Game Efficiency through Linear Programming Duality

Nguyễn Kim Thắng*
IBISC, University Paris-Saclay, France

November 20, 2018

Abstract

The efficiency of a game is typically quantified by the price of anarchy (PoA), defined as the worst ratio of the value of an equilibrium — solution of the game — and that of an optimal outcome. Given the tremendous impact of tools from mathematical programming in the design of algorithms and the similarity of the price of anarchy and different measures such as the approximation and competitive ratios, it is intriguing to develop a duality-based method to characterize the efficiency of games.

In the paper, we present an approach based on linear programming duality to study the efficiency of games. We show that the approach provides a general recipe to analyze the efficiency of games and also to derive concepts leading to improvements. The approach is particularly appropriate to bound the PoA. Specifically, in our approach the dual programs naturally lead to competitive PoA bounds that are (almost) optimal for several classes of games. The approach indeed captures the smoothness framework and also some current non-smooth techniques/concepts. We show the applicability to the wide variety of games and environments, from congestion games to Bayesian welfare, from full-information settings to incomplete-information ones.

*Research supported by the ANR project OATA n° ANR-15-CE40-0015-01

1 Introduction

Algorithmic Game Theory — a domain at the intersection of Game Theory and Algorithms — has been extensively studied in the last two decades. The development of the domain, as well as those of many other research fields, have witnessed a common phenomenon: interesting notions, results have been flourished at the early stage, then deep methods, techniques have been established at a more mature stage leading to further achievements. In Algorithmic Game Theory, a representative illustration is the notion and results on the price of anarchy and the smoothness argument method [34]. In a game, the price of anarchy (PoA) [19] 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 of a game — in the same sense of approximation ratio in Approximation Algorithms and competitive ratio in Online Algorithms.

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 primal-dual and dual-fitting in Approximation Algorithms [45] and online primal-dual framework [8] in Online Algorithms. Given the similarity of the notions of PoA, approximation and competitive ratios, it is intriguing and also desirable to develop a method based on duality to characterize the PoA of games. In this paper, we present and aim at developing a framework based on linear programming duality to study the efficiency of games.

1.1 A primal-dual approach

In high-level, the approach follows the standard primal-dual or dual-fitting techniques in approximation/online algorithms. The approach consists of associating a game to an underlying optimization problem and formulate an integer program corresponding to the optimization problem. Next consider the linear program by relaxing the integer constraints and its dual LP. Note that until this step, no notion of game has been intervened. 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. The PoA is then bounded by the ratio between the primal objective (essentially, the cost of the Nash equilibrium) and the dual objective (a lower bound of the optimum cost by weak duality). This approach has been considered by Kulkarni and Mirrokni [21] for full-information games with convex objectives.

There are two crucial steps in the approach. First, by this method, the bound of PoA is at least as large as the integrality gap of the formulation. Hence, to prove optimal PoA one has to derive a formulation (of the corresponding optimization problem) whose the integrality gap matches to the optimal PoA. This is very similar to the issue of linear-programming-based approaches in Approximation/Online Algorithms. Note that this issue is a main obstacle in [21] in order to study non-convex objectives (see discussion in Section 1.3). The second crucial step is the construction of dual variables. The dual variables need to reflect the notion of Nash equilibria as well as their properties in order to relate to the costs of equilibria. Intuitively, to prove optimal bound on the PoA, the constructed dual variables must constitute an optimal dual solution.

To overcome these obstacles, in the paper we systematically consider *configuration* linear programs and a primal-dual approach. Given a problem (game), we first consider a natural formulation of the problem. Then, the approach consists of introducing exponential variables and constraints to the natural formulation to get a configuration LP. The additional constraints have intuitive and simple interpretations: one constraint guarantees that the game admits exactly one outcome and the other constraint ensures that if a player uses a strategy then this strategy must be a component of the outcome. As the result, the configuration LPs significantly improve the integrality gap over that of the natural formulations.

The configuration LPs have been considered in approximation algorithms and to the best of our knowledge, the main approach is rounding. Here, to study the efficiency of games, we consider a primal-dual approach. The primal-dual approach is very appropriate to study the PoA through the mean of configuration LPs. In the dual of the configuration programs, the dual constraints naturally lead to the construction of dual variables and the PoA bounds. Intuitively, one dual constraint corresponds exactly to the definition of Nash equilibrium and the other dual constraint settles the PoA bounds. Note that our approach gives stronger formulations and leads to more general results than that in [21] (see Section 1.3 for a discussion in more details).

1.2 Overview of Results

We illustrate 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 very much in the sense of primal-dual methods in designing algorithms. Moreover, under the lens of LP duality, the notion of smooth games in both full-information settings [34] and incomplete-information settings [35, 41], the recent notion of no-envy learning [12] 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.

1.2.1 Smooth 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 [34]. Roughgarden [34] has introduced the smoothness framework, which became quickly a standard technique, and showed that every (λ, μ)-smooth game has a PoA of at most $\lambda/(1-\mu)$. Through the duality approach, we show that in terms of techniques to study the PoA for complete information settings, the 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 [34].

Informal Theorem 1 *The primal-dual approach captures the smoothness framework in full-information settings.*

1.2.2 Congestion Games

We consider fundamental classes of *congestion games* in which we revisit and unify results in the atomic, 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 (Informal Theorem 1), 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 [38]. Along the line toward the optimal PoA bound for non-atomic congestion games, the equilibrium characterization by a variational inequality is at the core of the analyses [38, 11, 10]. In our proof, we establish the optimal PoA directly by the means of LP duality. By the LP duality as the unified approach, one can clearly observe that the non-atomic setting is a version of the atomic setting in large games (in the sense of [15]) 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 [24], we manage to recover and extend a resource augmentation result related to non-atomic setting [37].

Informal Theorem 2 *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 [36] has presented a *local* smoothness property, a refinement of the smoothness framework, and proved that every (λ, μ)-local-smooth splittable game has a PoA of $\lambda/(1 - \mu)$. 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 [36]. 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 [36]. Note that the matching lower bound is given in [36] and that holds even for pure equilibria.

Definition 1 *A cost function $\ell : \mathbb{R}^+ \to \mathbb{R}^+$ is* (λ, μ)-dual-smooth *if for every vectors $\boldsymbol{u} = (u_1, \ldots, u_n)$ and $\boldsymbol{v} = (v_1, \ldots, v_n)$,*

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

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.*

Informal Theorem 3 *The price of anarchy of coarse correlated equilibria of a splittable congestion game G is at most $\inf_{(\lambda,\mu)} \lambda/(1 - \mu)$ 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.*

1.2.3 Welfare Maximization

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 [34], has been introduced by Roughgarden [35], Syrgkanis and Tardos [41]. This powerful notion has been widely used to study the PoA of Bayes-Nash equilibria (see the recent survey [39]). 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 4 *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. [14] 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 [14] 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 [14]. From this point of view, the primal-dual approach helps to identify the key steps in settling the PoA bounds.

Informal Theorem 5 ([14]) *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 [12] have very recently introduced *no-envy learning dynamic* — a novel concept of learning in auctions. Note that when players have fractionally sub-additive (XOS) valuations[1], 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 6 ([12]) *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 has a long and rich literature [20] and sequentially selling items leads to complex issues in analyzing PoA. Leme et al. [23], Syrgkanis and Tardos [40] 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 advertizing slots one-by-one in the non-increasing order of click-though-rates (from the most attractive to the least one). At each

[1] A 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 $v(S) = \max_{w^\ell \in \mathcal{W}} \sum_{j \in S} w^\ell_j \; \forall S \subset [m]$. The class XOS is a subset of sub-additive functions and is a superset of sub-modular functions.

step, players submit bid 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 PoA bound for the sponsored search problem.

Informal Theorem 7 *The PoA of sequential second-price auctions for the sponsored search problem is at most 2.*

Note that among all auction formats for the sponsored search problem, the best known PoA guarantee [9] is 2.927 which has been achieved in generalized second price (GSP) auctions. 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 in GSP auctions for the sponsored search problem. Consequently, this result shows that the efficiency of sequential auctions is not necessarily worse than the GSP ones and using primal-dual approach, analyzing sequential auctions is not necessarily harder than analyzing GSP ones neither.

Building upon the resilient 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 [40]. That also answers an question raised in [40] 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 8 *The PoA of sequential first-price auctions for the matching market problem is at most 2.*

1.3 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 [21]. In their approach, Kulkarni and Mirrokni [21] 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 [21] 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 [21] 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 [21] needs convex objective functions. In term of approaches based on mathematical programs in approximation algorithms, we have come up with stronger formulations than those in [21] — 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 [21] dealt only with full-information settings. A question has been raised in a the recent survey [39] is whether the framework in [21] could be extended to incomplete-information settings. Our primal-dual approach tends to answer that question.

The connection between LP duality and the PoA have been previously considered by Nadav and Roughgarden [29] and Bilo [6]. 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). However, this approach encounters also the integrality-gap obtacle 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 [26]. Makarychev and Sviridenko [26] 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 [34]. 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 [34, 5]). Subsequently, Roughgarden and Schoppmann [36] 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 [36] and in opinion formation games [4].

The smoothness framework has been extended to incomplete-information environments by Roughgarden [35], Syrgkanis and Tardos [41]. It has successfully yielded tight PoA bounds for several widely-used auction formats. We recommend the reader to a very recent survey [39] 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. [14] have proved that the PoA is constant while the smooth argument gives only logarithmic guarantees. An interesting open direction, as raised in [39], 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 [43]. Vohra [43] 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.

1.4 Organization of Paper

In the paper, we emphasize on presenting the primal-dual approach as an unified tool. For this purpose, along the paper we show the natural development of our framework from full-information games to incomplete-information ones, we reformulate known results and reveal their primal-dual natures. For improvements in concrete problems, results have been stated in Informal Theorems 3, 7, 8 and can be found in Sections 3.3, 4.3.1, 4.3.2, respectively. The structure of the paper is the following. In Section 2, we revisit smooth games. In Section 3, we consider congestion games. In Section 4, we study the problem of welfare maximization in Bayesian setting. The models, definitions and related work are given in the beginning of each (sub-)section.

2 Smooth Games under the Lens of Duality

In this section, we consider smooth games [34] in the point of view of configuration LPs and duality. 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* $\boldsymbol{s} = (s_1, \ldots, s_n)$. The cost $C_i(\boldsymbol{s})$ of player i is a function of the strategy profile $\boldsymbol{s}$ — the chosen strategies of all players. A *pure Nash equilibrium* is a strategy profile $\boldsymbol{s}$ such that no player can decrease its cost via a unilateral deviation; that is, for every player i and every strategy $s'_i \in \mathcal{S}_i$,

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

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

A *mixed Nash equilibrium* [30] of a game is a product distribution $\boldsymbol{\sigma} = \sigma_1 \times \ldots \times \sigma_n$ where σ_i is a probability distribution over the strategy set of player i such that no player can decrease its expected cost under $\boldsymbol{\sigma}$ via a unilateral deviation:

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

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

A *correlated equilibrium* [1] of a game is a joint probability distribution $\boldsymbol{\sigma}$ over the strategy profile of the game such that

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

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

Finally, a *coarse correlated equilibrium* [28] of a game is a joint probability distribution $\boldsymbol{\sigma}$ over the strategy profile of the game such that

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

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.

The notion of smooth games and robust price of anarchy are given in [34]. A game with a joint cost objective function $C(\boldsymbol{s}) = \sum_{i=1}^{n} C_i(\boldsymbol{s})$ is (λ, μ)-*smooth* if for every two outcomes $\boldsymbol{s}$ and $\boldsymbol{s}^*$,

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

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

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

Theorem 1 ([34]) *For every game G with robust PoA $\rho(G)$, every coarse correlated equilibrium $\boldsymbol{\sigma}$ of G and every strategy profile $\boldsymbol{s}^*$,*

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

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 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.

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

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 $\boldsymbol{\sigma}$, define dual variables as follows:

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

Informally, up to some constant factors depending on λ and μ, α_i is the cost of player i in equilibrium $\boldsymbol{\sigma}$, $-\beta$ stands for the cost of the game in equilibrium $\boldsymbol{\sigma}$ 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 $\boldsymbol{\sigma}$. We 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 $\boldsymbol{s}^*$ be the strategy profile corresponding to configuration A. Note that $\mathbb{E}_{\boldsymbol{s}\sim\boldsymbol{\sigma}}[C_i(\boldsymbol{s}^*)] = C_i(\boldsymbol{s}^*)$. The dual constraint reads

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

which is the definition of (λ,μ)-smoothness by arranging the terms and removing the 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 $\boldsymbol{\sigma}$ and the dual objective of the corresponding dual variables. The cost of equilibrium $\boldsymbol{\sigma}$ is $\mathbb{E}_{\boldsymbol{s}\sim\boldsymbol{\sigma}}[C(\boldsymbol{s})]$ while the dual objective of the constructed dual variables is

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

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

Remark. Having shown in [34], Theorem 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 Congestion Games

3.1 Atomic Congestion Games

Model. Atomic congestion games were defined by Rosenthal [32]. In this section, we consider atomic weighted congestion games, a generalized version of the standard congestion game. In a game, we are given a ground set E of resources, a set of n players with strategy sets $\mathcal{S}_1, \ldots, \mathcal{S}_n \subseteq 2^E$ and weights $w_1, \ldots, w_n$ and a cost function $\ell_e : \mathbb{R}^+ \to \mathbb{R}^+$ 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 $\boldsymbol{s} = (s_1, \ldots, s_n)$ where $s_i \in \mathcal{S}_i$ for each player i, we say that $w_e(\boldsymbol{s}) = \sum_{i:e \in s_i} w_i$ is the *load* induced on e by $\boldsymbol{s}$. The cost of a player i is defined as $C_i(\boldsymbol{s}) = \sum_{e:e \in s_i} w_i \cdot \ell_e(w_e)$ where w_e is the load on resource e induced by profile $\boldsymbol{s}$. The total cost of the game in profile $\boldsymbol{s}$ is $C(\boldsymbol{s}) = \sum_{i=1}^n C_i(\boldsymbol{s}) = \sum_{e:e \in s_i} w_e(\boldsymbol{s}) \cdot \ell_e\big(w_e(\boldsymbol{s})\big)$.

The PoA of atomic congestion games has been a extensively studied topic in algorithmic game theory. Most notably, Roughgarden [34] proved that the smoothness argument gave tight bounds for (unweighted) atomic congestion games. For the weighted setting, Bhawalkar et al. [5] showed that the smoothness framework also gave tight bounds for large classes of congestion games.

In this section, we reprove the upper bound [34, 5] on the PoA in atomic congestion games. The result is proved by the same duality approach described in Section 2, but we keep representing here for the following purposes. 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 term 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 $\ell_e : \mathbb{R}^+ \to \mathbb{R}^+$ for a resource e is (λ, μ)-*resource-smooth* if for every 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 2 ([34, 5]) *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 $z_{eT} = 1$.

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

Dual Variables. Fix parameters λ and μ. Given a coarse correlated equilibrium $\boldsymbol{\sigma}$, define corresponding dual variables as follows.

$$
\alpha_i = \frac{1}{\lambda}\mathbb{E}_{\boldsymbol{s}\sim\boldsymbol{\sigma}}[C_i(\boldsymbol{s})], \quad \beta_e := -\frac{\mu}{\lambda}\mathbb{E}_{\boldsymbol{s}\sim\boldsymbol{\sigma}}\Bigg[\sum_{i: e \in s_i} w_i \ell_e(w_e(\boldsymbol{s}))\Bigg], \quad \gamma_{i,e} = \frac{1}{\lambda}\mathbb{E}_{\boldsymbol{s}\sim\boldsymbol{\sigma}}\big[w_i \cdot \ell_e\big(w_e(\boldsymbol{s}_{-i}) + w_i\big)\big]
$$

where $w_e(\boldsymbol{s}_{-i}) = \sum_{i' \neq i, e \in s_{i'}} w_{i'}$. Informally, up to some constant factors, α_i is the cost of player i in equilibrium $\boldsymbol{\sigma}$, $-\beta_e$ 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

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

The inequality holds since without expectation and by linearity of expectation (and also $\mathbb{E}_{\boldsymbol{s}\sim\boldsymbol{\sigma}}\big[w(T) \cdot \ell_e(w(T))\big] = w(T)\ell_e(w(T))$), 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 $\sum_i \alpha_i = \sum_i \frac{1}{\lambda}\mathbb{E}_{\boldsymbol{s}\sim\boldsymbol{\sigma}}[C_i(\boldsymbol{s})] = \mathbb{E}_{\boldsymbol{s}\sim\boldsymbol{\sigma}}\big[\frac{1}{\lambda}\sum_e w_e(\boldsymbol{s})\ell_e(w_e(\boldsymbol{s}))\big]$. Therefore,

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

Hence, PoA $\leq \lambda/(1-\mu)$. □

3.2 Nonatomic Congestion Games

Model. Non-atomic congestion games were defined by Roughgarden and Tardos [38], motivated by the non-atomic routing games of Wardrop [44] and Beckmann et al. [2] and the congestion games of Rosenthal [32]. 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 $\boldsymbol{f} = (f_{s_{ij}})$ 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 $\boldsymbol{f}$. That is, $f_e := \sum_{i=1}^{n} \sum_{e \in s_{ij}} f_{s_{ij}}$. Each resource has a non-decreasing cost function $\ell_e : \mathbb{R}^+ \to \mathbb{R}^+$. With respect to a strategy distribution $\boldsymbol{f}$, players of type i selecting strategy $s_{ij} \in \mathcal{S}_i$ incurs a cost $C_{s_{ij}}(\boldsymbol{f}) = \sum_{e \in s_{ij}} \ell_e(f_e)$. A strategy distribution $\boldsymbol{f}$ is an *pure equilibrium* if for each player type i and strategy $s_{ij}, s_{ij'} \in \mathcal{S}_i$ with $f_{s_{ij}} > 0$,

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

The more general equilibrium concept such as mixed, correlated and coarse correlated equilibria, are defined similarly as in Section 2. The social cost of a strategy distribution $\boldsymbol{f}$ is

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

For non-atomic congestion games, tight bounds on the PoA for almost all classes of cost function have been given in [38]. 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. [2]. This argument is explained in [11, 10]. Moreover, the connection between smoothness arguments and PoA bounds for non-atomic congestion games was revealed in [11].

3.2.1 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 [7]. Hence, the robust PoA is indeed the PoA of pure Nash equilibrium. However, as we do not use the equilibrium characterization from [2], 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* $\xi(\mathcal{L})$ for $\mathcal{L}$ is defined as

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

Theorem 3 ([38]) *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 $\xi(\mathcal{L})$.*

Proof
Formulation. Denote a finite set of multiples of ϵ as $\{a_0, a_1, \ldots, a_m\}$ 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 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 type i distributes 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.

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

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

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 $\boldsymbol{\sigma}$, define the corresponding dual variables as follows.

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

The dual variables have similar interpretations as previous analysis. Variable α_i is the total cost of resources in a strategy used by player type i in equilibrium $\boldsymbol{\sigma}$ 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 $\boldsymbol{\sigma}$.

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

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

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

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

where in the equalities, we use the definition of dual variables. Note that the term $\big(f_e \ell_e(f_e) + v_e \ell_e(v_e) - v_e \ell_e(f_e)\big) \geq 0$ 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 $\mathbb{E}_{\boldsymbol{f} \sim \boldsymbol{\sigma}}\big[\sum_e f_e \ell_e(f_e)\big]$. Hence, the ratio between primal and dual is at most

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

which is bounded by $\xi(\mathcal{L})$ where $\mathcal{L}$ is the class of cost functions on resources in the game. ☐

Remark. The proofs of Theorem 2 and Theorem 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 [15], 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 Resource Augmentation in Non-Atomic Congestion Games

Roughgarden and Tardos [37] proved that in every non-atomic selfish routing game, the cost of an equilibrium is upper bounded by that of an optimal solution that routes 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. [24] 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 $(\mathcal{D}_r)$.

$$\min \sum_{e,T_e} w(T_e)\ell_e(w(T_e))z_{e,T_e} \qquad (\mathcal{P}_r)$$

$$\sum_{j,k} a_k x_{ijk} = (1+r)\cdot w_i \quad \forall i$$

$$\sum_{T_e} z_{e,T_e} = 1 \quad \forall e$$

$$\sum_{T_e:(i,k)\in T_e} z_{e,T_e} = \sum_{j:e\in s_{ij}} x_{ijk} \quad \forall (i,k), e$$

$$x_{ijk}, z_{e,T_e} \in \{0,1\} \quad \forall i,j,e,T_e$$

$$\max \sum_i (1+r)w_i\alpha_i + \sum_e \beta_e \qquad (\mathcal{D}_r)$$

$$a_k\alpha_i \leq \sum_{e:e\in s_{ij}} \gamma_{i,k,e} \quad \forall i,k,j$$

$$\beta_e + \sum_{(i,k)\in T_e} \gamma_{i,k,e} \leq w(T_e)\ell_e(w(T_e)) \quad \forall e, T_e$$

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 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 that of an optimal solution in $(G, (1+r)w, \ell)$.*

Proof Given a coarse correlated equilibrium $\boldsymbol{\sigma}$ 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. As the dual constraints of $(\mathcal{D}_r)$ and $(\mathcal{D}_0)$ are the same, the construction in the proof of Theorem 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 $\boldsymbol{\sigma}$, which is $\mathbb{E}_{\boldsymbol{f}\sim\boldsymbol{\sigma}}\left[\sum_e f_e\ell_e(f_e)\right]$. The former is

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

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 in non-decreasing). Hence, we deduce that the objective of $(\mathcal{D}_r)$ is at least r times the cost of equilibrium $\boldsymbol{\sigma}$. □

3.3 Splittable Congestion Games

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

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

The proof of this equilibrium characterization can be found in [17]. Again, the more general concepts of mixed, correlated and coarse correlated equilibria are defined similarly as in Section 2. In the game, the social cost is defined as $C(\boldsymbol{u}) := \sum_{i=1}^{n} C_i(\boldsymbol{u}) = \sum_e u_e \ell_e(u_e)$.

The PoA bounds has been recently established for a large class of cost functions by Roughgarden and Schoppmann [36]. 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 [36], the works of Cominetti et al. [10] and Harks [17] 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. Motivating 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 [36, Section 5]. A class of cost function $\mathcal{L}$ 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 for non-atomic congestion game. Denote a finite set of multiples of ϵ as $\{a_0, a_1, \ldots, a_m\}$ where $a_k = k \cdot \epsilon$ and $m = \max_{i=1}^{n} w_i/\epsilon$. 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 $s_{ij} \in \mathcal{S}_i$ where $e \in s_{ij}$. Intuitively, a configuration of a resource is a strategy profile of a game restricted on the resource. Let x_{ijk} be variable indicating whether player i distributes an amount a_k of its weight to strategy $s_{ij} \in S_i$. 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 $s_{ij} \in S_i$ 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.

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

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

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 2 *A cost function* $\ell : \mathbb{R}^+ \to \mathbb{R}^+$ *is* (λ, μ)-dual-smooth *if for every vectors* $\boldsymbol{u} =$

$(u_1, \ldots, u_n)$ *and* $\boldsymbol{v} = (v_1, \ldots, v_n)$*,*

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

where $u = \sum_i u_i$ *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 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 $\boldsymbol{\sigma}$, define corresponding dual variables as follows.

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

The dual variables have similar interpretations as 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 players i' other than i follows 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

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

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

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

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

4 Efficiency in Welfare Maximization

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 $\boldsymbol{a} = (a_1, \ldots, a_n) \in \mathcal{A} = \mathcal{A}_1 \times \ldots \times \mathcal{A}_n$, the auctioneer decides an outcome $o(\boldsymbol{a})$ among the set of feasible outcomes $\mathcal{O}$. Each player i has a *private 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 $o(\boldsymbol{a})$ of the game is determined by the action profile $\boldsymbol{a}$, the utility of a player i is denoted as $u_i(\boldsymbol{a}; v_i)$. We are interested in auctions that in general consist of an allocation rule and a payment rule. Given an action profile $\boldsymbol{a} = (a_1, \ldots, a_n)$, the auctioneer decides an allocation and a payment $p_i(\boldsymbol{a})$ for each player i. Then, the *utility* of player i with valuation v_i, following the quasi-linear utility model, is defined as $u_i(\boldsymbol{a}; v_i) = v_i - p_i(\boldsymbol{a})$. The *social welfare* of an auction is defined as the total utility of all participants (the players and the auctioneer): $\textsc{Sw}(\boldsymbol{a}; \boldsymbol{v}) = \sum_{i=1}^n u_i(\boldsymbol{a}; v_i) + \sum_{i=1}^n p_i(\boldsymbol{a})$.

In the paper, we consider incomplete-information settings. In contrast to the full-information settings where private valuations are deterministically determined, in incomplete-informations settings the valuation vectors $\boldsymbol{v}$ (in which each component is the valuation of a player) is drawn from a publicly known distribution $\boldsymbol{F}$ with density function $\boldsymbol{f}$. 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 $\sigma_i : \mathcal{V}_i \to \Delta(\mathcal{A}_i)$ from a valuation $v_i \in \mathcal{V}_i$ to a distribution over actions $\sigma_i(v_i) \in \Delta(\mathcal{A}_i)$.

Definition 3 (Bayes-Nash equilibrium) *A strategy profile* $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ *is a* Bayes-Nash equilibrium (BNE) *if for every player i, for every valuation* $v_i \in \mathcal{V}_i$*, and for every action* $a'_i \in \mathcal{A}_i$*:*

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

For a vector $\boldsymbol{w}$, we use $\boldsymbol{w}_{-i}$ to denote the vector $\boldsymbol{w}$ with the i-th component removed. Besides, $\boldsymbol{F}_{-i}(v_i)$ 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

$$\inf_{\boldsymbol{F}, \boldsymbol{\sigma}} \frac{\mathbb{E}_{\boldsymbol{v} \sim \boldsymbol{F}} \left[\mathbb{E}_{\boldsymbol{a} \sim \boldsymbol{\sigma}(\boldsymbol{v})} [\textsc{Sw}(\boldsymbol{a}; \boldsymbol{v})] \right]}{\mathbb{E}_{\boldsymbol{v} \sim \boldsymbol{F}} \left[\textsc{Opt}(\boldsymbol{v}) \right]}$$

where the infimum is taken over Bayes-Nash equilibria $\boldsymbol{\sigma}$ and $\textsc{Opt}(\boldsymbol{v})$ is the optimal welfare with valuation profile $\boldsymbol{v}$.

In the paper, we consider discrete settings 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.

4.1 Smooth Auctions

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. Smooth auctions have been defined by Roughgarden [35] and Syrgkanis and Tardos [41]. The definitions are slightly different but both are inspired by the original smoothness argument [34] and all known smoothness-based proofs can be equivalently analyzed by one of these definitions. In this section, we consider the definition of smooth auctions in [35] 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 [41].

Definition 4 ([35]) *For parameters $\lambda, \mu \geq 0$, an auction is (λ, μ)-*smooth *if for every valuation profile $\boldsymbol{v} = (v_1, \ldots, v_n)$, there exists action distribution $D_1^*(\boldsymbol{v}), \ldots, D_n^*(\boldsymbol{v})$ over $\mathcal{A}_1, \ldots, \mathcal{A}_n$ such that, for every action profile $\boldsymbol{a}$,*

$$\sum_i \mathbb{E}_{a_i^* \sim D_i^*(\boldsymbol{v})}\big[u_i(a_i^*, \boldsymbol{a}_{-i}; v_i)\big] \geq \lambda \cdot \mathrm{Sw}(\boldsymbol{a}^*; \boldsymbol{v}) - \mu \cdot \mathrm{Sw}(\boldsymbol{a}; \boldsymbol{v}) \tag{1}$$

Theorem 6 ([35]) *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) such that $(i, a_i) \in A$ means 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 $\boldsymbol{v}$, let $z_A(\boldsymbol{v})$ be the variable indicating the probability that the chosen configuration (action profile) is A. For each configuration A and valuation profile $\boldsymbol{v}$, the auctioneer outcomes an allocation and a payment and that results in a social welfare denoted as $c_A(\boldsymbol{v})$. In the other words, if $\boldsymbol{a}$ is the action profile corresponding to the configuration A then $c_A(\boldsymbol{v})$ is in fact $\mathrm{Sw}(\boldsymbol{a}; \boldsymbol{v})$. Consider the following formulation and its dual.

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

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

In the primal, 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 $(v_i, \boldsymbol{v}_{-i})$, the probability that the configuration A contains (i, a_i) must be $f_{-i}(\boldsymbol{v}_{-i}) \cdot x_{i,a_i}(v_i)$. The primal objective is the expected welfare of the auction.

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

$$\alpha_i(v_i) := \frac{1}{\lambda}\mathbb{E}_{\boldsymbol{v}_{-i}}\big[\mathbb{E}_{\boldsymbol{b}\sim\boldsymbol{\sigma}(v_i,\boldsymbol{v}_{-i})}[u_i(\boldsymbol{b}; v_i)]\big],$$
$$\beta(\boldsymbol{v}) := \frac{\mu}{\lambda}\mathbb{E}_{\boldsymbol{b}\sim\boldsymbol{\sigma}(\boldsymbol{v})}\big[\mathrm{Sw}(\boldsymbol{b}; \boldsymbol{v})\big],$$
$$\gamma_{i,a_i}(\boldsymbol{v}) := \frac{1}{\lambda}\mathbb{E}_{\boldsymbol{b}_{-i}\sim\boldsymbol{\sigma}_{-i}(\boldsymbol{v}_{-i})}[u_i(a_i, \boldsymbol{b}_{-i}; v_i)].$$

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

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

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

This is exactly the definition that $\boldsymbol{\sigma}$ is a Bayes-Nash equilibrium.

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

$$\frac{\mu}{\lambda}\mathbb{E}_{\boldsymbol{b}\sim\boldsymbol{\sigma}(\boldsymbol{v})}\big[\mathrm{Sw}(\boldsymbol{b}; \boldsymbol{v})\big] + \sum_{(i,a_i)\in A} \frac{1}{\lambda}\mathbb{E}_{\boldsymbol{b}_{-i}\sim\boldsymbol{\sigma}_{-i}(\boldsymbol{v}_{-i})}[u_i(a_i, \boldsymbol{b}_{-i}; v_i)] \geq \mathrm{Sw}(\boldsymbol{a}; \boldsymbol{v}). \tag{2}$$

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

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

Hence, by taking expectation over $\boldsymbol{\sigma}(\boldsymbol{v})$, Inequality (2) follows.

Price of Anarchy. The welfare of equilibrium $\boldsymbol{\sigma}$ is $\mathbb{E}_{\boldsymbol{v}}\mathbb{E}_{\boldsymbol{b}\sim\boldsymbol{\sigma}(\boldsymbol{v})}\big[\mathrm{Sw}(\boldsymbol{b}; \boldsymbol{v})\big]$ while the dual objective of the constructed dual variables is

$$\begin{aligned}&\sum_{i,v_i} f_i(v_i) \cdot \frac{1}{\lambda}\mathbb{E}_{\boldsymbol{v}_{-i}}\big[\mathbb{E}_{\boldsymbol{b}\sim\boldsymbol{\sigma}(\boldsymbol{v})}[u_i(\boldsymbol{b}; v_i)]\big] + \sum_{\boldsymbol{v}} f(\boldsymbol{v}) \cdot \frac{\mu}{\lambda}\mathbb{E}_{\boldsymbol{b}\sim\boldsymbol{\sigma}(\boldsymbol{v})}\big[\mathrm{Sw}(\boldsymbol{b}; \boldsymbol{v})\big] \\ &\quad \leq \frac{1+\mu}{\lambda} \cdot \mathbb{E}_{\boldsymbol{v}}\mathbb{E}_{\boldsymbol{b}\sim\boldsymbol{\sigma}(\boldsymbol{v})}\big[\mathrm{Sw}(\boldsymbol{b}; \boldsymbol{v})\big]\end{aligned}$$

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

Remark. The notion of (λ, μ)-smooth auctions due to Syrgkanis and Tardos [41] is defined similarly as Definition 4 but now the parameter $\mu \geq 1$ and Inequality (1) is replaced by the following inequality:

$$\sum_i \mathbb{E}_{a_i^* \sim D_i^*(\boldsymbol{v})}\big[u_i(a_i^*, \boldsymbol{a}_{-i}; v_i)\big] \geq \lambda \cdot \textsc{Opt}(\boldsymbol{v}) - \mu \cdot \mathrm{R}(\boldsymbol{a}) \tag{3}$$

where $\mathrm{R}(\boldsymbol{a})$ is the total payment of players if the action profile is $\boldsymbol{a}$. Note that, in order to bound the price of anarchy, Inequality (3) can be replaced by a weaker one, which is:

$$\sum_i \mathbb{E}_{a_i^* \sim D_i^*(\boldsymbol{v})}\big[u_i(a_i^*, \boldsymbol{a}_{-i}; v_i)\big] \geq \lambda \cdot \mathrm{Sw}(\boldsymbol{a}^*; \boldsymbol{v}) - \mu \cdot \mathrm{R}(\boldsymbol{a}) \tag{4}$$

Using the same proof structure of Theorem 6, we can prove that the price of anarchy is at most λ/μ [41]. Specifically, define dual variables α and γ as previous and

$$\beta(\boldsymbol{v}) = \frac{\mu}{\lambda} \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{\sigma}(\boldsymbol{v})}\big[\mathrm{R}(\boldsymbol{b})\big]$$

The feasibility follows the definitions of Bayes-Nash equilibria and smooth auctions, in particular Inequality (4). To bound the price of anarchy, as $\mu \geq 1$, we have

$$\sum_{i,v_i} f_i(v_i) \cdot \frac{1}{\lambda} \mathbb{E}_{\boldsymbol{v}_{-i}}\big[\mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{\sigma}(\boldsymbol{v})}[u_i(\boldsymbol{b}; v_i)]\big] + \sum_{\boldsymbol{v}} f(\boldsymbol{v}) \cdot \frac{\mu}{\lambda} \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{\sigma}(\boldsymbol{v})}\big[\mathrm{R}(\boldsymbol{b})\big] \leq \frac{\mu}{\lambda} \cdot \mathbb{E}_{\boldsymbol{v}} \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{\sigma}(\boldsymbol{v})}\big[\mathrm{Sw}(\boldsymbol{b}; \boldsymbol{v})\big]$$

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

4.2 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 $v_i : 2^{[m]} \to \mathbb{R}^+$ over different subsets of items $S \subset 2^{[m]}$. For simplicity, we denote $v_i(S)$ as v_{iS}. The valuation profile $\boldsymbol{v} = (v_1, \ldots, v_n)$ is drawn from a *product* distribution $\boldsymbol{F}$. 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 *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.

4.2.1 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 every 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. [18] and Bhawalkar and Roughgarden [3], respectively. Recently, Feldman et al. [14] 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. [14] and show that the duality approach captures the non-smooth technique in [14].

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

$$x_{ij}(v_i) := \mathbb{E}_{\boldsymbol{v}_{-i}\sim \boldsymbol{F}_{-i}}\big[\overline{x}_{ij}(v_i, \boldsymbol{v}_{-i})\big], \qquad z_{iS}(v_i) := \mathbb{E}_{\boldsymbol{v}_{-i}\sim \boldsymbol{F}_{-i}}\big[\overline{z}_{iS}(v_i, \boldsymbol{v}_{-i})\big]$$

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

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

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

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

$$\alpha_i(v_i) := 2f_i(v_i)\cdot \mathbb{E}_{\boldsymbol{v}_{-i}\sim \boldsymbol{F}_{-i}}\Big[\mathbb{E}_{\boldsymbol{\sigma}}\big[u_i\big(\boldsymbol{\sigma}(v_i, \boldsymbol{v}_{-i}), v_i\big)\big]\Big] = 2f_i(v_i)\cdot \mathbb{E}_{\boldsymbol{b}\sim \boldsymbol{B}(v_i)}\left[u_i\big(\boldsymbol{b}, v_i\big)\right]$$

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,

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

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

The following lemma shows the feasibility of the variables. The main core of the proof relies on an argument in [14].

Lemma 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 [14] (or see [33, Lemma 1.3] for another clear exposition), for any set of items S, there exists an action b_i^* such that

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

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

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

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

Theorem 7 ([14]) *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 $i^*(j) \in \arg\max_i \mathbb{E}_{\boldsymbol{v}_{-i}\sim\boldsymbol{F}_{-i}}\left[\max_{k\neq i} b_{kj}\right]$. Hence,

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

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

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

Fix a random choice of profile $\boldsymbol{v}$ and $\boldsymbol{\sigma}$ (so the bid profile $\boldsymbol{b}$ 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. □

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 $2(1+r)$.

4.2.2 Connection between Primal-Dual and No-Envy Learning

Very recently, Daskalakis and Syrgkanis [12] 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 [12] 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 [12]. We first define the *online learning problem*. In the online learning problem, at each step t, the player chooses a bid vector $b^t = (b^t_1, \dots, b^t_m)$ where b^t_j 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 $\theta^t = (\theta^t_1, \dots, \theta^t_m)$. The player wins the set $S^*(b^t, \theta^t) = \{j : b^t_j \geq \theta^t_j\}$ and gets reward:

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

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

Definition 5 ([12]) *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:*

$$\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\big[u(b^t, \theta^t)\big] \geq \max_{S \subset [m]} \left(\frac{1}{r} \cdot v(S) - \sum_{j \in S} \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\big[\theta^t_j\big] \right) - \epsilon(T) \tag{5}$$

where the no-envy rate $\epsilon(T) \to 0$ *while* $T \to \infty$*. 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 8 ([12]) *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 $\frac{1}{2r}\textsc{Opt} - n \cdot \epsilon(T)$ where $\textsc{Opt}$ is the expected optimal welfare.*

Proof Let b^t_i be the bid vector of player i where b^t_{ij} is the bid of player i on item j in step t. In a S2A the threshold $\theta^t_{ij} = \max_{k \neq i} b^t_{kj}$. Consider the same primal and dual LPs in Section 4.2.1.

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 4.2.1) as follows.

$$\alpha_i(v_i) := r \cdot f_i(v_i) \cdot \mathbb{E}_{\boldsymbol{v}_{-i} \sim \boldsymbol{F}_{-i}} \left[\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}_{\boldsymbol{b}^t(v_i, \boldsymbol{v}_{-i})} \left[u_i\big(b_i^t, \theta_i^t\big) \right] \right] + r \cdot \epsilon(T)$$

$$\gamma_{i,j}(v_i) := r \cdot f_i(v_i) \cdot \mathbb{E}_{\boldsymbol{v}_{-i} \sim \boldsymbol{F}_{-i}} \left[\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}_{\boldsymbol{b}^t(v_i, \boldsymbol{v}_{-i})} \left[\theta_{ij}^t \right] \right] = r \cdot f_i(v_i) \cdot \mathbb{E}_{\boldsymbol{v}_{-i} \sim \boldsymbol{F}_{-i}} \left[\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}_{\boldsymbol{b}_{-i}^t(\boldsymbol{v}_{-i})} \left[\theta_{ij}^t \right] \right]$$

$$\beta_j := r \cdot \max_i \max_{v_i} \mathbb{E}_{\boldsymbol{v}_{-i} \sim \boldsymbol{F}_{-i}} \left[\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}_{\boldsymbol{b}^t(v_i, \boldsymbol{v}_{-i})} \left[\theta_{ij}^t \right] \right] = r \cdot \max_i \mathbb{E}_{\boldsymbol{v}_{-i} \sim \boldsymbol{F}_{-i}} \left[\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}_{\boldsymbol{b}_{-i}^t(\boldsymbol{v}_{-i})} \left[\theta_{ij}^t \right] \right]$$

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

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

This inequality follows immediately from the definition of r-approximate no-envy learning algorithms (specifically, Inequality (5)) by simplifying and rearranging terms. (Note that $\mathbb{E}_{\boldsymbol{v}_{-i} \sim \boldsymbol{F}_{-i}}[f_i(v_i) \cdot v_{iS}] = f_i(v_i) \cdot v_{iS}$).

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

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

Besides, in the dual objective,

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

$$\sum_{j} \beta_j \leq r \cdot \mathbb{E}_{\boldsymbol{v}} \left[\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}_{\boldsymbol{b}^t(\boldsymbol{v})} \big[v_i\big(S^*(b_i^t, \theta_i^t)\big) \big] \right]$$

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

4.3 Sequential Auctions

4.3.1 Sequential Second Price Auctions in Sponsored Search

Model. In the sponsored search problem, there are n players and n slots. Each player i has a *private valuation* v_i, representing its valuation per click. We use $\boldsymbol{v} = (v_1, \ldots, v_n)$ 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* $\beta_1 \geq \beta_2 \geq \ldots \geq \beta_n$. 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 [16, 31] 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 $\boldsymbol{F}$ 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 polices 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. [27] from optimization perspective and was proposed simultaneously by Edelman et al. [13] and Varian [42] from game theoretical viewpoint (see [22, 25] for surveys on the topic). Recently, Caragiannis et al. [9] have proved the PoA 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 [34]. The study of PoA in sequential auctions has been initiated by Leme et al. [23]. 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 [40] extended the results to incomplete-informations settings and gave constant bounds for both auctions. Leme et al. [23], Syrgkanis and Tardos [40] 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 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 and its dual. 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.

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

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

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

Dual Variables. For fixed (v_i, α_i), denote $\boldsymbol{B}(v_i, \alpha_i)$ the distribution of the equilibrium bid $\boldsymbol{b} = \boldsymbol{\sigma}\big((v_i, \boldsymbol{v}_{-i}), (\alpha_i, \boldsymbol{\alpha}_{-i})\big)$. Recall that $\boldsymbol{b} = (b_1, \ldots, b_n)$ 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 $\boldsymbol{B}_{-i}$ the distribution of the equilibrium bid $\boldsymbol{b}_{-i} = \boldsymbol{\sigma}_{-i}\big((v_i, \boldsymbol{v}_{-i}), (\alpha_i, \boldsymbol{\alpha}_{-i})\big) = \boldsymbol{\sigma}_{-i}\big(\boldsymbol{v}_{-i}, \boldsymbol{\alpha}_{-i}\big)$ where the last equality is due to the independence of distributions F_i's. Define the dual variables as follows.

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

Note that $\pi^{-1}(\boldsymbol{b}_{-i}, j)$ is the winner of slot j in the round selling slot j assuming that player i do 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:

$$\mathbb{E}_{\boldsymbol{b}\sim \boldsymbol{B}(v_i,\alpha_i)} \left[\beta_{\pi(\boldsymbol{b},i)} \cdot \alpha_i v_i + \beta_j \cdot \alpha_{\pi^{-1}(\boldsymbol{b}_{-i},j)} b_{\pi^{-1}(\boldsymbol{b}_{-i},j),j}\right] \geq \beta_j \cdot \alpha_i v_i \tag{6}$$

We prove this inequality through a choice of a hypothetical deviation of player i and use the assumption that $\boldsymbol{\sigma}$ is a Bayes-Nash equilibrium. We first make some observations. Consider a fixed valuation profile $\boldsymbol{v}_{-i}$, a fixed quality factor profile $\boldsymbol{\alpha}_{-i}$ and a realization of (mixed) equilibrium $\boldsymbol{\sigma}\big((v_i, \boldsymbol{v}_i), (\alpha_i, \boldsymbol{\alpha}_{-i})\big)$, denoted as $\boldsymbol{b} = (b_1, \ldots, b_n)$. 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 $\boldsymbol{\sigma}$ 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 (6) 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 to any slot then submit v_i. As $\boldsymbol{\sigma}$ is a Bayes-Nash equilibrium, the utility of player i is at least that induced by this deviation. Specifically,

$$\mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i,\alpha_i)} \big[u_i(\boldsymbol{b}) \big] \geq \mathbb{E}_{\boldsymbol{b}_{-i} \sim \boldsymbol{B}_{-i}} \mathbb{E}_{b'_i} \Big[u_i(b'_i, \boldsymbol{b}_{-i}) \Big]$$

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

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

$$\mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i,\alpha_i)} \Big[u_i(\boldsymbol{b}) \big| \text{Case 2 or Case 3} \Big] \geq \mathbb{E}_{\boldsymbol{b}_{-i} \sim \boldsymbol{B}_{-i}} \mathbb{E}_{b'_i} \Big[u_i(b'_i, \boldsymbol{b}_{-i}) \big| \text{Case 2 or Case 3} \Big] \tag{7}$$

Note that if Case 3 holds then player i gets slot j with the payment $\frac{\alpha_{\pi^{-1}(\boldsymbol{b}_{-i},j)} b_{\pi^{-1}(\boldsymbol{b}_{-i},j),j}}{\alpha_i}$. So

$$\mathbb{E}_{\boldsymbol{b}_{-i} \sim \boldsymbol{B}_{-i}} \mathbb{E}_{b'_i} \Big[u_i(b'_i, \boldsymbol{b}_{-i}) \big| \text{Case 3} \Big] = \mathbb{E}_{\boldsymbol{b}_{-i} \sim \boldsymbol{B}_{-i}} \left[\beta_j \cdot \alpha_i \left(v_i - \frac{\alpha_{\pi^{-1}(\boldsymbol{b}_{-i},j)} b_{\pi^{-1}(\boldsymbol{b}_{-i},j),j}}{\alpha_i} \right) \Big| \text{Case 3} \right] \tag{8}$$

We are now ready to prove the inequality (6). We have

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i,\alpha_i)} \left[\beta_{\pi(\boldsymbol{b},i)} \cdot \alpha_i v_i + \beta_j \cdot \alpha_{\pi^{-1}(\boldsymbol{b}_{-i},j)} b_{\pi^{-1}(\boldsymbol{b}_{-i},j),j} \right] \\
&= \sum_{\ell=1,2,3} \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i,\alpha_i)} \left[\beta_{\pi(\boldsymbol{b},i)} \cdot \alpha_i v_i + \beta_j \cdot \alpha_{\pi^{-1}(\boldsymbol{b}_{-i},j)} b_{\pi^{-1}(\boldsymbol{b}_{-i},j),j} \big| \text{Case } \ell \right] \\
&\geq \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i,\alpha_i)} \left[\beta_j \cdot \alpha_i v_i \big| \text{Case 1} \right] + \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i,\alpha_i)} \left[\beta_{\pi(\boldsymbol{b},i)} \cdot \alpha_i v_i + \beta_j \cdot \alpha_{\pi^{-1}(\boldsymbol{b}_{-i},j)} b_{\pi^{-1}(\boldsymbol{b}_{-i},j),j} \big| \text{Case 2 or 3} \right] \\
&\geq \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i,\alpha_i)} \left[\beta_j \cdot \alpha_i v_i \big| \text{Case 1} \right] + \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i,\alpha_i)} \left[u_i(\boldsymbol{b}) + \beta_j \cdot \alpha_{\pi^{-1}(\boldsymbol{b}_{-i},j)} b_{\pi^{-1}(\boldsymbol{b}_{-i},j),j} \big| \text{Case 2 or 3} \right] \\
&\geq \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i,\alpha_i)} \left[\beta_j \cdot \alpha_i v_i \big| \text{Case 1} \right] + \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i,\alpha_i)} \left[u_i\big((b'_i, \boldsymbol{b}_{-i})\big) + \beta_j \cdot \alpha_{\pi^{-1}(\boldsymbol{b}_{-i},j)} b_{\pi^{-1}(\boldsymbol{b}_{-i},j),j} \big| \text{Case 2 or 3} \right] \\
&\geq \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i,\alpha_i)} \left[\beta_j \cdot \alpha_i v_i \big| \text{Case 1 or 2} \right] + \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i,\alpha_i)} \left[u_i\big((b'_i, \boldsymbol{b}_{-i})\big) + \beta_j \cdot \alpha_{\pi^{-1}(\boldsymbol{b}_{-i},j)} b_{\pi^{-1}(\boldsymbol{b}_{-i},j),j} \big| \text{Case 3} \right] \\
&\geq \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i,\alpha_i)} \left[\beta_j \cdot \alpha_i v_i \big| \text{Case 1 or 2} \right] \\
&\qquad + \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i,\alpha_i)} \left[\beta_j \cdot \alpha_i v_i - \beta_j \cdot \alpha_{\pi^{-1}(\boldsymbol{b}_{-i},j)} b_{\pi^{-1}(\boldsymbol{b}_{-i},j),j} + \beta_j \cdot \alpha_{\pi^{-1}(\boldsymbol{b}_{-i},j)} b_{\pi^{-1}(\boldsymbol{b}_{-i},j),j} \big| \text{Case 3} \right] \\
&= \mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{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 follows the assumption of Case 1: $\beta_{\pi(i)} \geq \beta_j$. The second inequality holds since the utility $u_i(\boldsymbol{b}) \leq \beta_{\pi(\boldsymbol{b},i)} \cdot \alpha_i v_i$. The third inequality is due to (7). 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 (8). Hence, the constructed dual variables form a dual feasible solution.

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

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

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

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

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

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

Moreover,

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

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

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

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

Remark. The non-overbidding assumption can be relaxed in the same way as the remark in Section 4.2.1. 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 $(1+r)$.

4.3.2 Sequential First Price Auctions in Matching Markets

Model. In the matching market problem, there are n players and m items. Each player i has *private unit-demand valuation* $v_i : 2^{[m]} \to \mathbb{R}$ 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 every $j < j'$ 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. [23] proved that the sequential auctions via the second price mechanisms may lead to unbounded inefficiency. The authors [23] 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 [40] 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 [23, 40] and the primal-dual approach enables the improvement. The proof follows similar structure as the one in Section 4.3.1; 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 the player 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 an item is assigned to at most one player, respectively.

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

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

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

Dual Variables. For a fixed valuation v_i, denote $\boldsymbol{B}(v_i)$ the distribution of the equilibrium bid $\boldsymbol{b} = \boldsymbol{\sigma}(v_i, \boldsymbol{v}_{-i})$. Recall that $\boldsymbol{b} = (b_1, \ldots, b_n)$ 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 $\boldsymbol{B}_{-i}$ the distribution of the equilibrium bid $\boldsymbol{b}_{-i} = \boldsymbol{\sigma}_{-i}(v_i, \boldsymbol{v}_{-i}) = \boldsymbol{\sigma}_{-i}(\boldsymbol{v}_{-i})$ where the last equality is due to the independence of distributions. Define the dual variables as follows.

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

Note that $\pi^{-1}(\boldsymbol{b}_{-i}, j)$ 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:

$$\mathbb{E}_{\boldsymbol{b} \sim \boldsymbol{B}(v_i)} \left[v_{i,\pi(\boldsymbol{b},i)} + \sum_{j \in S} b_{\pi^{-1}(\boldsymbol{b}_{-i},j),j} \right] \geq v_{iS} \tag{9}$$

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 $\boldsymbol{v}_{-i}$ and a realization of (mixed) equilibrium $\boldsymbol{\sigma}(v_i, \boldsymbol{v}_i)$, denoted as $\boldsymbol{b} = (b_1, \ldots, b_n)$. Now the assignment π of items to players is completely determined. Let $T = \pi(\boldsymbol{b}, i)$. 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 (by 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 of player i. Player i draws a random sample of a valuation profile $\boldsymbol{w}_{-i} \in \boldsymbol{F}_{-i}$ and determine the winning set $T = \pi(\boldsymbol{\sigma}(v_i, \boldsymbol{w}_{-i}), i)$ 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,

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

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

$$\mathbb{E}_{\boldsymbol{v}_{-i} \sim \boldsymbol{B}_{-i}} \mathbb{E}_{\boldsymbol{\sigma}} \left[b_{\pi^{-1}\left(\boldsymbol{\sigma}_{-i}(\boldsymbol{v}_{-i}), j^*(S)\right), j^*(S)} \middle| \text{Case 2 or 3} \right] \geq \mathbb{E}_{\boldsymbol{v}_{-i} \sim \boldsymbol{B}_{-i}} \mathbb{E}_{\boldsymbol{\sigma}} \left[v_{i,j^*(S)} - v_{i,j^*(T)} \middle| \text{Case 2 or 3} \right] \tag{10}$$

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

We are now ready to prove the inequality (9). We have

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{b}\sim\boldsymbol{B}(v_i)}\Big[v_{i,\pi(\boldsymbol{b},i)} + \sum_{j\in S} b_{\pi^{-1}(\boldsymbol{b}_{-i},j),j}\Big] \\
&\qquad \geq \sum_{\ell=1,2,3} \mathbb{E}_{\boldsymbol{b}\sim\boldsymbol{B}(v_i)}\Big[v_{i,\pi(\boldsymbol{b},i)} + b_{\pi^{-1}(\boldsymbol{b}_{-i},j^*(S)),j^*(S)} \big| \text{Case } \ell\Big] \\
&\qquad \geq \mathbb{E}_{\boldsymbol{b}\sim\boldsymbol{B}(v_i)}\big[v_{iS}\big|\text{Case 1}\big] + \mathbb{E}_{\boldsymbol{b}\sim\boldsymbol{B}(v_i)}\big[v_{i,\pi(\boldsymbol{b},i)} + b_{\pi^{-1}(\boldsymbol{b}_{-i},j^*(S)),j^*(S)}\big|\text{Case 2 or 3}\big] \\
&\qquad \geq \mathbb{E}_{\boldsymbol{b}\sim\boldsymbol{B}(v_i)}\big[v_{iS}\big|\text{Case 1}\big] + \mathbb{E}_{\boldsymbol{b}\sim\boldsymbol{B}(v_i)}\big[v_{i,\pi(\boldsymbol{b},i)} + \big(v_{i,j^*(S)} - v_{i,j^*(T)}\big)\big|\text{Case 2 or 3}\big] \\
&\qquad = v_{iS}
\end{aligned}$$

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 (10). Hence, the constructed dual variables form a dual feasible solution.

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

$$\sum_{i,v_i} y_i(v_i) = \sum_i \mathbb{E}_{v_i\sim F_i}\mathbb{E}_{\boldsymbol{b}\sim\boldsymbol{B}(v_i)}\Big[v_{i,\pi(\boldsymbol{b},i)}\Big] = \mathbb{E}_{\boldsymbol{b}}\Big[\sum_i v_{i,\pi(\boldsymbol{b},i)}\Big].$$

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

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

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

$$z_j = \mathbb{E}_{b_{i^*}}\mathbb{E}_{\boldsymbol{b}_{-i^*}\sim\boldsymbol{B}_{-i^*}}\Big[b_{\pi^{-1}(\boldsymbol{b}_{-i^*},j),j}\Big] = \mathbb{E}_{\boldsymbol{b}}\Big[b_{\pi^{-1}(\boldsymbol{b}_{-i^*},j),j}\Big] \leq \mathbb{E}_{\boldsymbol{b}}\big[b_{\pi^{-1}(\boldsymbol{b},j),j}\big]$$

Summing over all items j, we get

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

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. □

5 Conclusion

In the paper, we have presented a primal-dual approach to study the efficiency of games. We have shown the applicability of the approach on a wide variety of settings and have given simple and improved analyses for several problems in settings of different natures. Beyond concrete results, the main point of the paper is to illuminate the potential of the primal-dual approach. In this approach, the PoA-bound analyses now can be done similarly as the analyses of LP-based algorithms in Approximation/Online Algorithms. We hope that linear programming and duality would bring new ideas and techniques, from well-developed domains such as approximation, online algorithms, etc to algorithmic game theory, not only for the analyses and the understanding of current games but also for the design of new games (auctions) and new concepts leading to improved efficiency.

References

[1] Robert J Aumann. Subjectivity and correlation in randomized strategies. *Journal of mathematical Economics*, 1(1):67–96, 1974.

[2] Martin Beckmann, CB McGuire, and Christopher B Winsten. Studies in the economics of transportation. Technical report, 1956.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] Roberto Cominetti, José R Correa, and Nicolás E Stier-Moses. The impact of oligopolistic competition in networks. *Operations Research*, 57(6):1421–1437, 2009.

[11] 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.

[12] 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.

[13] 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.

[14] 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.

[15] 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.

[16] Drew Fudenberg and Jean Tirole. Game theory, 1991.

[17] Tobias Harks. Stackelberg strategies and collusion in network games with splittable flow. *Theory of Computing Systems*, 48(4):781–802, 2011.

[18] 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.

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

[20] Vijay Krishna. *Auction theory.* Academic press, 2009.

[21] 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.

[22] Sébastien Lahaie, David M Pennock, Amin Saberi, and Rakesh V Vohra. Sponsored search auctions. *Algorithmic game theory*, pages 699–716, 2007.

[23] 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.

[24] 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.

[25] 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.

[26] 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.

[27] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *Journal of the ACM (JACM)*, 54(5):22, 2007.

[28] 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.

[29] 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.

[30] John F Nash. Equilibrium points in n-person games. *Proc. Nat. Acad. Sci. USA*, 36(1):48–49, 1950.

[31] Hans Peters. *Game theory: A Multi-leveled approach*. Springer, 2015.

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

[33] Tim Roughgarden. Frontiers in mechanism design. Lecture 17, 2014.

[34] Tim Roughgarden. Intrinsic robustness of the price of anarchy. *Journal of the ACM (JACM)*, 62(5):32, 2015.

[35] Tim Roughgarden. The price of anarchy in games of incomplete information. *ACM Transactions on Economics and Computation*, 3(1):6, 2015.

[36] Tim Roughgarden and Florian Schoppmann. Local smoothness and the price of anarchy in splittable congestion games. *Journal of Economic Theory*, 156:317–342, 2015.

[37] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2):236–259, 2002.

[38] Tim Roughgarden and Éva Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 47(2):389–403, 2004.

[39] Tim Roughgarden, Vasilis Syrgkanis, and Eva Tardos. The price of anarchy in auctions. *Journal of Artificial Intelligence Research*, 59:59–101, 2017.

[40] Vasilis Syrgkanis and Eva Tardos. Bayesian sequential auctions. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 929–944. ACM, 2012.

[41] 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.

[42] Hal R Varian. Position auctions. *international Journal of industrial Organization*, 25(6): 1163–1178, 2007.

[43] Rakesh V Vohra. *Mechanism design: a linear programming approach*, volume 47. Cambridge University Press, 2011.

[44] John Glen Wardrop. Some theoretical aspects of road traffic research. *Proceedings of the institution of civil engineers*, 1(3):325–362, 1952.

[45] David P Williamson and David B Shmoys. *The design of approximation algorithms*. Cambridge university press, 2011.