and Tardos [118] proposed a bluffing deviation, where a player pretends to play as in equilibrium, until the right moment when the player deviates to acquire some item. This hypothetical deviation gives rise to useful inequalities to bound the PoA.

In this section, we show a PoA bound of 2. To our knowledge, this is the best-known PoA guarantee among all auctions of different formats for the sponsored search problem. In the analysis, the dual variables are intuitively constructed such that they correspond to the player utilities and player payments. In order to show the feasibility of dual variables, we also use the idea of bluffing deviations. These deviations, coupling with the assumption of equilibrium, lead to useful inequalities which are served to prove the feasibility. The primal-dual approach indeed enables the improvement as well as a fairly simple proof.

Formulation. For player i with valuation v_i and quality factor α_i , let $x_{ij}(v_i, \alpha_i)$ be a variable indicating the interim assignment of slot j to player i . Recall that F_i is the distribution of (v_i, α_i) . Consider the following relaxation of the sponsored search problem. In the primal relaxation, the first constraint says that a player receives at most one slot and the second one ensures that one slot is assigned to at most one player.

$$\begin{aligned} \max \sum_{i,j} \mathbb{E}_{(v_i, \alpha_i) \sim F_i} & \left[\beta_j \alpha_i v_i \cdot x_{ij}(v_i, \alpha_i) \right] \\ & \sum_j x_{ij}(v_i, \alpha_i) \leq 1 && \forall i, v_i, \alpha_i \\ \sum_i \sum_{(v_i, \alpha_i)} f_i(v_i, \alpha_i) x_{ij}(v_i, \alpha_i) & \leq 1 && \forall j \\ & x_{ij}(v_i, \alpha_i) \geq 0 && \forall i, j, v_i, \alpha_i \end{aligned}$$

The dual is the following.

$$\begin{aligned} \min \sum_i \sum_{(v_i, \alpha_i)} y_i(v_i, \alpha_i) + \sum_j z_j \\ y_i(v_i, \alpha_i) + f_i(v_i, \alpha_i) z_j & \geq f_i(v_i, \alpha_i) \cdot \beta_j \alpha_i v_i && \forall i, j, v_i, \alpha_i \\ y_i(v_i, \alpha_i), z_j & \geq 0 && \forall i, j, v_i, \alpha_i \end{aligned}$$

Theorem 3.9 *For every sequential second-price auction setting, the expected welfare of every perfect Bayesian equilibrium is at least half the maximum welfare.*

Proof Fix a Bayes-Nash equilibrium σ . Let $\pi(\sigma(v, \alpha), i)$ be the random variable indicating the slot that player i receives in the equilibrium $\sigma(v, \alpha)$ given the valuation profile v and the quality factor profile α . Whenever σ and (v, α) are clear in the context, we simply write $\pi(\sigma(v, \alpha), i)$ as $\pi(i)$. Inversely, let $\pi^{-1}(\sigma(v, \alpha), j)$ be the winner of slot j in profile $\sigma(v, \alpha)$. Note that $\pi^{-1}(\sigma(v, \alpha), j)$ is also a random variable.

Dual Variables. For fixed (v_i, α_i) , denote $B(v_i, \alpha_i)$ the distribution of the equilibrium bid $\mathbf{b} = \sigma((v_i, v_{-i}), (\alpha_i, \alpha_{-i}))$. Recall that $\mathbf{b} = (b_1, \dots, b_n)$ where b_i is a bid vector over bids b_{ij} — the equilibrium bid that player i submits in the round selling slot j . Moreover, denote B_{-i} the distribution of the equilibrium bid $\mathbf{b}_{-i} =$

$\sigma_{-i}((v_i, \mathbf{v}_{-i}), (\alpha_i, \boldsymbol{\alpha}_{-i})) = \sigma_{-i}(\mathbf{v}_{-i}, \boldsymbol{\alpha}_{-i})$ where the last equality is due to the independence of distributions F_i 's. Define the dual variables as follows.

$$y_i(v_i, \alpha_i) := f_i(v_i, \alpha_i) \cdot \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[\beta_{\pi(\mathbf{b}, i)} \cdot \alpha_i v_i \right],$$

$$z_j := \max_i \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}} \left[\beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b}_{-i}, j)} b_{\pi^{-1}(\mathbf{b}_{-i}, j)} \right]$$

Note that $\pi^{-1}(\mathbf{b}_{-i}, j)$ is the winner of slot j in the round selling slot j assuming that player i does not participate to this round.

Feasibility. Fix a player i with valuation v_i and quality factor α_i , and a slot j . We show that the dual constraint corresponding to i, j, v_i, α_i is satisfied. By the dual variable definitions and the independence of distributions, it is equivalent to prove that:

$$\mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[\beta_{\pi(\mathbf{b}, i)} \cdot \alpha_i v_i + \beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b}_{-i}, j)} b_{\pi^{-1}(\mathbf{b}_{-i}, j)} \right] \geq \beta_j \cdot \alpha_i v_i \quad (3.4)$$

We prove this inequality through a choice of a hypothetical deviation of player i and use the assumption that σ is a Bayes-Nash equilibrium. We first make some observations. Consider a fixed valuation profile \mathbf{v}_{-i} , a fixed quality factor profile $\boldsymbol{\alpha}_{-i}$ and a realization of (mixed) equilibrium $\sigma((v_i, \mathbf{v}_i), (\alpha_i, \boldsymbol{\alpha}_{-i}))$, denoted as $\mathbf{b} = (b_1, \dots, b_n)$. Now the assignment π of slots to players is completely determined. There are three different cases.

Case 1: Player i receives some slot $\pi(i) \leq j$. Then $\beta_{\pi(i)} \cdot \alpha_i v_i \geq \beta_j \cdot \alpha_i v_i$ since $\beta_{\pi(i)} \geq \beta_j$.

Case 2: $\pi(i) > j$ and $\alpha_{\pi^{-1}(j)} b_{\pi^{-1}(j), j} \geq \alpha_i v_i$. Then $\beta_j \cdot \alpha_{\pi^{-1}(j)} \cdot b_{\pi^{-1}(j)} \geq \beta_j \cdot \alpha_i v_i$.

Case 3: $\pi(i) > j$ and $\alpha_{\pi^{-1}(j)} b_{\pi^{-1}(j), j} < \alpha_i v_i$. Note that in this case in the round j , player i could have submitted a bid without violating the no-overbidding property such that the corresponding effective bid is infinitesimal larger than $\alpha_{\pi^{-1}(j)} b_{\pi^{-1}(j), j}$ and could have received slot j .

We are now choosing a bid deviation in order to prove the dual constraint based on the fact that σ is a Bayes-Nash equilibrium. Intuitively, the different cases above suggest the following deviation. For the first two cases, the term inside the expectations in the left-hand-side of (3.4) is already larger than the right-hand-side (so no need to deviate). Hence, the deviation is necessary only in Case 3.

Formally, we define the (mixed) deviation b'_i as follows. First, player i follows the equilibrium strategy b_i . If until the allocation step of slot j , player i has not received any slot then he submits v_i . As σ is a Bayes-Nash equilibrium, the utility of player i is at least that induced by this deviation. Specifically,

$$\mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} [u_i(\mathbf{b})] \geq \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}} \mathbb{E}_{b'_i} \left[u_i(b'_i, \mathbf{b}_{-i}) \right],$$

where, for short, we write $u_i(\mathbf{b}) = u_i(\mathbf{b}; v_i, \alpha_i)$ since (v_i, α_i) is fixed.

By definition of the deviation b'_i , player i follows the same equilibrium strategy b_i if Case 1 happens. Therefore, the above inequality is equivalent to

$$\begin{aligned} & \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[u_i(\mathbf{b}) \mid \text{Case 2 or Case 3} \right] \\ & \geq \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}} \mathbb{E}_{b'_i} \left[u_i(b'_i, \mathbf{b}_{-i}) \mid \text{Case 2 or Case 3} \right] \end{aligned} \quad (3.5)$$

Note that if Case 3 holds then player i gets slot j with the payment

$$\frac{\alpha_{\pi^{-1}(\mathbf{b}_{-i}, j)} b_{\pi^{-1}(\mathbf{b}_{-i}, j), j}}{\alpha_i}.$$

So

$$\begin{aligned} & \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}} \mathbb{E}_{b'_i} \left[u_i(b'_i, \mathbf{b}_{-i}) \mid \text{Case 3} \right] \\ & = \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}} \left[\beta_j \cdot \alpha_i \left(v_i - \frac{\alpha_{\pi^{-1}(\mathbf{b}_{-i}, j)} b_{\pi^{-1}(\mathbf{b}_{-i}, j), j}}{\alpha_i} \right) \mid \text{Case 3} \right] \end{aligned} \quad (3.6)$$

We are now ready to prove the inequality (3.4). We have

$$\begin{aligned} & \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[\beta_{\pi(\mathbf{b}, i)} \cdot \alpha_i v_i + \beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b}_{-i}, j)} b_{\pi^{-1}(\mathbf{b}_{-i}, j), j} \right] \\ & = \sum_{\ell=1,2,3} \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[\beta_{\pi(\mathbf{b}, i)} \cdot \alpha_i v_i + \beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b}_{-i}, j)} b_{\pi^{-1}(\mathbf{b}_{-i}, j), j} \mid \text{Case } \ell \right] \\ & \geq \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[\beta_j \cdot \alpha_i v_i \mid \text{Case 1} \right] \\ & \quad + \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[\beta_{\pi(\mathbf{b}, i)} \cdot \alpha_i v_i + \beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b}_{-i}, j)} b_{\pi^{-1}(\mathbf{b}_{-i}, j), j} \mid \text{Case 2 or 3} \right] \\ & \geq \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[\beta_j \cdot \alpha_i v_i \mid \text{Case 1} \right] \\ & \quad + \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[u_i(\mathbf{b}) + \beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b}_{-i}, j)} b_{\pi^{-1}(\mathbf{b}_{-i}, j), j} \mid \text{Case 2 or 3} \right] \\ & \geq \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[\beta_j \cdot \alpha_i v_i \mid \text{Case 1} \right] \\ & \quad + \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[u_i((b'_i, \mathbf{b}_{-i})) + \beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b}_{-i}, j)} b_{\pi^{-1}(\mathbf{b}_{-i}, j), j} \mid \text{Case 2 or 3} \right] \\ & \geq \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[\beta_j \cdot \alpha_i v_i \mid \text{Case 1 or 2} \right] \\ & \quad + \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[u_i((b'_i, \mathbf{b}_{-i})) + \beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b}_{-i}, j)} b_{\pi^{-1}(\mathbf{b}_{-i}, j), j} \mid \text{Case 3} \right] \\ & \geq \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[\beta_j \cdot \alpha_i v_i \mid \text{Case 1 or 2} \right] \\ & \quad + \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[\beta_j \cdot \alpha_i v_i - \beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b}_{-i}, j)} b_{\pi^{-1}(\mathbf{b}_{-i}, j), j} + \beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b}_{-i}, j)} b_{\pi^{-1}(\mathbf{b}_{-i}, j), j} \mid \text{Case 3} \right] \\ & = \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} \left[\beta_j \cdot \alpha_i v_i \right] = \beta_j \cdot \alpha_i v_i \end{aligned}$$

The first inequality uses the assumption of Case 1: $\beta_{\pi(i)} \geq \beta_j$. The second inequality holds since the utility $u_i(\mathbf{b}) \leq \beta_{\pi(\mathbf{b}, i)} \cdot \alpha_i v_i$. The third inequality is due to (3.5). The fourth inequality follows the assumption of Case 2: $\alpha_{\pi^{-1}(j)} v_{\pi^{-1}(j)} \geq \alpha_i v_i$. The last inequality follows (3.6). Hence, the constructed dual variables form a dual feasible solution.

Bounding primal and dual. Let \mathbf{B} be the distribution of the equilibrium bid $\mathbf{b} = \sigma(\mathbf{v})$. The expected welfare of equilibrium σ is $\mathbb{E}_{\mathbf{b} \sim \mathbf{B}} [\sum_i \beta_{\pi(\mathbf{b},i)} \alpha_i v_i]$. By the definition of dual variables, we have

$$\sum_{i, (v_i, \alpha_i)} y_i(v_i, \alpha_i) = \sum_i \mathbb{E}_{(v_i, \alpha_i) \sim F_i} \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i, \alpha_i)} [\beta_{\pi(\mathbf{b},i)} \alpha_i v_i] = \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} \left[\sum_i \beta_{\pi(\mathbf{b},i)} \alpha_i v_i \right].$$

Besides, consider a slot j and let i^* be the player such that

$$z_j = \mathbb{E}_{\mathbf{b}_{-i^*} \sim \mathbf{B}_{-i^*}} \left[\beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b}_{-i^*,j})} b_{\pi^{-1}(\mathbf{b}_{-i^*,j}),j} \right].$$

As the right-hand side is independent of $b_{i^*,j}$, we have

$$z_j = \mathbb{E}_{b_{i^*}} \mathbb{E}_{\mathbf{b}_{-i^*} \sim \mathbf{B}_{-i^*}} \left[\beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b}_{-i^*,j})} b_{\pi^{-1}(\mathbf{b}_{-i^*,j}),j} \right] = \mathbb{E}_{\mathbf{b}} \left[\beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b}_{-i^*,j})} b_{\pi^{-1}(\mathbf{b}_{-i^*,j}),j} \right].$$

Moreover,

$$z_j \leq \mathbb{E}_{\mathbf{b}} \left[\beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b},j)} b_{\pi^{-1}(\mathbf{b},j),j} \right] \leq \mathbb{E}_{\mathbf{b}} \left[\beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b},j)} v_{\pi^{-1}(\mathbf{b},j)} \right].$$

The first inequality holds since the effective bid of the slot- j -winner in round j including all players is larger than that in case player i^* does not participate. The last inequality is due to the non-overbidding property. Summing over all j , we have

$$\sum_j z_j \leq \mathbb{E}_{\mathbf{b}} \left[\sum_j \beta_j \cdot \alpha_{\pi^{-1}(\mathbf{b},j)} v_{\pi^{-1}(\mathbf{b},j)} \right] = \mathbb{E}_{\mathbf{b}} \left[\sum_i \beta_{\pi(\mathbf{b},i)} \alpha_i v_i \right].$$

Thus, the dual objective value is at most twice the expected welfare of the equilibrium. \square

Remark. The non-overbidding assumption can be relaxed in the same way as the remark in Section 3.4.1. Specifically, if players are allowed to bid up to a constant r times their valuations (hence, the utility of a winning player may be negative) then the PoA is at most $(1+r)$.

3.5.2 Sequential First Price Auctions in Matching Markets

Model. In the matching market problem, there are n players and m items. Each player i has a *private unit-demand valuation* $v_i : 2^{[m]} \rightarrow \mathbb{R}$ defined as $v_{iS} := \max_j v_{ij}$ where v_{ij} is the valuation of player i on item j . Note that in the sponsored search problem $v_{ij} \geq v_{ij'}$ for slots $j < j'$ and for every player i , while in the matching market problem it might be that for some items j, j' and some players i, i' , $v_{ij} > v_{ij'}$ and $v_{i'j} < v_{i'j'}$. The valuation vector v_i is drawn from a publicly known distribution F_i . In the model, we assume that the distributions F_i 's are mutually independent. An *outcome* is an assignment of items to players.

In the auction, the auctioneer sells items sequentially one-by-one via the first price mechanisms. At the consideration of item j , the auctioneer collects all the bids b_{ij} on item j from all players. We also assume that the non-overbidding property, meaning that $b_{ij} \leq v_i$ for all i and j . The auctioneer then assigns item j to the player with highest bid. Note that, in contrast to the sponsored search problem, a player may receive multiple items. The payment of the winning player is simply the winning bid. The *utility* of player i is $(v_{iS} - \sum_{j \in S} b_{ij})$ where S is its allocated items. The

social welfare of the outcome is $\sum_{i,j} v_{iS}$ where the sum is taken over all players i and their corresponding allocated items S .

Related work about sequential auctions have been summarized in the previous section. For the matching market problem, Leme et al. [90] proved that the sequential auctions via the second price mechanisms may lead to unbounded inefficiency. The authors [90] then considered the sequential first price auctions and showed that in full-information settings, the PoA is at most 2 and 4 for pure and mixed Nash equilibria. Subsequently, Syrgkanis and Tardos [118] extended the results to incomplete-information settings. They proved a Bayesian PoA bound $2e/(e-1)$ for matching markets with independent valuations. They also raised a question whether the difference of PoA bounds between the full-information settings and the incomplete-information ones is necessary.

In this section, we answer this question by showing that the (mixed) Bayesian PoA is at most 2. In the proof, we use similar bluffing deviations as in [90, 118] and the primal-dual approach enables the improvement. The proof follows similar structure as the one in Section 3.5.1; however, there is a subtle difference compared to the sponsored search problem. In the latter, each player receives at most one item (slot) so in constructing the hypothetical deviation, it is sufficient to design a deviation in which the player gets one item, improves its utility and then leaves the game (bids 0 in subsequent rounds). In the matching market problem, a player may receive multiple items hence the player would deviate in such a way that he receives only the highest valuable item without receiving (so paying for) items allocated in previous rounds. However, such deviations may lead to completely different outcomes and the equilibrium structure could be very complex to analyze. Therefore, we do not reason directly on the utility of players in deviation. Instead, we explore the connection between the winning bid and the player valuation. Consequently, the argument works only for the sequential auctions via the first price mechanisms (but not via the second price mechanisms).

Formulation. For every player i , every valuation v_i and every set of items S , let $x_{iS}(v_i)$ be a variable indicating the interim assignment of S to player i . Consider the following formulation and its dual. In the primal, the first and second constraints are relaxations of the facts that a player receives a set of items and that an item is assigned to at most one player, respectively.

$$\begin{aligned}
\max \quad & \sum_{i,S} \mathbb{E}_{v_i \sim F_i} [v_{iS} \cdot x_{iS}(v_i)] & \min \quad & \sum_i \sum_{v_i} y_i(v_i) + \sum_j z_j \\
& \sum_S x_{iS}(v_i) \leq 1 \quad \forall i, v_i & & y_i(v_i) + f_i(v_i) \sum_{j \in S} z_j \geq f_i(v_i) \cdot v_{iS} \\
& \sum_i \sum_{v_i} f_i(v_i) \sum_{S: j \in S} x_{iS}(v_i) \leq 1 \quad \forall j & & \\
& x_{iS}(v_i) \geq 0 \quad \forall i, j, v_i & & y_i(v_i), z_j \geq 0 \quad \forall i, j, v_i
\end{aligned}$$

Theorem 3.10 *For every sequential first-price auction, the expected welfare of every perfect Bayesian equilibrium is at least half the maximum welfare.*

Proof Fix a Bayes-Nash equilibrium σ . Let $\pi(\sigma(v), i)$ be the random variable indicating the set of items allocated to player i in the equilibrium given the valuation profile v . Inversely, let $\pi^{-1}(\sigma(v), j)$ be the winner of item j . Note that $\pi^{-1}(\sigma(v), j)$ is also a random variable.

Dual Variables. For a fixed valuation v_i , denote $\mathbf{B}(v_i)$ the distribution of the equilibrium bid $\mathbf{b} = \boldsymbol{\sigma}(v_i, \mathbf{v}_{-i})$. Recall that $\mathbf{b} = (b_1, \dots, b_n)$ where b_i is a bid vector over b_{ij} — the equilibrium bid that player i submits in the round selling item j for $1 \leq j \leq m$. Moreover, denote \mathbf{B}_{-i} the distribution of the equilibrium bid $\mathbf{b}_{-i} = \boldsymbol{\sigma}_{-i}(v_i, \mathbf{v}_{-i}) = \boldsymbol{\sigma}_{-i}(\mathbf{v}_{-i})$ where the last equality is due to the independence of distributions. Define the dual variables as follows.

$$\begin{aligned} y_i(v_i) &:= f_i(v_i) \cdot \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i)} \left[v_{i, \pi(\mathbf{b}, i)} \right], \\ z_j &:= \max_i \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}} \left[b_{\pi^{-1}(\mathbf{b}_{-i}, j), j} \right]. \end{aligned}$$

Note that $\pi^{-1}(\mathbf{b}_{-i}, j)$ is the winner of item j assuming that player i does not participate to this round.

Feasibility. Fix a player i with valuation v_i and a set of items S . We show that the dual constraint corresponding to i, S, v_i is satisfied. By the dual variable definitions and the independence of distributions, it is equivalent to prove that:

$$\mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i)} \left[v_{i, \pi(\mathbf{b}, i)} + \sum_{j \in S} b_{\pi^{-1}(\mathbf{b}_{-i}, j), j} \right] \geq v_i S. \quad (3.7)$$

We prove this inequality through a choice of a hypothetical deviation of player i and use the assumption that $\boldsymbol{\sigma}$ is a Nash-Bayes equilibrium. For any set of items U , let $j^*(U) \in U$ be an item such that $v_{j^*} = \max_{j \in U} v_{ij} = v_{iU}$. We first make some observations. Consider a fixed valuation profile \mathbf{v}_{-i} and a realization of (mixed) equilibrium $\boldsymbol{\sigma}(v_i, \mathbf{v}_{-i})$, denoted as $\mathbf{b} = (b_1, \dots, b_n)$. Now the assignment π of items to players is completely determined. Let $T = \pi(\mathbf{b}, i)$. There are three different cases.

Case 1: $v_{i, j^*(T)} \geq v_{i, j^*(S)}$.

Case 2: $v_{i, j^*(T)} < v_{i, j^*(S)}$ (so $j^*(S) \notin T$) and the round of $j^*(S)$ is before the round of $j^*(T)$. In this case, $b_{\pi^{-1}(j^*(S)), j^*(S)} \geq v_{i, j^*(S)} - v_{i, j^*(T)}$ since otherwise i could have improved its utility by submitting a bid of value $(v_{i, j^*(S)} - v_{i, j^*(T)})$ and stop playing the remaining rounds (i.e., submitting bids 0).

Case 3: $v_{i, j^*(T)} < v_{i, j^*(S)}$ (so $j^*(S) \notin T$) and the round of $j^*(T)$ is before the round of $j^*(S)$. Again, in this case, $b_{\pi^{-1}(j^*(S)), j^*(S)} \geq v_{i, j^*(S)} - v_{i, j^*(T)}$ by the same argument.

The cases suggest the following (mixed) deviation b'_i for player i . Player i draws a random sample of a valuation profile $\mathbf{w}_{-i} \in \mathbf{F}_{-i}$ and determines the winning set $T = \pi(\boldsymbol{\sigma}(v_i, \mathbf{w}_{-i}), i)$ and also item $j^*(T)$. If $v_{i, j^*(T)} \geq v_{i, j^*(S)}$ then player i follows the equilibrium strategy b_i . Otherwise, player i first follows strategy b_i until the round of item $j^*(S)$. In the round of $j^*(S)$, bid $b'_{i, j^*(S)} = v_{i, j^*(S)} - v_{i, j^*(T)}$ and in the subsequent rounds, bid 0.

As $\boldsymbol{\sigma}$ is a Bayes-Nash equilibrium, the utility of player i is at least that induced by this deviation. Specifically,

$$\mathbb{E}_{\mathbf{b} \sim \mathbf{B}(v_i)} [u_i(\mathbf{b})] = \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{B}_{-i}} \mathbb{E}_{\boldsymbol{\sigma}} \left[u_i(b_i, \boldsymbol{\sigma}_{-i}(\mathbf{v}_{-i})) \right] \geq \mathbb{E}_{\mathbf{w}_{-i} \sim \mathbf{B}_{-i}} \mathbb{E}_{\boldsymbol{\sigma}} \left[u_i(b'_i, \boldsymbol{\sigma}_{-i}(\mathbf{w}_{-i})) \right]$$

where since v_i is fixed, for short, we write $u_i(\mathbf{b}) = u_i(\mathbf{b}; v_i)$. By definition of the deviation b'_i , player i follows the same equilibrium strategy b_i if Case 1 happens.

Therefore, the above inequality implies

$$\begin{aligned} & \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{B}_{-i}} \mathbb{E}_{\boldsymbol{\sigma}} \left[b_{\pi^{-1}(\boldsymbol{\sigma}_{-i}(\mathbf{v}_{-i}), j^*(S)), j^*(S)} \mid \text{Case 2 or 3} \right] \\ & \geq \mathbb{E}_{\mathbf{v}_{-i} \sim \mathbf{B}_{-i}} \mathbb{E}_{\boldsymbol{\sigma}} \left[v_{i, j^*(S)} - v_{i, j^*(T)} \mid \text{Case 2 or 3} \right], \end{aligned} \quad (3.8)$$

where $T = \pi(\boldsymbol{\sigma}_{-i}(\mathbf{v}_i, \mathbf{v}_{-i}), i)$ the set of items allocated to i .

We are now ready to prove the inequality (3.7). We have

$$\begin{aligned} & \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(\mathbf{v}_i)} \left[v_{i, \pi(\mathbf{b}, i)} + \sum_{j \in S} b_{\pi^{-1}(\mathbf{b}_{-i}, j), j} \right] \\ & \geq \sum_{\ell=1,2,3} \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(\mathbf{v}_i)} \left[v_{i, \pi(\mathbf{b}, i)} + b_{\pi^{-1}(\mathbf{b}_{-i}, j^*(S)), j^*(S)} \mid \text{Case } \ell \right] \\ & \geq \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(\mathbf{v}_i)} [v_{iS} \mid \text{Case 1}] + \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(\mathbf{v}_i)} \left[v_{i, \pi(\mathbf{b}, i)} + b_{\pi^{-1}(\mathbf{b}_{-i}, j^*(S)), j^*(S)} \mid \text{Case 2 or 3} \right] \\ & \geq \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(\mathbf{v}_i)} [v_{iS} \mid \text{Case 1}] + \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(\mathbf{v}_i)} \left[v_{i, \pi(\mathbf{b}, i)} + (v_{i, j^*(S)} - v_{i, j^*(T)}) \mid \text{Case 2 or 3} \right] \\ & = v_{iS} \end{aligned}$$

The first inequality holds since $j^*(S) \in S$ and the bids are non-negative. The second inequality holds due to the assumption of Case 1. The third inequality follows Inequality (3.8). Hence, the constructed dual variables form a dual feasible solution.

Bounding primal and dual. By the definition of dual variables, we have

$$\sum_{i, v_i} y_i(v_i) = \sum_i \mathbb{E}_{v_i \sim F_i} \mathbb{E}_{\mathbf{b} \sim \mathbf{B}(\mathbf{v}_i)} \left[v_{i, \pi(\mathbf{b}, i)} \right] = \mathbb{E}_{\mathbf{b}} \left[\sum_i v_{i, \pi(\mathbf{b}, i)} \right].$$

Besides, consider an item j and let i^* be a player such that

$$z_j = \mathbb{E}_{\mathbf{b}_{-i^*} \sim \mathbf{B}_{-i^*}} \left[b_{\pi^{-1}(\mathbf{b}_{-i^*}, j), j} \right].$$

As the right-hand side is independent of b_{i^*} , we have

$$z_j = \mathbb{E}_{b_{i^*}} \mathbb{E}_{\mathbf{b}_{-i^*} \sim \mathbf{B}_{-i^*}} \left[b_{\pi^{-1}(\mathbf{b}_{-i^*}, j), j} \right] = \mathbb{E}_{\mathbf{b}} \left[b_{\pi^{-1}(\mathbf{b}_{-i^*}, j), j} \right] \leq \mathbb{E}_{\mathbf{b}} \left[b_{\pi^{-1}(\mathbf{b}, j), j} \right].$$

Summing over all items j , we get

$$\sum_j z_j \leq \mathbb{E}_{\mathbf{b}} \left[\sum_j b_{\pi^{-1}(\mathbf{b}, j), j} \right] = \mathbb{E}_{\mathbf{b}} \left[\sum_i \sum_{j \in \pi(\mathbf{b}, i)} b_{ij} \right] \leq \mathbb{E}_{\mathbf{b}} \left[\sum_i v_{i, \pi(\mathbf{b}, i)} \right],$$

where the last inequality is due to non-overbidding property. Thus, the dual objective value is at most twice the expected welfare of the equilibrium. \square

Chapter 4

Efficient Online Learning Algorithms and Auction Design

In this chapter, we present a general efficient online learning algorithm and characterize its regret bound based on the concavity parameters. We give an online algorithm which at any step has access to the gradient of the black-box functions (Section 4.1.1) and derive an algorithm without this assumption (Section 4.1.2). Subsequently, using the general algorithm, we provide the performance of fictitious plays in Section 4.2 and revenue maximization in multi-dimensional environments in Section 4.3.

4.1 Framework of Online Learning

We say that a function $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}$ is α_Φ -strongly convex w.r.t $\|\cdot\|$ if

$$\Phi(\mathbf{x}') \geq \Phi(\mathbf{x}) + \langle \nabla \Phi(\mathbf{x}), \mathbf{x}' - \mathbf{x} \rangle + \frac{\alpha_\Phi}{2} \|\mathbf{x}' - \mathbf{x}\|^2.$$

Given a strictly convex function $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}$, define the *Bregman divergence*

$$D_\Phi(\mathbf{x} \|\mathbf{x}') := \Phi(\mathbf{x}) - \Phi(\mathbf{x}') - \langle \nabla \Phi(\mathbf{x}'), \mathbf{x} - \mathbf{x}' \rangle$$

The following lemma generalizes the Pythagorean theorem (proof can be found in [14] for example).

Lemma 4.1 (Generalized Pythagorean Property) *Given a convex body $\mathcal{K} \subset \mathbb{R}^n$. Let $\mathbf{x} \in \mathcal{K}$ and $\mathbf{y}' \in \mathbb{R}^n$. Let \mathbf{y} be the projection of \mathbf{y}' on \mathcal{K} , defined as $\mathbf{y} = \arg \min_{\bar{\mathbf{y}} \in \mathcal{K}} D_\Phi(\bar{\mathbf{y}} \|\mathbf{y}')$. Then $D_\Phi(\mathbf{x} \|\mathbf{y}) \leq D_\Phi(\mathbf{x} \|\mathbf{y}')$.*

4.1.1 Regret of (λ, μ) -Concave Functions

In this section, we assume that at every time step, the online algorithm has access to the gradient of the functions. The algorithm is the standard mirror descent.

Mirror descent. Let Φ be a α_Φ -strongly convex function w.r.t $\|\cdot\|$. At time step t , define $\nabla^t = -\nabla F^t(\mathbf{x}^t)$ and denote $\boldsymbol{\theta}^t = \nabla \Phi(\mathbf{x}^t)$. The algorithm selects the decision \mathbf{x}^{t+1} as follows.

$$\begin{aligned} \boldsymbol{\vartheta}^{t+1} &= \boldsymbol{\theta}^t - \eta \cdot \nabla^t \\ \mathbf{y}^{t+1} &= \nabla \Phi^*(\boldsymbol{\vartheta}^{t+1}) \\ \mathbf{x}^{t+1} &= \arg \min_{\mathbf{x} \in \mathcal{K}} D_\Phi(\mathbf{x} \|\mathbf{y}^{t+1}) \end{aligned}$$

An equivalent description is

$$\mathbf{x}^{t+1} = \arg \max_{\mathbf{x} \in \mathcal{K}} \{ \langle \eta \nabla F^t(\mathbf{x}^t), \mathbf{x} - \mathbf{x}^t \rangle - D_{\Phi}(\mathbf{x} \parallel \mathbf{x}^t) \}. \quad (4.1)$$

The regret bound of the mirror descent follows the standard technique. Here we give the proof based on the potential argument of Bansal and Gupta [14] and derive the bound based on the concavity parameters.

Theorem 4.1 *If F^t is (λ, μ) -concave for every $1 \leq t \leq T$, then the mirror descent algorithm achieves $(\frac{\lambda}{\mu}, R(T))$ -regret where*

$$R(T) = \frac{1}{\mu \cdot \eta} D_{\Phi}(\mathbf{x}^* \parallel \mathbf{x}^1) + \frac{\eta}{\mu \cdot 2\alpha_{\Phi}} \sum_{t=1}^T \|\nabla^t\|_*^2.$$

If $\|\nabla^t\|_* \leq L$ for $1 \leq t \leq T$ (i.e., F^t is L -Lipschitz w.r.t $\|\cdot\|$) and $D_{\Phi}(\mathbf{x}^* \parallel \mathbf{x}^1)$ is bounded by G^2 then by choosing $\eta = \frac{G}{L} \sqrt{\frac{2\alpha_{\Phi}}{T}}$,

$$R(T) \leq \frac{GL}{\mu} \sqrt{2\alpha_{\Phi} T}.$$

Proof Define the potential as

$$\Psi^t = \frac{1}{\eta} D_{\Phi}(\mathbf{x}^* \parallel \mathbf{x}^t).$$

It is proved in [14] that:

$$\begin{aligned} D_{\Phi}(\mathbf{x}^* \parallel \mathbf{x}^{t+1}) - D_{\Phi}(\mathbf{x}^* \parallel \mathbf{x}^t) &\leq \eta \langle \nabla^t, \mathbf{x}^* - \mathbf{x}^t \rangle + \frac{\eta^2}{2\alpha_{\Phi}} \|\nabla^t\|_*^2 \\ &= -\eta \langle \nabla F^t, \mathbf{x}^* - \mathbf{x}^t \rangle + \frac{\eta^2}{2\alpha_{\Phi}} \|\nabla^t\|_*^2. \end{aligned} \quad (4.2)$$

For completeness we give the proof of this inequality. We have

$$\begin{aligned} \eta (\Psi^{t+1} - \Psi^t) &= D_{\Phi}(\mathbf{x}^* \parallel \mathbf{x}^{t+1}) - D_{\Phi}(\mathbf{x}^* \parallel \mathbf{x}^t) \\ &\leq D_{\Phi}(\mathbf{x}^* \parallel \mathbf{y}^{t+1}) - D_{\Phi}(\mathbf{x}^* \parallel \mathbf{x}^t) \\ &= \Phi(\mathbf{x}^*) - \Phi(\mathbf{y}^{t+1}) - \underbrace{\langle \nabla \Phi(\mathbf{y}^{t+1}), \mathbf{x}^* - \mathbf{y}^{t+1} \rangle}_{\vartheta^{t+1}} - \Phi(\mathbf{x}^*) + \Phi(\mathbf{x}^t) + \underbrace{\langle \nabla \Phi(\mathbf{x}^t), \mathbf{x}^* - \mathbf{x}^t \rangle}_{\theta^t} \\ &= \Phi(\mathbf{x}^t) - \Phi(\mathbf{y}^{t+1}) - \langle \vartheta^{t+1}, \mathbf{x}^t - \mathbf{y}^{t+1} \rangle - \langle \vartheta^{t+1} - \theta^t, \mathbf{x}^* - \mathbf{x}^t \rangle \\ &= \Phi(\mathbf{x}^t) - \Phi(\mathbf{y}^{t+1}) - \langle \theta^t, \mathbf{x}^t - \mathbf{y}^{t+1} \rangle + \langle \eta \nabla^t, \mathbf{x}^t - \mathbf{y}^{t+1} \rangle + \langle \eta \nabla^t, \mathbf{x}^* - \mathbf{x}^t \rangle \\ &\leq -\frac{\alpha_{\Phi}}{2} \|\mathbf{y}^{t+1} - \mathbf{x}^t\|^2 + \eta \langle \nabla^t, \mathbf{x}^t - \mathbf{y}^{t+1} \rangle + \eta \langle \nabla^t, \mathbf{x}^* - \mathbf{x}^t \rangle \\ &= -\frac{\alpha_{\Phi}}{2} \|\mathbf{y}^{t+1} - \mathbf{x}^t\|^2 + \frac{1}{\alpha_{\Phi}} \langle \eta \nabla^t, \alpha_{\Phi} (\mathbf{x}^t - \mathbf{y}^{t+1}) \rangle + \eta \langle \nabla^t, \mathbf{x}^* - \mathbf{x}^t \rangle \\ &\leq \frac{\eta^2}{2\alpha_{\Phi}} \|\nabla^t\|_*^2 + \eta \langle \nabla^t, \mathbf{x}^* - \mathbf{x}^t \rangle \end{aligned}$$

where the first inequality is due to the generalized Pythagorean property (Lemma 4.1); the fourth equality follows the update rule $\vartheta^{t+1} = \theta^t - \eta \cdot \nabla^t$; the second inequality holds since Φ is α_{Φ} -strongly convex; and in the last inequality, we use Cauchy-Schwarz inequality $\langle \mathbf{a}, \mathbf{b} \rangle \leq \|\mathbf{b}\| \|\mathbf{a}\|_* \leq \|\mathbf{b}\|^2/2 + \|\mathbf{a}\|_*^2/2$.

Using the bound of the potential change due to Inequality (4.2), we get

$$\begin{aligned}
\sum_{t=1}^T (\lambda F^t(\mathbf{x}^*) - \mu F^t(\mathbf{x}^t)) &\leq \Psi_1 + \sum_{t=1}^T \left[\lambda F^t(\mathbf{x}^*) - \mu F^t(\mathbf{x}^t) + \Psi^{t+1} - \Psi^t \right] \\
&\leq \Psi_1 + \sum_{t=1}^T \left[\underbrace{\lambda F^t(\mathbf{x}^*) - \mu F^t(\mathbf{x}^t) - \langle \nabla F^t(\mathbf{x}^t), \mathbf{x}^* - \mathbf{x}^t \rangle}_{\leq 0 \text{ since } F^t \text{ is } (\lambda, \mu)\text{-concave}} + \frac{\eta}{2\alpha_\Phi} \|\nabla_t\|_*^2 \right] \\
&\leq \frac{1}{\eta} D_\Phi(\mathbf{x}^* \|\mathbf{x}^1) + \frac{\eta}{2\alpha_\Phi} \sum_{t=1}^T \|\nabla_t\|_*^2 \tag{4.3}
\end{aligned}$$

If the norms $\|\nabla^t\|_*$ are bounded by L and $D_\Phi(\mathbf{x}^* \|\mathbf{x}^1)$ is bounded by G^2 then

$$\sum_{t=1}^T F^t(\mathbf{x}^t) \geq \frac{\lambda}{\mu} \sum_{t=1}^T F^t(\mathbf{x}^*) - \frac{1}{\mu \cdot \eta} G^2 - \frac{\eta}{\mu \cdot 2\alpha_\Phi} T L^2$$

Choose $\eta = \frac{G}{L} \sqrt{\frac{2\alpha_\Phi}{T}}$, we deduce that the algorithm is $(\frac{\lambda}{\mu}, R(T))$ -regret where $R(T) = O(\frac{GL}{\mu} \sqrt{2\alpha_\Phi T})$. \square

4.1.2 Multilinear Extension of Discretization

In this section, we consider the domain $\mathcal{K} = [0, 1]^n$. In order to apply the mirror descent algorithm, it is crucial to compute the gradient of the black-box function whereas we can only query its values. In order to bypass this issue, we consider a discretization of $[0, 1]^n$ and the multilinear extensions of the functions on these discrete points.

Let f be a function $f : [0, 1]^n \rightarrow \mathbb{R}$. Consider a lattice $\mathcal{L} = \{0, 2^{-M}, 2 \cdot 2^{-M}, \dots, \ell \cdot 2^{-M}, \dots, 1\}^n$ where $0 \leq \ell \leq 2^M$ for some large parameter M as a discretization of $[0, 1]^n$. Note that each $x_i \in \{0, 2^{-M}, 2 \cdot 2^{-M}, \dots, \ell \cdot 2^{-M}, \dots, 1\}$ can be uniquely decomposed as $x_i = \sum_{j=0}^M 2^{-j} y_{ij}$ where $y_{ij} \in \{0, 1\}$. Hence, there is a bijective correspondence between $\mathbf{x} = (x_1, \dots, x_n) \in \mathcal{L}$ and $\mathbf{y} = (y_{10}, \dots, y_{1M}, \dots, y_{n0}, \dots, y_{nM}) \in \{0, 1\}^{n \times (M+1)}$ where $x_i = \sum_{j=0}^M 2^{-j} y_{ij}$. Therefore, the restriction of function f to the discrete set \mathcal{L} can be represented as $\tilde{f} : \{0, 1\}^{n \times (M+1)} \rightarrow \mathbb{R}$ such that $\tilde{f}(\mathbf{y}) = f(\mathbf{x})$ where $x_i = \sum_{j=0}^M 2^{-j} y_{ij}$.

Consider a multilinear extension $F : [0, 1]^{n \times (M+1)} \rightarrow \mathbb{R}$ of \tilde{f} defined as follows.

$$F(\mathbf{z}) := \sum_{S \subset [n \times (M+1)]} \tilde{f}(\mathbf{1}_S) \prod_{(i,j) \in S} z_{ij} \prod_{(i,j) \notin S} (1 - z_{ij})$$

By the definition, $F(\mathbf{z})$ can be seen as $\mathbb{E}[\tilde{f}(\mathbf{1}_S)]$ where the $(ij)^{th}$ -coordinate of $\mathbf{1}_S$ equals 1 (i.e., $(\mathbf{1}_S)_{ij} = 1$) with probability z_{ij} . We remark some properties of the multilinear function F :

$$\frac{\partial F}{\partial z_{ij}}(\mathbf{z}) = F(\mathbf{z}_{-ij}, 1) - F(\mathbf{z}_{-ij}, 0),$$

and in general,

$$\frac{\partial F}{\partial z_{ij}}(\mathbf{z}) = \frac{F(\mathbf{z}_{-ij}, z_{ij}^*) - F(\mathbf{z}_{-ij}, z_{ij})}{z_{ij}^* - z_{ij}} \quad \forall 0 \leq z_{ij}^* \neq z_{ij} \leq 1$$

where these equalities follow the fact that F is linear with respect to any z_{ij} . Therefore,

$$\langle \nabla F(z), z^* - z \rangle = \sum_{i=1}^n \sum_{j=0}^M [F(z_{-ij}, z_{ij}^*) - F(z_{-ij} z_{ij})] \quad (4.4)$$

Evaluation. Given a function f , it requires an exponential number of queries to f in order to evaluate exactly the linear extension F . However, one can evaluate F approximately up to any precision by the following lemma (proof can be found in [123]).

Lemma 4.2 Let $\mathbf{1}_{S_1}, \dots, \mathbf{1}_{S_k}$ be independent random vectors in $\{0, 1\}^{n \times (M+1)}$ where for every $1 \leq h \leq k$, element (i, j) appears independently in S_h with probability z_{ij} . Then,

$$\left| \frac{1}{k} \sum_{h=1}^k \tilde{f}(\mathbf{1}_{S_h}) - F(z) \right| \leq \epsilon |\max_S \tilde{f}(\mathbf{1}_S)|$$

with probability at least $1 - e^{-k\epsilon^2/4}$.

Online Learning Algorithm. We consider the following algorithm for the general problem where the domain $\mathcal{K} = [0, 1]^n$.

Algorithm 3 Algorithm for the general problem with $\mathcal{K} = [0, 1]^n$.

- 1: Initially, let z^1 be an arbitrary point in $[0, 1]^{n \times (M+1)}$.
- 2: **for** $t = 1$ to T **do**
- 3: Round z^t to a random solution $x^t \in \mathcal{L}$ corresponding to the point $\mathbf{1}_{S^{t+1}} \in \{0, 1\}^{n \times (M+1)}$ such that element (i, j) appears in S^{t+1} with probability z_{ij}^{t+1} .
- 4: Play x^t .
- 5: Let $F^t : [0, 1]^{n \times (M+1)} \rightarrow \mathbb{R}$ be the multilinear relaxation of the discretization of f^t (defined above).
- 6: Update and get a solution $z^{t+1} \in [0, 1]^{n \times (M+1)}$ by applying the mirror descent framework on F^t (Section 4.1.1). Specifically,

$$z^{t+1} = \arg \max_{z \in [0, 1]^{n \times (M+1)}} \{ \langle \eta \nabla F^t(z^t), z - z^t \rangle - D_{\Phi}(z \| z^t) \}.$$

Note that ∇F^t can be computed using Equation (4.4).

- 7: **end for**
-

Theorem 4.2 Given an arbitrarily small constant ϵ and let $f^t : [0, 1]^n \rightarrow \mathbb{R}$ be the reward functions for $1 \leq t \leq T$. Let F^t be the multilinear extension of the discretization of f^t based on a lattice \mathcal{L} (defined earlier). Assume that for every $x \in [0, 1]^n$, there exists $\bar{x} \in \mathcal{L}$ such that $|f^t(x^*) - f^t(\bar{x})| \leq \epsilon$. Moreover, assume that for every $1 \leq t \leq T$, the multilinear extension F^t is (λ, μ) -concave and $\|\nabla F^t\|_*$ is bounded by L and $D_{\Phi}(\cdot \| \cdot)$ is bounded by G^2 . Then, the online randomized algorithm (described above) achieves

$$\sum_{t=1}^T \mathbb{E}[f^t(x^t)] \geq \frac{\lambda}{\mu} \sum_{t=1}^T f^t(x^*) - O\left(\frac{GL}{\mu} \sqrt{2\alpha_{\Phi} T} + \frac{\lambda}{\mu} T\epsilon\right)$$

Moreover, the algorithm makes a polynomial number of value queries to functions f^t 's.

Proof Let $\mathbf{x}^* \in \arg \max_{\mathbf{x} \in [0,1]^n} \sum_{t=1}^T f^t(\mathbf{x})$ be the best solution in hindsight. By the assumption of \mathcal{L} , \mathbf{x}^* is close to some $\bar{\mathbf{x}} \in \mathcal{L}$, i.e., $|f^t(\mathbf{x}^*) - f^t(\bar{\mathbf{x}})| \leq \epsilon$. Let $\mathbf{1}_{\bar{s}} \in \{0,1\}^{n \times (M+1)}$ be the point corresponding to $\bar{\mathbf{x}}$. Note that $F^t(\mathbf{1}_{\bar{s}}) = \tilde{f}^t(\mathbf{1}_{\bar{s}})$ for every t . By Theorem 4.1, we have

$$\sum_{t=1}^T F^t(\mathbf{z}^t) \geq \frac{\lambda}{\mu} \sum_{t=1}^T F^t(\mathbf{1}_{\bar{s}}) - O\left(\frac{GL}{\mu} \sqrt{2\alpha_{\Phi} T}\right)$$

Hence, we deduce the regret bound:

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}[f^t(\mathbf{x}^t)] &= \sum_{t=1}^T \mathbb{E}[\tilde{f}^t(\mathbf{1}_{S^t})] \geq \frac{\lambda}{\mu} \sum_{t=1}^T \tilde{f}^t(\mathbf{1}_{\bar{s}}) - O\left(\frac{GL}{\mu} \sqrt{2\alpha_{\Phi} T}\right) \\ &= \frac{\lambda}{\mu} \sum_{t=1}^T f^t(\bar{\mathbf{x}}) - O\left(\frac{GL}{\mu} \sqrt{2\alpha_{\Phi} T}\right) \\ &\geq \frac{\lambda}{\mu} \sum_{t=1}^T [f^t(\mathbf{x}^*) - \epsilon] - O\left(\frac{GL}{\mu} \sqrt{2\alpha_{\Phi} T}\right) \\ &\geq \frac{\lambda}{\mu} \sum_{t=1}^T f^t(\mathbf{x}^*) - O\left(\frac{GL}{\mu} \sqrt{2\alpha_{\Phi} T} + \frac{\lambda}{\mu} T\epsilon\right). \end{aligned}$$

We consider the complexity of the algorithm. At time t , the algorithm needs to compute \mathbf{z}^{t+1} such that

$$\mathbf{z}^{t+1} = \arg \min_{\mathbf{z} \in \mathcal{D}} \{ \langle -\eta \nabla F^t(\mathbf{z}^t), \mathbf{z} - \mathbf{z}^t \rangle + D_{\Phi}(\mathbf{z} \|\mathbf{z}^t) \}$$

and round \mathbf{z}^{t+1} to \mathbf{x}^{t+1} . By the property of multilinear extension, in particular Equation (4.4), one can evaluate the term $\langle \nabla F^t(\mathbf{z}^t), \mathbf{z} - \mathbf{z}^t \rangle$ by computing $F^t(\mathbf{z}^t)$ and $n(M+1)$ terms $F^t(\mathbf{z}^t_{-ij}, z_{ij})$ for $1 \leq i \leq n$ and $0 \leq j \leq M$. These values can be computed approximately up to high precision with a polynomial numbers of value queries to function f^t . For example, given an arbitrary $\epsilon' > 0$, applying Lemma 4.2 with $k = \frac{8 \log T}{\epsilon'^2}$ one can approximate the value of F to error ϵ' with probability $1 - T^{-2}$. Besides, one can apply the standard independent rounding in Step 2 in order to round \mathbf{z}^{t+1} to \mathbf{x}^{t+1} . Hence, the algorithm needs to make only a polynomial numbers of value queries at each time step. \square

4.2 Applications to Fictitious Play in Smooth Auctions

We consider adaptive dynamics in auctions. In the setting, there is an underlying auction \mathbf{o} and there are n players, each player i has a set of actions \mathcal{A}_i (that can be arbitrarily large but finite) and a valuation function v_i taking values in $[0,1]$. In each time step $1 \leq t \leq T$, each player i selects a strategy which is a distribution in $\Delta(\mathcal{A}_i)$ according to some adaptive dynamic. The strategy profile at time t is denoted as $\boldsymbol{\sigma}^t \in \Delta(\mathcal{A})$. Given the strategy profile $\boldsymbol{\sigma}^t$, the auction induces a social welfare $\text{SW}(\mathbf{o}, \boldsymbol{\sigma}^t) := \mathbb{E}_{\mathbf{a} \sim \boldsymbol{\sigma}^t} [\text{SW}(\mathbf{o}(\mathbf{a}); \mathbf{v})]$. In this setting, we study the performance of adaptive dynamics, especially the ones which are not guaranteed to fulfill the vanishing regret condition, and eventually design dynamics/auctions with performance guarantee. Among others, *fictitious play* is an interesting, widely-studied dynamic which attracts a lot of attention in the community. In this section, we will study the performance of a version of fictitious play in smooth auctions.

Valuation-Oriented Fictitious Play. Consider the Perturbed Discrete Time Fictitious Play (PDTFP). Initially, each player chooses some arbitrary action. At time $t + 1$, given a strategy profile σ^t where $\sigma_i^t \in \Delta(\mathcal{A}_i)$ and perturbations $N_i^t : \Delta(\mathcal{A}_i) \rightarrow \mathbb{R}$ for $1 \leq i \leq n$, player i selects a mixed strategy σ_i^{t+1} such that

$$\sigma_i^{t+1} \in \arg \max_{\sigma_i \in \Delta(\mathcal{A}_i)} \mathbb{E}_{a_i \sim \sigma_i} \mathbb{E}_{\mathbf{a}_{-i}^t \sim \sigma_{-i}^t} [v_i(\mathbf{a}_{-i}^t, a_i)] - \frac{1}{\eta} N_i^t(\sigma_i)$$

Equivalently,

$$\sigma_i^{t+1} \in \arg \max_{\sigma_i \in \Delta(\mathcal{A}_i)} \mathbb{E}_{a_i \sim \sigma_i} \mathbb{E}_{\mathbf{a}^t \sim \sigma^t} [v_i(\mathbf{a}_{-i}^t, a_i) - v_i(\mathbf{a}^t)] - \frac{1}{\eta} N_i^t(\sigma_i) \quad (4.5)$$

since $\mathbb{E}[v_i(\mathbf{a}^t)]$ is already determined. One common example of perturbations is the *relative entropy* (or *Kullback-Leibler divergence*), defined as

$$N_i^t(\sigma_i) = \sum_{a_i \in \mathcal{A}_i} \sigma_i(a_i) \log \frac{\sigma_i(a_i)}{\sigma_i^t(a_i)}$$

which is the Bregman divergence with the negative entropy function

$$\Phi_i(\sigma_i) = \sum_{a_i \in \mathcal{A}_i} \sigma_i(a_i) \log \sigma_i(a_i).$$

Let V_i be the multilinear extension of the valuation v_i of player i (construction in Section 4.1.2) where now the corresponding lattice is the set of pure strategies \mathcal{A} . Note that the social welfare is the sum of all player valuations. Given an action profile \mathbf{a}^t , define $\nabla^t(\mathbf{a}^t) : \mathbb{R}^n \rightarrow \mathbb{R}$ such as

$$\langle \nabla^t(\mathbf{a}^t), \mathbf{x} \rangle = \sum_{i=1}^n \frac{\partial V_i(\mathbf{a})}{\partial a_i} \cdot x_i.$$

As V_i is the multilinear extension of v_i , for every action \mathbf{a}^* we have

$$\langle \nabla^t(\mathbf{a}^t), \mathbf{a}^* - \mathbf{a}^t \rangle = \sum_{i=1}^n [V_i(a_i^*, \mathbf{a}_{-i}^t) - V_i(a_i^t, \mathbf{a}_{-i}^t)].$$

The PDTFP dynamic can be cast as the mirror descent algorithm. By Equation (4.5) — the update rules of PDTFP dynamic — at every time step t , strategy profile $\sigma^{t+1} = (\sigma_1^{t+1}, \dots, \sigma_n^{t+1})$ is exactly the solution of the mirror descent update Equation (4.1):

$$\sigma^{t+1} \in \arg \max_{\sigma \in \Delta(\mathcal{A})} \mathbb{E}_{\mathbf{a} \sim \sigma} \mathbb{E}_{\mathbf{a}^t \sim \sigma^t} [\langle \nabla^t(\mathbf{a}^t), \mathbf{a} - \mathbf{a}^t \rangle - \frac{1}{\eta} D_\Phi(\mathbf{a} \parallel \mathbf{a}^t)],$$

where Φ is a strongly convex function such that $\mathbb{E}[D_\Phi(\mathbf{a} \parallel \mathbf{a}^t)] = \sum_{i=1}^n N_i^t(\sigma_i)$. Again, if the perturbations N_i^t are relative entropy functions then $\Phi(\sigma) = \sum_{i=1}^n \Phi_i(\sigma_i) = \sum_{i=1}^n \sum_{a_i \in \mathcal{A}_i} \sigma_i(a_i) \log \sigma_i(a_i)$.

Note that the PDTFP dynamic associated to a specific choice of entropy function is usually called *smooth fictitious play* [61] (or *logit dynamic*). Benaïm and Faure [21] have provided an explicit example [21, Example 1.2] showing that this dynamic does not always admit vanishing regret (consistent in their terms).

In the following, we prove the regret bound of the PDTFP dynamics (without the vanishing regret condition) in smooth auctions.

Theorem 4.3 *If the underlying auction \mathfrak{o} is a (λ, μ) -smooth and $D_{\Phi}(\cdot \| \cdot)$ is bounded by G^2 then the PDTFP dynamic achieves $\left(\frac{\lambda}{1+\mu}, R(T)\right)$ -regret where where $R(T) = O\left(\frac{G\sqrt{T}}{1+\mu}\right)$. In particular, $R(T) = O\left(\frac{\sqrt{T \log(n|A|)}}{1+\mu}\right)$ if the perturbation is the relative entropy function.*

Proof The analysis follows closely the one of Theorem 4.1 with some modifications. For simplicity, without loss of generality, assume that the distributions $\bar{D}_1, \dots, \bar{D}_n$ in the definition of smooth auctions (Definition 1.9) give rise to a pure strategy profile $\bar{\mathbf{a}}$ and at any time step t , the PDTFP dynamic outputs a pure profile \mathbf{a}^t . The analysis remains the same for general distributions/mixed profiles by putting additional expectations into some formula.

As the underlying auction is (λ, μ) -smooth, given a fixed valuation profile \mathbf{v} , there exists a strategy profile $\bar{\mathbf{a}}$ such that for any profile \mathbf{a} , it holds that

$$\sum_{i=1}^n u_i(\bar{a}_i, \mathbf{a}_{-i}; v_i) \geq \lambda \cdot \text{OPT}(\mathbf{v}) - \mu \cdot \text{REV}(\mathbf{a}),$$

where $\text{OPT}(\mathbf{v})$ stands for the optimal welfare given the valuation profile \mathbf{v} . We first derive an useful inequality based on the smoothness of the auction. We have

$$\begin{aligned} \langle \nabla^t(\mathbf{a}^t), \bar{\mathbf{a}} - \mathbf{a}^t \rangle &= \sum_i [V_i(\bar{a}_i, \mathbf{a}_{-i}^t; \mathbf{v}) - V_i(a_i, \mathbf{a}_{-i}^t; \mathbf{v})] \\ &= \sum_i [u_i(\bar{a}_i, \mathbf{a}_{-i}^t; v_i) + p_i(\bar{a}_i, \mathbf{a}_{-i}^t; v_i)] - \text{SW}(\mathbf{a}^t; \mathbf{v}) \\ &\geq \lambda \cdot \text{OPT}(\mathbf{v}) - \mu \cdot \text{REV}(\mathbf{a}^t; \mathbf{v}) - \text{SW}(\mathbf{a}^t; \mathbf{v}) \\ &\geq \lambda \cdot \text{OPT}(\mathbf{v}) - (1 + \mu) \cdot \text{SW}(\mathbf{a}^t). \end{aligned} \quad (4.6)$$

The first inequality follows the (λ, μ) -smoothness and the non-negativity of payments p_i 's. The second inequality is obvious since the revenue is always smaller than the welfare. We remark that Inequality (4.6) is similar to (but not the same as) the notation of (\cdot, \cdot) -concavity since it can written as

$$\langle \nabla \text{SW}(\mathbf{a}^t), \bar{\mathbf{a}} - \mathbf{a}^t \rangle \geq \lambda \cdot \text{SW}(\mathbf{a}^*) - (1 + \mu) \cdot \text{SW}(\mathbf{a}^t)$$

where \mathbf{a}^* is the optimal strategy. Hence, there would be a connection between concavity and smoothness.

Define the potential as $\Psi^t = \frac{1}{\eta} D_{\Phi}(\bar{\mathbf{a}} \| \mathbf{a}^t)$. Note that here we use the Bregman divergence from the strategy $\bar{\mathbf{a}}$ (induced by the smooth auction) to \mathbf{a}^t instead of the Bregman divergence from the optimal strategy \mathbf{a}^* to \mathbf{a}^t (as in Theorem 4.1). By the same arguments as for proving Inequality (4.2), we have

$$\eta(\Psi^{t+1} - \Psi^t) = D_{\Phi}(\bar{\mathbf{a}} \| \mathbf{a}^{t+1}) - D_{\Phi}(\bar{\mathbf{a}} \| \mathbf{a}^t) \leq -\eta \langle \nabla^t(\mathbf{a}^t), \bar{\mathbf{a}} - \mathbf{a}^t \rangle + \frac{\eta^2}{2\alpha_{\Phi}} \|\nabla^t(\mathbf{a}^t)\|_*^2.$$

Given the valuation profile \mathbf{v} , let \mathbf{a}^* be the action that gives the optimal welfare, i.e., $\text{SW}(\mathbf{a}^*; \mathbf{v}) = \text{OPT}(\mathbf{v})$. Using the same arguments as in the proof of Theorem 4.1,

we have

$$\begin{aligned}
\sum_{t=1}^T (\lambda \text{Sw}(\mathbf{a}^*) - (1 + \mu) \text{Sw}(\mathbf{a}^t)) &\leq \Psi^1 + \sum_{t=1}^T \left[\lambda \text{Sw}(\mathbf{a}^*) - (1 + \mu) \text{Sw}(\mathbf{a}^t) + \Psi^{t+1} - \Psi^t \right] \\
&\leq \Psi^1 + \sum_{t=1}^T \left[\underbrace{\lambda \text{OPT}(\mathbf{v}) - (1 + \mu) \text{Sw}(\mathbf{a}^t) - \langle \nabla^t(\mathbf{a}^t), \bar{\mathbf{a}} - \mathbf{a}^t \rangle}_{\leq 0 \text{ by Inequality (4.6)}} + \frac{\eta}{2\alpha_\Phi} \|\nabla^t(\mathbf{a}^t)\|_*^2 \right] \\
&\leq \frac{1}{\eta} D_\Phi(\bar{\mathbf{a}} \|\mathbf{a}^1) + \frac{\eta}{2\alpha_\Phi} \sum_{t=1}^T \|\nabla^t(\mathbf{a}^t)\|_*^2.
\end{aligned}$$

Thus,

$$\sum_{t=1}^T \text{Sw}(\mathbf{a}^t) \geq \frac{\lambda}{1 + \mu} \sum_{t=1}^T \text{Sw}(\mathbf{a}^*) - \frac{1}{(1 + \mu)\eta} D_\Phi(\bar{\mathbf{a}} \|\mathbf{a}^1) - \frac{\eta}{(1 + \mu)2\alpha_\Phi} \sum_{t=1}^T \|\nabla^t(\mathbf{a}^t)\|_*^2$$

Note that if player valuations are in the range $[0, 1]$, then

$$\|\nabla^t(\mathbf{a}^t)\|_* \leq \|\nabla^t(\mathbf{a}^t)\|_\infty \leq 1.$$

By the theorem assumptions, $D_\Phi(\bar{\mathbf{a}} \|\mathbf{a}^1) \leq G^2$. Hence, choosing $\eta = O(G/\sqrt{T})$, the PDTFP dynamic achieves $(\frac{\lambda}{1+\mu}, R(T))$ -regret where $R(T) = O(\frac{G\sqrt{T}}{1+\mu})$.

Consider the particular PDTFP dynamic with relative entropy perturbation. Function $\Phi(\boldsymbol{\sigma})$ is $\alpha_\Phi = \frac{1}{2\ln 2}$ -strongly convex (due to Pinsker's inequality). Moreover,

$$D_\Phi(\bar{\mathbf{a}} \|\mathbf{a}^1) \leq \max_i \log(n|\mathcal{A}_i|) \leq \log(n|\mathcal{A}|).$$

Therefore, choosing $\eta = O(1/\sqrt{T \log(n|\mathcal{A}|)})$, the PDTFP dynamic with relative entropy perturbation achieves $(\frac{\lambda}{1+\mu}, R(T))$ -regret where $R(T) = O(\frac{\sqrt{T \log(n|\mathcal{A}|)}}{1+\mu})$. \square

4.3 Online Simultaneous Second-Price Auctions with Reserve Prices

In this section, we are interested in the objective of maximizing the revenue. In the setting, there are n bidders and m items to be sold to these bidders. At each time step $t = 1, 2, \dots, T$, the auctioneer selects reserve prices $\mathbf{r}_i^t = (r_{i1}^t, \dots, r_{im}^t)$ for each bidder i where r_{ij} is the reserve price of item j for bidder i . Subsequently, every bidder i picks a bid vector $\mathbf{b}_i^t = (b_{i1}^t, \dots, b_{im}^t)$ where b_{ij}^t is the bid of bidder i on item $1 \leq j \leq m$. Then the auction for each item $1 \leq j \leq m$ works as follows: (1) remove all bidders i with $b_{ij}^t < r_{ij}^t$; (2) run the second-price auction on the remaining bidders to determine the winner of item j ; (3) charge the winner of item j the larger of r_{ij}^t and the second highest bid among the bids b_{ij}^t of remaining bidders.

Denote the revenue of selling item j as $\text{REV}_j(\mathbf{r}^t, \mathbf{b}^t)$ where $\mathbf{b}^t = (b_1^t, \dots, b_n^t)$ and $\mathbf{r}^t = (r_1^t, \dots, r_n^t)$. The revenue of the auctioneer at time step t is $\text{REV}(\mathbf{r}^t, \mathbf{b}^t) = \sum_j \text{REV}_j(\mathbf{r}^t, \mathbf{b}^t)$. The goal of the auctioneer is to achieve the total revenue approximately close to that achieved by the best fixed reserve-price auction.

In the setting, by scaling assume that every bid satisfies $0 \leq b_{ij} \leq 1$ for every i, j , and make the same assumption for every reserve price. The convex domain $\mathcal{K} = [0, 1]^{n \times m}$. Consider the set of values $\{\ell \cdot 2^{-M} : 0 \leq \ell \leq 2^M\}$ for some large

parameter M as a discretization of $[0, 1]$. Observe that for any reserve price vector \mathbf{r} , $|\text{REV}(\mathbf{r}, \mathbf{b}) - \text{REV}(\bar{\mathbf{r}}, \mathbf{b})| \leq m \cdot 2^{-M}$ where $\bar{\mathbf{r}}$ is a reserve price vector such that \bar{r}_{ij} is the largest multiple of 2^{-M} smaller than r_{ij} for every i, j . Therefore, one can approximate the revenue up to any arbitrary precision by restricting the reserve price values in the discretization (by choosing large enough M).

Consider the lattice $\mathcal{L} = \{\ell \cdot 2^{-M} : 0 \leq \ell \leq 2^M\}^{n \times m}$. Note that $\mathcal{K} \cap \mathcal{L} = \mathcal{L}$. We slightly abuse notation by denoting $\text{REV}_j(\mathbf{1}_S, \mathbf{b})$ as $\text{REV}_j(\mathbf{r}, \mathbf{b})$ where $\mathbf{1}_S$ is the point in the lattice corresponding to the reserve price \mathbf{r} . Following the construction of multilinear extension on the lattice \mathcal{L} described in Section 4.1.2, the multilinear extension of the revenue is defined as follows. Given a bid vector \mathbf{b} , $\overline{\text{REV}}(\cdot, \mathbf{b}) : [0, 1]^{n \times m \times (M+1)} \rightarrow \mathbb{R}$ such that

$$\overline{\text{REV}}(\mathbf{z}, \mathbf{b}) = \sum_{S \subset [n \times m \times (M+1)]} \left(\sum_{j=1}^m \text{REV}_j(\mathbf{1}_S, \mathbf{b}) \right) \prod_{(i,j,k) \in S} z_{ijk} \prod_{(i,j,k) \notin S} (1 - z_{ijk})$$

Online Reserve-Price Algorithm. Initially, let $\mathbf{z}^1 = \mathbf{r}^1$ be an arbitrary feasible reserve-price. At each time step $t > 1$, compute a random reserve price \mathbf{r}^{t+1} using the online algorithm in Section 4.1.2. Specifically, compute

$$\mathbf{z}^{t+1} = \arg \max_{\mathbf{z} \in [0,1]^{n \times m \times (M+1)}} \left\{ \eta \langle \nabla \overline{\text{REV}}(\mathbf{z}^t, \mathbf{b}^t), \mathbf{z} - \mathbf{z}^t \rangle - D_\Phi(\mathbf{z} \parallel \mathbf{z}^t) \right\}$$

where Φ is the negative entropy function. Then, round \mathbf{z}^{t+1} to a random reserve-price $\mathbf{r}^{t+1} \in \mathcal{L}$. In other words, $\mathbf{r}^{t+1} \in \mathcal{L}$ corresponds to a random point $\mathbf{1}_{S^{t+1}} \in \mathcal{L}$ such that element $(i, j, k) \in S^t$ with probability z_{ijk}^{t+1} . Return \mathbf{r}^{t+1} or $\mathbf{0}$ each with probability $1/2$.

Note that the convex domain $\mathcal{K} = [0, 1]^{n \times m}$ so this algorithm is polynomial. Specifically, an offline algorithm (function) in this setting is $\text{REV}_j(\mathbf{1}_S, \mathbf{b})$, given reserve prices $\mathbf{1}_S$ and bids \mathbf{b} , it returns the revenue of selling item j . The algorithm needs to request only a polynomial number of value queries at each time step (Theorem 4.2).

Analysis. In order to analyze the performance of this algorithm, we study the properties of some related functions and then derive the regret bound for the algorithm.

Fix a bid vector \mathbf{b} . Let \mathbf{r}_j be a vector consisting of reserve prices on item j , i.e., $\mathbf{r}_j = (r_{1j}, \dots, r_{nj})$. As \mathbf{b} is fixed and the selling procedure of each item depends only on the reserve prices to the item, so for simplicity denote $\text{REV}_j(\mathbf{r}, \mathbf{b})$ as $\text{REV}_j(\mathbf{r}_j)$ and $\text{REV}(\mathbf{r}, \mathbf{b})$ as $\text{REV}(\mathbf{r})$. Define a function $h_j : \{0, 1\}^{n \times (M+1)} \rightarrow \mathbb{R}$ such that $h_j(\mathbf{1}_T) = \max\{\text{REV}_j(\mathbf{1}_T), \text{REV}_j(\mathbf{1}_\emptyset)\} = \max\{\text{REV}_j(\mathbf{r}), \text{REV}_j(\mathbf{0})\}$ where \mathbf{r}_j is the reserve price corresponding to $\mathbf{1}_T$ for $T \subset [n \times (M+1)]$. Let $H_j : [0, 1]^{n \times (M+1)} \rightarrow \mathbb{R}$ be the multilinear extension of h_j . Moreover, define $H : [0, 1]^{n \times m \times (M+1)} \rightarrow \mathbb{R}$ as the multilinear extension of $\max\{\text{REV}(\mathbf{r}), \text{REV}(\mathbf{0})\}$ defined as

$$H(\mathbf{z}) = \sum_{S \subset [n \times m \times (M+1)]} \max\{\text{REV}(\mathbf{1}_S), \text{REV}(\mathbf{1}_\emptyset)\} \prod_{(i,j,k) \in S} z_{ijk} \prod_{(i,j,k) \notin S} (1 - z_{ijk})$$

Lemma 4.3 *It holds that $H(\mathbf{z}) = \sum_{j=1}^m H_j(\mathbf{z}_j)$ where \mathbf{z}_j is the restriction of \mathbf{z} to the coordinate related to item j .*

Proof As items are sold separately,

$$H(\mathbf{z}) = \sum_{S \subset [n \times m \times (M+1)]} \left(\sum_{j=1}^m h_j(\mathbf{1}_A) \right) \prod_{(i,j,k) \in S} z_{ijk} \prod_{(i,j,k) \notin S} (1 - z_{ijk})$$

where $A \subset [n \times (M+1)]$ is the restriction of S on coordinates related to item j . Therefore,

$$\begin{aligned} H(\mathbf{z}) &= \sum_{j=1}^m \sum_{U \subset [n \times (m-1) \times M]} \underbrace{\left[\sum_{A \subset [n \times (M+1)]} h_j(\mathbf{1}_A) \prod_{(i,k) \in A} z_{ijk} \prod_{(i,k) \notin A} (1 - z_{ijk}) \right]}_{\text{independent of } U \text{ since the allocation of } j \text{ depends only on bids to item } j} \\ &\quad \cdot \prod_{(i,j',k) \in U} z_{ij'k} \prod_{(i,j',k) \notin U, j' \neq j} (1 - z_{ij'k}) \\ &= \sum_{j=1}^m \left[\sum_{A \subset [n \times (M+1)]} h_j(\mathbf{1}_A) \prod_{(i,k) \in A} z_{ijk} \prod_{(i,k) \notin A} (1 - z_{ijk}) \right] \\ &\quad \cdot \underbrace{\sum_{U \subset [n \times (m-1) \times (M+1)]} \prod_{(i,j',k) \in U} z_{ij'k} \prod_{(i,j',k) \notin U, j' \neq j} (1 - z_{ij'k})}_{=1} \\ &= \sum_{j=1}^m \left[\sum_{A \subset [n \times (M+1)]} h_j(\mathbf{1}_A) \prod_{(i,k) \in A} z_{ijk} \prod_{(i,k) \notin A} (1 - z_{ijk}) \right] = \sum_{j=1}^m H_j(\mathbf{z}_j) \end{aligned}$$

□

We will prove that H is $(1, 1)$ -concave. By Lemma 4.3, it is sufficient to prove that property for every function H_j .

Lemma 4.4 *Function H_j is $(1, 1)$ -concave.*

Proof We prove that the inequality of $(1, 1)$ -concavity holds for all points in the lattice. As the multilinear extension can be seen as the expectation over these points, the lemma will follow. Fix a bid profile $\mathbf{b}_j = (b_{1j}, \dots, b_{nj})$. Without loss of generality, assume that $b_{1j} \geq b_{2j} \geq \dots \geq b_{nj}$. Let \mathbf{r}_j and \mathbf{r}_j^* be two arbitrary reserve price vectors. We will show that

$$\begin{aligned} \sum_{i=1}^n \left[\max\{\text{REV}_j(\mathbf{r}_{-i,j}, r_{ij}^*), \text{REV}_j(\mathbf{0})\} - \max\{\text{REV}_j(\mathbf{r}_j), \text{REV}_j(\mathbf{0})\} \right] \\ \geq \max\{\text{REV}_j(\mathbf{r}_j^*), \text{REV}_j(\mathbf{0})\} - \max\{\text{REV}_j(\mathbf{r}_j), \text{REV}_j(\mathbf{0})\} \end{aligned} \quad (4.7)$$

where $\mathbf{r}_{-i,j}$ stands for the reserve price vectors on item j without the reserve price of bidder i .

Observe that the revenue $\max\{\text{REV}_j(\mathbf{r}'_j), \text{REV}_j(\mathbf{0})\}$ for every reserve price \mathbf{r}'_j is at least the second highest bid b_{2j} (that is obtained in $\text{REV}_j(\mathbf{0})$). In particular, for any reserve price \mathbf{r}'_j such that the auction either (1) removes the first bidder (with highest bid) or (2) removes the second bidder and $r'_{1j} \leq b_{2j}$, the revenue

$$\max\{\text{REV}_j(\mathbf{r}'_j), \text{REV}_j(\mathbf{0})\} = \text{REV}_j(\mathbf{0}).$$

Hence, $\max\{\text{REV}_j(\mathbf{r}'_j), \text{REV}_j(\mathbf{0})\} \neq \text{REV}_j(\mathbf{0})$ if and only if $b_{2j} < r'_{1j} \leq b_{1j}$.

By these observations, we deduce that

$$\max\{\text{REV}_j(\mathbf{r}_{-ij}, r_{ij}^*), \text{REV}_j(\mathbf{0})\} \neq \max\{\text{REV}_j(\mathbf{r}_j), \text{REV}_j(\mathbf{0})\}$$

if and only if $i = 1$ and

- either $b_{2j} \leq r_{1j} \neq r_{1j}^* \leq b_{1j}$;
- or $r_{1j}^* \in (b_{2j}, b_{1j}]$ but $r_{1j} \notin (b_{2j}, b_{1j}]$;
- or inversely $r_{1j} \in (b_{2j}, b_{1j}]$ but $r_{1j}^* \notin (b_{2j}, b_{1j}]$.

Thus, proving Inequality (4.7) is equivalent to showing that

$$\begin{aligned} \max\{\text{REV}_j(\mathbf{r}_{-1j}, r_{1j}^*), \text{REV}_j(\mathbf{0})\} - \max\{\text{REV}_j(\mathbf{r}_j), \text{REV}_j(\mathbf{0})\} \\ \geq \max\{\text{REV}_j(\mathbf{r}_j^*), \text{REV}_j(\mathbf{0})\} - \max\{\text{REV}_j(\mathbf{r}_j), \text{REV}_j(\mathbf{0})\}. \end{aligned}$$

Case 1: $b_{2j} \leq r_{1j} \neq r_{1j}^* \leq b_{1j}$. In this case, both sides are equal to $r_{1j}^* - r_{1j}$.

Case 2: $r_{1j}^* \in (b_{2j}, b_{1j}]$ but $r_{1j} \notin (b_{2j}, b_{1j}]$. In this case, both sides are equal to $r_{1j}^* - b_{2j}$.

Case 3: $r_{1j} \in (b_{2j}, b_{1j}]$ but $r_{1j}^* \notin (b_{2j}, b_{1j}]$. In this case, both sides are equal to $b_{2j} - r_{1j}$.

Case 4: the complementary of all previous cases. In this case, both sides are equal to 0.

Therefore, Inequality (4.7) holds and so the lemma follows. \square

Theorem 4.4 *The online reserve price algorithm achieves $(1/2, O(m\sqrt{nm}\sqrt{T} \log T))$ -regret.*

Proof Consider an imaginary algorithm which is similar to our online reserve price algorithm but at every step t , its gain is $\max\{\text{REV}_j(\mathbf{r}^t), \text{REV}_j(\mathbf{0})\}$. (This algorithm is called imaginary since one cannot decide which reserve price between \mathbf{r}^t and $\mathbf{0}$ is better when the bid vector is not known.) We verify the conditions of Theorem 4.2. The discretization satisfies the condition that for any given bids \mathbf{b} , for any reserve price \mathbf{r} , there exists a reserve price $\bar{\mathbf{r}}$ in the lattice which gives $|\text{REV}(\mathbf{r}, \mathbf{b}) - \text{REV}(\bar{\mathbf{r}}, \mathbf{b})| \leq m \cdot 2^{-M}$. For arbitrary $\epsilon > 0$, choose the parameter $M = \log(mT/\epsilon)$, $|\text{REV}(\mathbf{r}, \mathbf{b}) - \text{REV}(\bar{\mathbf{r}}, \mathbf{b})| \leq \epsilon/T$. Besides, bids are in the range $[0, 1]$, then

$$\|\nabla H^t\|_* \leq m \|\nabla H_j^t\|_\infty \leq m.$$

For Φ is the negative entropy function,

$$D_\Phi(\bar{\mathbf{z}}\|\mathbf{z}^1) \leq O(\log(2^{nm(M+1)})) = O(nm \log(mT))$$

with the chosen parameter M . Finally, Lemma 4.4 shows the $(1, 1)$ -concavity of H . Therefore, applying Theorem 4.2, the imaginary algorithm achieves the regret bound of $(1, R)$ where

$$R = O\left(m\sqrt{nm}\sqrt{T} \log\left(\frac{mT}{\epsilon}\right) + \epsilon\right) = O\left(m\sqrt{nm}\sqrt{T} \log T\right)$$

if one chooses $\epsilon = m$.

As the online reserve price algorithm selects at every step t either \mathbf{r}^t or $\mathbf{0}$ with probability $1/2$, the revenue of the algorithm is at least half that of the imaginary algorithm. The theorem follows. \square

Chapter 5

Conclusion

In this thesis, we have presented primal-dual approaches as unified techniques in order to design competitive algorithms for online problems, to analyze the efficiency of games and to study the dynamics of online learning processes. Designing online algorithms and online learning algorithms or analyzing the PoA are essentially reduced to constructing/studying the smoothness or concavity parameters (also with other standard properties/conditions). We have shown the applicability of the approaches on a wide variety of settings and have given simple and improved analyses for several problems in settings of different natures. The approaches have brought new ideas not only for the analyses and the understanding of current games but also for the design of new algorithms/games/dynamics and new concepts leading to improved efficiency.

Theoretically, non-convexity is considered as a strong barrier in optimization. However, it has been observed that in many settings such as deep learning, non-convex problems have been efficiently solved by different methods (for example, stochastic gradient descent, etc). As showed in the thesis, the notions of smoothness and (λ, μ) -concavity shed some light on the study of non-convex problems. We hope that our approach would contribute some elements towards the understanding of non-convex problems. Studying non-convex problems constitutes a direction of my research in the next years.

A major challenge in Learning is the design of solutions which are robust in dynamically evolving environments. An important research agenda is the design of efficient (polynomial time) algorithms with performance guarantee in adversarial non-stationary, non-stochastic environments by applying optimization methods that learn from experience and observations. As mentioned earlier, designing such algorithms is not possible in general adversarial environments. However, efficient online learning may be achievable in well-structured settings with regularity conditions. Characterizing conditions, or in general discovering the hidden regularity or hidden structures, under which efficient online learning algorithms exist is in my current research agenda.

Appendix A

Smoothness Parameters

A.1 Smoothness Parameters of Polynomials

We show technical lemmas in order to determine smoothness parameters for polynomials with non-negative coefficients. The following lemma has been proved in [45].

Lemma A.1 ([45]) *Let k be a positive integer. Let $0 < a(k) \leq 1$ be a function on k . Then, for any $x, y > 0$, it holds that*

$$y(x+y)^k \leq \frac{k}{k+1} a(k) x^{k+1} + b(k) y^{k+1}$$

where α is some constant and

$$b(k) = \begin{cases} \Theta \left(\alpha^k \cdot \left(\frac{k}{\log ka(k)} \right)^{k-1} \right) & \text{if } \lim_{k \rightarrow \infty} (k-1)a(k) = \infty, \quad (\text{A.1a}) \\ \Theta \left(\alpha^k \cdot k^{k-1} \right) & \text{if } (k-1)a(k) \text{ are bounded } \forall k, \quad (\text{A.1b}) \\ \Theta \left(\alpha^k \cdot \frac{1}{ka(k)^k} \right) & \text{if } \lim_{k \rightarrow \infty} (k-1)a(k) = 0. \quad (\text{A.1c}) \end{cases}$$

Lemma A.2 *For any sequences of non-negative real numbers $\{a_1, a_2, \dots, a_n\}$ and $\{b_1, b_2, \dots, b_n\}$ and for any polynomial g of degree k with non-negative coefficients, it holds that*

$$\sum_{i=1}^n \left[g \left(b_i + \sum_{j=1}^i a_j \right) - g \left(\sum_{j=1}^i a_j \right) \right] \leq \lambda(k) \cdot g \left(\sum_{i=1}^n b_i \right) + \mu(k) \cdot g \left(\sum_{i=1}^n a_i \right)$$

where $\mu(k) = \frac{k-1}{k}$ and $\lambda(k) = \Theta(k^{k-1})$. The same inequality holds for $\mu(k) = \frac{k-1}{k \ln k}$ and $\lambda(k) = \Theta((k \ln k)^{k-1})$.

Proof We first prove for $\mu(k) = \frac{k-1}{k}$ and $\lambda(k) = \Theta(k^{k-1})$. Let $g(z) = g_0 z^k + g_1 z^{k-1} + \dots + g_k$ with $g_t \geq 0 \forall t$. The lemma holds since it holds for every z^t for $0 \leq t \leq k$.

Specifically,

$$\begin{aligned}
\sum_{i=1}^n \left[g\left(b_i + \sum_{j=1}^i a_j\right) - g\left(\sum_{j=1}^i a_j\right) \right] &= \sum_{t=1}^k g_{k-t} \cdot \sum_{i=1}^n \left[\left(b_i + \sum_{j=1}^i a_j\right)^t - \left(\sum_{j=1}^i a_j\right)^t \right] \\
&\leq \sum_{t=1}^k g_{k-t} \cdot \left[t \cdot b_i \cdot \left(b_i + \sum_{j=1}^i a_j\right)^{t-1} \right] \\
&\leq \sum_{t=1}^k g_{k-t} \cdot \left[\lambda(t) \left(\sum_{i=1}^n b_i\right)^t + \mu(t) \left(\sum_{i=1}^n a_i\right)^t \right] \\
&\leq \lambda(k) \cdot g\left(\sum_{i=1}^n b_i\right) + \mu(k) \cdot g\left(\sum_{i=1}^n a_i\right)
\end{aligned} \tag{A.2}$$

The first inequality follows the convex inequality $(x + y)^{k+1} - x^{k+1} \leq (k + 1)y(x + y)^k$. The second inequality follows Lemma A.1 (Case A.1b and $a(k) = 1/(k + 1)$). The last inequality holds since $\mu(t) \leq \mu(k)$ and $\lambda(t) \leq \lambda(k)$ for $t \leq k$.

The case $\mu(k) = \frac{k-1}{k \ln k}$ and $\lambda(k) = \Theta((k \ln k)^{k-1})$ is proved similarly. The only different step is in inequality (A.2). In fact, applying Lemma A.1 (Case A.1c and $a(k) = \frac{1}{(k+1) \ln k}$), one gets the lemma inequality for $\mu(k) = \frac{k-1}{k \ln k}$ and $\lambda(k) = \Theta((k \ln k)^{k-1})$. \square

Appendix B

Concavity Parameters

B.1 Concavity Parameters of Multilinear Extension of Submodular Functions

Let $f : \{0,1\}^n \rightarrow \mathbb{R}^+$ be a monotone submodular function. We determine the concavity parameter of the multilinear extension F of f . By notation, $\mathbf{x} \leq \mathbf{y}$ iff $x_i \leq y_i$ for all $1 \leq i \leq n$. Recall the following useful properties of F which is due to the monotonicity and the diminishing return property of f , respectively.

$$F(\mathbf{x}) \leq F(\mathbf{y}), \quad \nabla F(\mathbf{x}) \geq \nabla F(\mathbf{y}) \quad \forall \mathbf{x} \leq \mathbf{y} \quad (\text{B.1})$$

The following theorem has been proved in [70]. We represent it here for completeness.

Lemma B.1 ([70]) *Let F be the multilinear extension of a monotone submodular function. Then, for any vectors \mathbf{x}, \mathbf{y} , it holds that*

$$\langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \geq F(\mathbf{y}) - 2F(\mathbf{x})$$

Proof For any vectors $\mathbf{x} \leq \mathbf{z}$, using Inequality (B.1), we have

$$\begin{aligned} F(\mathbf{z}) - F(\mathbf{x}) &= \int_0^1 \langle \mathbf{z} - \mathbf{x}, \nabla F(\mathbf{x} + t(\mathbf{z} - \mathbf{x})) \rangle dt \\ &\leq \int_0^1 \langle \mathbf{z} - \mathbf{x}, \nabla F(\mathbf{x}) \rangle dt = \langle \mathbf{z} - \mathbf{x}, \nabla F(\mathbf{x}) \rangle \end{aligned}$$

Therefore,

$$F(\mathbf{x} \vee \mathbf{y}) - F(\mathbf{x}) \leq \langle \mathbf{x} \vee \mathbf{y} - \mathbf{x}, \nabla F(\mathbf{x}) \rangle \quad (\text{B.2})$$

Similarly, for any vectors $\mathbf{x} \geq \mathbf{z}$ we have

$$\begin{aligned} F(\mathbf{z}) - F(\mathbf{x}) &= \int_0^1 \langle \mathbf{z} - \mathbf{x}, \nabla F(\mathbf{x} + t(\mathbf{z} - \mathbf{x})) \rangle dt \\ &\geq \int_0^1 \langle \mathbf{z} - \mathbf{x}, \nabla F(\mathbf{x}) \rangle dt = \langle \mathbf{z} - \mathbf{x}, \nabla F(\mathbf{x}) \rangle \end{aligned}$$

Therefore,

$$F(\mathbf{x} \wedge \mathbf{y}) - F(\mathbf{x}) \leq \langle \mathbf{x} \wedge \mathbf{y} - \mathbf{x}, \nabla F(\mathbf{x}) \rangle \quad (\text{B.3})$$

Summing (B.2) and (B.3) and note that $(\mathbf{x} \vee \mathbf{y}) + (\mathbf{x} \wedge \mathbf{y}) = \mathbf{x} + \mathbf{y}$, we get

$$F(\mathbf{x} \vee \mathbf{y}) + F(\mathbf{x} \wedge \mathbf{y}) - 2F(\mathbf{x}) \leq \langle \mathbf{y} - \mathbf{x}, \nabla F(\mathbf{x}) \rangle$$

Hence,

$$F(\mathbf{y}) - 2F(\mathbf{x}) \leq \langle \mathbf{y} - \mathbf{x}, \nabla F(\mathbf{x}) \rangle$$

since $F(\mathbf{y}) \leq F(\mathbf{x} \vee \mathbf{y})$ (monotonicity of F) and $F(\mathbf{x} \wedge \mathbf{y}) \geq 0$. □

Bibliography

- [1] Ajit Agrawal, Philip Klein, and R Ravi. When trees collide: An approximation algorithm for the generalized steiner problem on networks. *SIAM Journal on Computing*, 24(3):440–456, 1995.
- [2] Susanne Albers. Energy-efficient algorithms. *Commun. ACM*, 53(5):86–96, 2010.
- [3] S. Anand, Naveen Garg, and Amit Kumar. Resource augmentation for weighted flow-time explained by dual fitting. In *Proc. 23rd ACM-SIAM Symposium on Discrete Algorithms*, pages 1228–1241, 2012.
- [4] Matthew Andrews, Spyridon Antonakopoulos, and Lisa Zhang. Minimum-cost network design with (dis)economies of scale. In *Proc. 51th Symposium on Foundations of Computer Science (FOCS)*, pages 585–592, 2010.
- [5] Matthew Andrews, Antonio Fernández Anta, Lisa Zhang, and Wenbo Zhao. Routing for power minimization in the speed scaling model. *IEEE/ACM Trans. Netw.*, 20(1):285–294, 2012.
- [6] Antonios Antoniadis, Sungjin Im, Ravishankar Krishnaswamy, Benjamin Moseley, Viswanath Nagarajan, Kirk Pruhs, and Cliff Stein. Hallucination helps: Energy efficient virtual circuit routing. In *Proc. 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1141–1153, 2014.
- [7] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1): 121–164, 2012.
- [8] Robert J Aumann. Subjectivity and correlation in randomized strategies. *Journal of mathematical Economics*, 1(1):67–96, 1974.
- [9] Yossi Azar, Umang Bhaskar, Lisa Fleischer, and Debmalya Panigrahi. Online mixed packing and covering. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 85–100, 2013.
- [10] Yossi Azar, Nikhil R. Devanur, Zhiyi Huang, and Debmalya Panigrahi. Speed scaling in the non-clairvoyant model. In *Proc. 27th ACM on Symposium on Parallelism in Algorithms and Architectures*, pages 133–142, 2015.
- [11] Yossi Azar, Niv Buchbinder, TH Hubert Chan, Shahar Chen, Ilan Reuven Cohen, Anupam Gupta, Zhiyi Huang, Ning Kang, Viswanath Nagarajan, Joseph Seffi Naor, and Debmalya Panigrahi. Online algorithms for covering and packing problems with convex objectives. In *Proc. Symposium on Foundations of Computer Science (FOCS)*, 2016.
- [12] Maria-Florina Balcan and Nicholas J. A. Harvey. Learning submodular functions. In *Proc. Machine Learning and Knowledge Discovery in Databases*, pages 846–849, 2012.

- [13] Evripidis Bampis, Alexander V. Kononov, Dimitrios Letsios, Giorgio Lucarelli, and Maxim Sviridenko. Energy efficient scheduling and routing via randomized rounding. In *Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, pages 449–460, 2013.
- [14] Nikhil Bansal and Anupam Gupta. Potential-Function Proofs for First-Order Methods. *arXiv:1712.04581*, 2017.
- [15] Nikhil Bansal, Tracy Kimbrel, and Kirk Pruhs. Speed scaling to manage energy and temperature. *J. ACM*, 54(1), 2007.
- [16] Nikhil Bansal, Niv Buchbinder, and Joseph Naor. Randomized competitive algorithms for generalized caching. *SIAM Journal on Computing*, 41(2):391–414, 2012.
- [17] Nikhil Bansal, Niv Buchbinder, and Joseph Seffi Naor. A primal-dual randomized algorithm for weighted paging. *Journal of the ACM*, 59(4):19, 2012.
- [18] Nikhil Bansal, Ho-Leung Chan, Dmitriy Katz, and Kirk Pruhs. Improved bounds for speed scaling in devices obeying the cube-root rule. *Theory of Computing*, 8(1):209–229, 2012.
- [19] Reuven Bar-Yehuda and Shimon Even. A linear-time approximation algorithm for the weighted vertex cover problem. *Journal of Algorithms*, 2(2):198–203, 1981.
- [20] Martin Beckmann, CB McGuire, and Christopher B Winsten. Studies in the economics of transportation. Technical report, 1956.
- [21] Michel Benaïm and Mathieu Faure. Consistency of vanishingly smooth fictitious play. *Mathematics of Operations Research*, 38(3):437–450, 2013.
- [22] Kshipra Bhawalkar and Tim Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 700–709. SIAM, 2011.
- [23] Kshipra Bhawalkar, Sreenivas Gollapudi, and Kamesh Munagala. Coevolutionary opinion formation games. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 41–50. ACM, 2013.
- [24] Kshipra Bhawalkar, Martin Gairing, and Tim Roughgarden. Weighted congestion games: the price of anarchy, universal worst-case examples, and tightness. *ACM Transactions on Economics and Computation*, 2(4):14, 2014.
- [25] Vittorio Bilo. A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games. In *International Workshop on Approximation and Online Algorithms*, pages 215–228, 2012.
- [26] Avrim Blum and Jason D Hartline. Near-optimal online auctions. In *Proc. 16th Symposium on Discrete Algorithms*, pages 1156–1163, 2005.
- [27] Avrim Blum, Eyal Even-Dar, and Katrina Ligett. Routing without regret: On convergence to nash equilibria of regret-minimizing algorithms in routing games. *Theory of Computing*, 6(1):179–199, 2010.

- [28] Avrim Blum, Anupam Gupta, Yishay Mansour, and Ankit Sharma. Welfare and profit maximization with production costs. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 77–86. IEEE, 2011.
- [29] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [30] George W Brown. Iterative solution of games by fictitious play. *Activity analysis of production and allocation*, 13(1):374–376, 1951.
- [31] Sebastien Bubeck, Michael B Cohen, James R Lee, Yin Tat Lee, and Aleksander Madry. k -server via multiscale entropic regularization. In *Proc. 50th Symposium on Theory of Computing*, 2018.
- [32] Niv Buchbinder and Joseph Naor. Online primal-dual algorithms for covering and packing. *Mathematics of Operations Research*, 34(2):270–286, 2009.
- [33] Niv Buchbinder and Joseph Naor. The design of competitive online algorithms via a primal-dual approach. *Foundations and Trends in Theoretical Computer Science*, 3(2-3):93–263, 2009.
- [34] Niv Buchbinder and Joseph Naor. Online primal-dual algorithms for covering and packing. *Mathematics of Operations Research*, 34(2):270–286, 2009.
- [35] Niv Buchbinder, Kamal Jain, and Joseph Seffi Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. In *European Symposium on Algorithms*, pages 253–264, 2007.
- [36] Niv Buchbinder, Joseph Seffi Naor, R Ravi, and Mohit Singh. Approximation algorithms for online weighted rank function maximization under matroid constraints. In *International Colloquium on Automata, Languages, and Programming*, pages 145–156, 2012.
- [37] Niv Buchbinder, Shahar Chen, and Joseph Seffi Naor. Competitive analysis via regularization. In *Proc. 25th Symposium on Discrete Algorithms*, pages 436–444, 2014.
- [38] Niv Buchbinder, Moran Feldman, and Roy Schwartz. Online submodular maximization with preemption. In *Proc. 26th Symposium on Discrete Algorithms*, pages 1202–1216, 2015.
- [39] Niv Buchbinder, Shahar Chen, Joseph Naor, and Ohad Shamir. Unified algorithms for online learning and competitive analysis. *Math. Oper. Res.*, 41(2): 612–625, 2016.
- [40] Ioannis Caragiannis. Better bounds for online load balancing on unrelated machines. In *Proc. 19th ACM-SIAM Symposium on Discrete Algorithms*, pages 972–981, 2008.
- [41] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, Maria Kyropoulou, Brendan Lucier, Renato Paes Leme, and Eva Tardos. Bounding the inefficiency of outcomes in generalized second price auctions. *Journal of Economic Theory*, 156:343–388, 2015.
- [42] Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge university press, 2006.

- [43] Nicolo Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. Regret minimization for reserve prices in second-price auctions. *IEEE Transactions on Information Theory*, 61(1):549–564, 2015.
- [44] Chandra Chekuri and Sanjeev Khanna. On multidimensional packing problems. *SIAM journal on computing*, 33(4):837–851, 2004.
- [45] Johanne Cohen, Christoph Dürr, and Nguyen Kim Thang. Smooth inequalities and equilibrium inefficiency in scheduling games. In *Internet and Network Economics*, pages 350–363, 2012.
- [46] Roberto Cominetti, José R Correa, and Nicolás E Stier-Moses. The impact of oligopolistic competition in networks. *Operations Research*, 57(6):1421–1437, 2009.
- [47] Michele Conforti and Gérard Cornuéjols. Submodular set functions, matroids and the greedy algorithm: tight worst-case bounds and some generalizations of the rado-edmonds theorem. *Discrete applied mathematics*, 7(3):251–274, 1984.
- [48] José R Correa, Andreas S Schulz, and Nicolás E Stier-Moses. A geometric approach to the price of anarchy in nonatomic congestion games. *Games and Economic Behavior*, 64(2):457–469, 2008.
- [49] Constantinos Daskalakis and Vasilis Syrgkanis. Learning in auctions: Regret is hard, envy is easy. In *57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 219–228, 2016.
- [50] Nikhil R. Devanur and Zhiyi Huang. Primal dual gives almost optimal energy efficient online algorithms. In *Proc. 25th ACM-SIAM Symposium on Discrete Algorithms*, 2014.
- [51] Nikhil R. Devanur and Kamal Jain. Online matching with concave returns. In *Proc. 44th ACM Symposium on Theory of Computing*, pages 137–144, 2012.
- [52] Miroslav Dudik, Nika Haghtalab, Haipeng Luo, Robert E Schapire, Vasilis Syrgkanis, and Jennifer Wortman Vaughan. Oracle-efficient online learning and auction design. In *Proc. 58th Symposium on Foundations of Computer Science (FOCS)*, pages 528–539, 2017.
- [53] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American economic review*, 97(1):242–259, 2007.
- [54] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 201–210. ACM, 2013.
- [55] Michal Feldman, Nicole Immorlica, Brendan Lucier, Tim Roughgarden, and Vasilis Syrgkanis. The price of anarchy in large games. In *Proc. 48th Symposium on Theory of Computing (STOC)*, pages 963–976, 2016.
- [56] Moran Feldman, Joseph Naor, and Roy Schwartz. A unified continuous greedy algorithm for submodular maximization. In *Symposium on Foundations of Computer Science (FOCS)*, pages 570–579, 2011.

- [57] Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *Proc. 27th Symposium on Discrete Algorithms*, pages 1014–1033, 2016.
- [58] Dimitris Fotakis. A primal-dual algorithm for online non-uniform facility location. *Journal of Discrete Algorithms*, 5(1):141–148, 2007.
- [59] Dimitris Fotakis. On the competitive ratio for online facility location. *Algorithmica*, 50(1):1–57, 2008.
- [60] Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- [61] Drew Fudenberg and David K Levine. Consistency and cautious fictitious play. *Journal of Economic Dynamics and Control*, 19(5-7):1065–1089, 1995.
- [62] Drew Fudenberg and David K Levine. *The theory of learning in games*. MIT press, 1998.
- [63] Drew Fudenberg and Jean Tirole. *Game theory*, 1991.
- [64] Michel X Goemans and David P Williamson. A general approximation technique for constrained forest problems. *SIAM Journal on Computing*, 24(2):296–317, 1995.
- [65] Michel X Goemans, Nicholas JA Harvey, Satoru Iwata, and Vahab Mirrokni. Approximating submodular functions everywhere. In *Proc. 20th Symposium on Discrete algorithms*, pages 535–544, 2009.
- [66] Gero Greiner, Tim Nonner, and Alexander Souza. The bell is ringing in speed-scaled multiprocessor scheduling. *Theory of Computing Systems*, 54(1):24–44, 2014.
- [67] Anupam Gupta, Ravishankar Krishnaswamy, and Kirk Pruhs. Online primal-dual for non-linear optimization with applications to speed scaling. In *Proc. 10th Workshop on Approximation and Online Algorithms*, pages 173–186, 2012.
- [68] James Hannan. Approximation to bayes risk in repeated play. *Contributions to the Theory of Games*, 3:97–139, 1957.
- [69] Tobias Harks. Stackelberg strategies and collusion in network games with splittable flow. *Theory of Computing Systems*, 48(4):781–802, 2011.
- [70] Hamed Hassani, Mahdi Soltanolkotabi, and Amin Karbasi. Gradient methods for submodular maximization. In *Advances in Neural Information Processing Systems*, pages 5841–5851, 2017.
- [71] Avinatan Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. Non-price equilibria in markets of discrete goods. In *Proc. 12th ACM Conference on Electronic Commerce*, pages 295–296, 2011.
- [72] Elad Hazan and Satyen Kale. Extracting certainty from uncertainty: Regret bounded by variation in costs. *Machine learning*, 80(2-3):165–188, 2010.
- [73] Elad Hazan and Satyen Kale. Online submodular minimization. *Journal of Machine Learning Research*, 13:2903–2922, 2012.

- [74] Elad Hazan and Tomer Koren. The computational power of optimization in online learning. In *Proc. 48th Symposium on Theory of Computing*, pages 128–141, 2016.
- [75] Elad Hazan et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- [76] AJ Hoffman. On simple combinatorial optimization problems. *Discrete Mathematics*, 106:285–289, 1992.
- [77] Zhiyi Huang and Anthony Kim. Welfare maximization with production costs: a primal dual approach. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 59–72. SIAM, 2015.
- [78] Sungjin Im, Janardhan Kulkarni, Kamesh Munagala, and Kirk Pruhs. Selfish-migrate: A scalable algorithm for non-clairvoyantly scheduling heterogeneous processors. In *Proc. 55th IEEE Symposium on Foundations of Computer Science*, 2014.
- [79] Sungjin Im, Nathaniel Kell, Janardhan Kulkarni, and Debmalya Panigrahi. Tight bounds for online vector scheduling. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 525–544. IEEE, 2015.
- [80] Rishabh K Iyer, Stefanie Jegelka, and Jeff A Bilmes. Curvature and optimal algorithms for learning and minimizing submodular functions. In *Advances in Neural Information Processing Systems*, pages 2742–2750, 2013.
- [81] Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- [82] Robert Kleinberg and Tom Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *Proc. 44th Symposium on Foundations of Computer Science*, pages 594–605, 2003.
- [83] Peter Kling and Peter Pietrzyk. Profitable scheduling on multiple speed-scalable processors. *ACM Transactions on Parallel Computing*, 2(3):19, 2015.
- [84] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. *Computer science review*, 3(2):65–69, 2009.
- [85] Andreas Krause and Daniel Golovin. Submodular function maximization. In *Tractability: Practical Approaches to Hard Problems*, pages 71–104. Cambridge University Press, 2014.
- [86] Vijay Krishna. *Auction theory*. Academic press, 2009.
- [87] Ravishankar Krishnaswamy, Viswanath Nagarajan, Kirk Pruhs, and Cliff Stein. Cluster before you hallucinate: approximating node-capacitated network design and energy efficient routing. In *Proc. 46th Symposium on Theory of computing*, pages 734–743, 2014.
- [88] Janardhan Kulkarni and Vahab Mirrokni. Robust price of anarchy bounds via lp and fenchel duality. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1030–1049. SIAM, 2015.
- [89] Sébastien Lahaie, David M Pennock, Amin Saberi, and Rakesh V Vohra. Sponsored search auctions. *Algorithmic game theory*, pages 699–716, 2007.

- [90] Renato Paes Leme, Vasilis Syrgkanis, and Éva Tardos. Sequential auctions and externalities. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 869–886. SIAM, 2012.
- [91] Nick Littlestone and Manfred K Warmuth. The weighted majority algorithm. *Information and computation*, 108(2):212–261, 1994.
- [92] Giorgio Lucarelli, Nguyen Kim Thang, Abhinav Srivastav, and Denis Trystram. Online non-preemptive scheduling in a resource augmentation model based on duality. In *European Symposium on Algorithms*, 2016.
- [93] Patrick Maillé, Evangelos Markakis, Maurizio Naldi, George D Stamoulis, and Bruno Tuffin. Sponsored search auctions: an overview of research with emphasis on game theoretic aspects. *Electronic Commerce Research*, 12(3):265–300, 2012.
- [94] Konstantin Makarychev and Maxim Sviridenko. Solving optimization problems with diseconomies of scale via decoupling. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 571–580. IEEE, 2014.
- [95] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *Journal of the ACM (JACM)*, 54(5):22, 2007.
- [96] Ishai Menache and Mohit Singh. Online caching with convex costs. In *Proceedings of the 27th ACM symposium on Parallelism in Algorithms and Architectures*, pages 46–54. ACM, 2015.
- [97] Adam Meyerson. Online facility location. In *Proceedings. 42nd IEEE Symposium on Foundations of Computer Science*, pages 426–431, 2001.
- [98] Hervé Moulin and J-P Vial. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3-4):201–221, 1978.
- [99] Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [100] Uri Nadav and Tim Roughgarden. The limits of smoothness: A primal-dual framework for price of anarchy bounds. In *International Workshop on Internet and Network Economics*, pages 319–326, 2010.
- [101] Viswanath Nagarajan and Xiangkun Shen. Online covering with sum of ℓ_q -norm objectives. In *Proc. 44th Colloquium on Automata, Languages and Programming (ICALP)*, 2017.
- [102] John F Nash. Equilibrium points in n-person games. *Proc. Nat. Acad. Sci. USA*, 36(1):48–49, 1950.
- [103] Kim Thang Nguyen. Lagrangian duality in online scheduling with resource augmentation and speed scaling. In *Proc. 21st European Symposium on Algorithms*, pages 755–766, 2013.
- [104] Christos H Papadimitriou and Kenneth Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Courier Corporation, 1998.

- [105] Hans Peters. *Game theory: A Multi-leveled approach*. Springer, 2015.
- [106] Robert W Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- [107] Tim Roughgarden. Frontiers in mechanism design. Lecture 17, 2014.
- [108] Tim Roughgarden. Intrinsic robustness of the price of anarchy. *Journal of the ACM*, 62(5):32, 2015.
- [109] Tim Roughgarden. The price of anarchy in games of incomplete information. *ACM Transactions on Economics and Computation*, 3(1):6, 2015.
- [110] Tim Roughgarden and Florian Schoppmann. Local smoothness and the price of anarchy in splittable congestion games. *Journal of Economic Theory*, 156:317–342, 2015.
- [111] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2):236–259, 2002.
- [112] Tim Roughgarden and Éva Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 47(2):389–403, 2004.
- [113] Tim Roughgarden and Joshua R Wang. Minimizing regret with multiple reserves. In *Proc. 2016 ACM Conference on Economics and Computation*, pages 601–616, 2016.
- [114] Tim Roughgarden, Vasilis Syrgkanis, and Eva Tardos. The price of anarchy in auctions. *Journal of Artificial Intelligence Research*, 59:59–101, 2017.
- [115] Shai Shalev-Shwartz and Yoram Singer. A primal-dual perspective of online learning algorithms. *Machine Learning*, 69(2-3):115–142, 2007.
- [116] Shai Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.
- [117] Maxim Sviridenko, Jan Vondrák, and Justin Ward. Optimal approximation for submodular and supermodular optimization with bounded curvature. *Mathematics of Operations Research*, 2017.
- [118] Vasilis Syrgkanis and Eva Tardos. Bayesian sequential auctions. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 929–944. ACM, 2012.
- [119] Vasilis Syrgkanis and Eva Tardos. Composable and efficient mechanisms. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 211–220. ACM, 2013.
- [120] Vasilis Syrgkanis, Akshay Krishnamurthy, and Robert Schapire. Efficient algorithms for adversarial contextual learning. In *International Conference on Machine Learning*, pages 2159–2168, 2016.
- [121] Hal R Varian. Position auctions. *international Journal of industrial Organization*, 25(6):1163–1178, 2007.
- [122] Rakesh V Vohra. *Mechanism design: a linear programming approach*, volume 47. Cambridge University Press, 2011.

-
- [123] Jan Vondrák. Polyhedral techniques in combinatorial optimization. Lecture notes, Nov 18 2010.
- [124] Jan Vondrák. Submodularity and curvature: The optimal algorithm. *RIMS Kokyuroku Bessatsu*, 2010.
- [125] John Glen Wardrop. Some theoretical aspects of road traffic research. *Proceedings of the institution of civil engineers*, 1(3):325–362, 1952.
- [126] David P Williamson. The primal-dual method for approximation algorithms. *Mathematical Programming*, 91(3):447–478, 2002.
- [127] David P Williamson and David B Shmoys. *The design of approximation algorithms*. Cambridge university press, 2011.