

Welfare Guarantees for Proportional Allocations

Ioannis Caragiannis^{1,2} · Alexandros A. Voudouris²

Published online: 23 March 2016

© Springer Science+Business Media New York 2016

Abstract According to the proportional allocation mechanism from the network optimization literature, users compete for a divisible resource – such as bandwidth – by submitting bids. The mechanism allocates to each user a fraction of the resource that is proportional to her bid and collects an amount equal to her bid as payment. Since users act as utility-maximizers, this naturally defines a proportional allocation game. Syrgkanis and Tardos (STOC 2013) quantified the inefficiency of equilibria in this game with respect to the social welfare and presented a lower bound of 26.8 % on the price of anarchy over coarse-correlated and Bayes-Nash equilibria in the full and incomplete information settings, respectively. In this paper, we improve this bound to 50 % over both equilibrium concepts. Our analysis is simpler and, furthermore, we argue that it cannot be improved by arguments that do not take the equilibrium structure into account. We also extend it to settings with budget constraints where we show the first constant bound (between 36 and 50 %) on the price of anarchy of the corresponding game with respect to an effective welfare benchmark that takes budgets into account.

A preliminary version of this paper appeared as [3]. This work was partially supported by the European Social Fund and Greek national funds through the research funding program Thales on “Algorithmic Game Theory”, and by the Caratheodory research grant E.114 from the University of Patras.

✉ Ioannis Caragiannis
caragian@ceid.upatras.gr

Alexandros A. Voudouris
voudouris@ceid.upatras.gr

¹ Computer Technology Institute and Press “Diophantus”, University of Patras, 26504 Rio, Greece

² Department of Computer Engineering and Informatics, University of Patras, 26504 Rio, Greece

Keywords Algorithmic game theory · Price of anarchy · Bayes-Nash equilibrium · Proportional allocation mechanism

1 Introduction

The *proportional allocation mechanism*, introduced by Kelly [12], is fundamental in the network optimization literature. According to this mechanism, a divisible resource — such as bandwidth of a communication link — is allocated to users as follows. Each user submits a bid to the mechanism; this corresponds to the user's *willingness-to-pay* for sharing the resource. The mechanism allocates to each user a fraction of the resource that is equal to the ratio of her bid over the total amount of bids. It also receives a payment from each user that is equal to her bid. This naturally defines a *proportional allocation game* among the users who act as players; each player has a (typically concave, non-negative, and non-decreasing) valuation function for the resource share she receives and aims to maximize her utility, i.e., her value for the resource share minus her payment to the mechanism. As it is typically the case in games, the *social welfare* (i.e., the total value of the players for the resource share they receive) at *equilibria* is, in general, suboptimal.

We aim to quantify this inefficiency of equilibria by bounding the *price of anarchy* [13] of proportional allocation games. Besides the well-known work of Johari and Tsitsiklis [10] who considered pure Nash equilibria in the full information setting, there has been surprisingly little focus on price of anarchy bounds over more general equilibrium concepts. The only exception we are aware of is the recent work of Syrgkanis and Tardos [23] who studied proportional allocation as part of a broader class of mechanisms. Motivated by their work, we present new bounds on the price of anarchy of proportional allocation under general equilibrium concepts, such as *coarse-correlated* equilibria in the full information setting and *Bayes-Nash* equilibria in the incomplete information setting. In particular, we prove that the social welfare at equilibrium is at least $1/2$ of the optimal social welfare. The bound holds for coarse-correlated and pure Bayes-Nash equilibria in the full information and Bayesian setting, respectively, and improves the bound of 26.8 % of [23]. The proof is conceptually simple and is obtained by bounding the utility of every player at equilibrium by the utility this player would have by deviating to a particular deterministic bid.

We also consider the scenario where players have budget constraints representing their *ability-to-pay*. Here, each player has a budget and is never allowed to bid above it. We assess the quality of equilibria in this case in terms of an *effective welfare* benchmark — proposed in previous work but further refined here — that takes budgets into account. We show that the effective welfare at equilibrium is at least a constant fraction of the optimal one. To the best of our knowledge, this is the first constant price of anarchy bound (in particular, between 36 and 50 %) with respect to this benchmark. Again, our proofs follow by considering a single deterministic deviation for each player, defined in a slightly different way compared to the deviation we consider in our bound on the social welfare.

Related Work The proportional allocation mechanism and its variations have received significant attention in the network optimization literature. Proportional allocation games have been considered in [8, 14, 15] where the existence and uniqueness conditions for pure Nash equilibria are proved. Variations of the mechanism with different definitions for the allocation rule or the payments have been considered in [16–18, 21] (see also the discussion in [9]).

Johari and Tsitsiklis [10] were the first who assessed the quality of proportional allocations in terms of the social welfare. They focus on pure Nash equilibria and proved a lower bound of $3/4$ on their price of anarchy. Their analysis is based on the important observation that a pure Nash equilibrium in a proportional allocation game is also a pure Nash equilibrium in a game where each player has a *linear* valuation function with slope equal to the derivative of the original valuation function at the share value they get at equilibrium. The optimal social welfare in the new game is not smaller than the original one and this allows them to consider the significantly simpler case of linear valuations in their analysis. Then, the price of anarchy bound is obtained by solving a linear program. An alternative proof to the result of [10] without using this argument is presented in [19] (see also [9]).

Unfortunately, this transformation does not apply to more general equilibrium concepts since the resource share each player receives is, in general, a random variable. This is a rather common difficulty that manifests itself in the analysis of games, as we depart from pure Nash equilibria and full information. In particular, Bayes-Nash equilibria have such an extremely rich structure that, typically, the price of anarchy analysis assesses their quality by rather ignoring this structure. Instead, it resorts to bounding the utility of each player by appropriately selected deviations which reveal a relation between the social welfare at equilibrium and the optimal social welfare. This approach has been used in a series of papers that mostly focus on auctions (e.g., see [1, 2, 4, 7, 11, 20, 23]) and is actually the approach we follow in the current paper as well.

Syrkkanis and Tardos [23] present a general analysis framework for the broad class of smooth mechanisms. Among other results, they show a price of anarchy lower bound of 26.8 % over coarse-correlated and mixed Bayes-Nash equilibria of proportional allocation games. In their analysis, they bound the utility of each player by the utility she would have by deviating to an appropriately defined *randomized* bid (an approach that has also been used in different contexts in [2, 11, 22, 24]) with a probability distribution that depends only on the optimal allocation and the valuation function of the player. In contrast, the deviating bid we consider depends on the bid strategies at equilibrium (this is in the same spirit as the recent analysis of Feldman et al. [7]) and, more interestingly, it is *deterministic*. In particular, it is defined as the product of the (expected) resource share a bidder receives in the optimal allocation and the expectation of bids of the other players at equilibrium.

Budget constraints are well-motivated in auction settings. In a slightly different context than ours, the effective welfare benchmark is considered by Dobzinski and Paes Leme, who call it *liquid* welfare in [6]. In proportional allocation, Syrkkanis and Tardos [23] prove that the social welfare at equilibrium is a constant fraction (specifically, equal to $2 - \sqrt{3} \approx 26.8\%$) of the optimal effective welfare. Note that

our guarantee is considerably stronger as we compare directly the effective welfare at equilibrium with its optimal value.

Roadmap The rest of the paper is structured as follows. We begin with preliminary definitions in Section 2. Our price of anarchy bounds in terms of the social welfare are proved in Section 3. There, we also argue that in order to improve our analysis, radically new ideas are required. The budget-constrained setting is studied in Section 4. We remark that we have not mentioned mixed Bayes-Nash equilibria in the above presentation of our results. Actually, we have observed that such equilibria coincide with pure Bayes-Nash ones even in the budget-constrained setting. We discuss related issues as well as additional open problems in Section 5.

2 Preliminaries

Each player (henceforth called *bidder*) i in a proportional allocation game has a concave¹ non-decreasing valuation function $v_i : [0, 1] \rightarrow \mathbb{R}^+$. A strategy for bidder i is simply a non-negative bid. Given a bid vector $\mathbf{b} = (b_1, b_2, \dots, b_n)$, with one bid per bidder, the proportional allocation mechanism allocates to each bidder a fraction of the resource that is proportional to the bid submitted by her. Denoting by d_i the resource share that is allocated to bidder i , it is $d_i = \frac{b_i}{\sum_j b_j}$. We often use the notation

B_{-i} to denote the sum of bids of all bidders besides i (hence, $d_i = \frac{b_i}{b_i + B_{-i}}$). The utility of bidder i from an allocation is simply the difference of her value for the fraction of the resource she gets minus her bid, i.e., $u_i(\mathbf{b}) = v_i(d_i) - b_i$.

A bid vector \mathbf{b} is a pure Nash equilibrium if the utility of all bidders is maximized, given the bid strategies of the other bidders. So, in a pure Nash equilibrium, no bidder has any incentive to deviate to another strategy. Denoting by (b'_i, \mathbf{b}_{-i}) the bid vector that is obtained from \mathbf{b} when bidder i unilaterally deviates to bid strategy b'_i , we can express this condition as $u_i(\mathbf{b}) \geq u_i(b'_i, \mathbf{b}_{-i})$.

The social welfare of an allocation d is the total value of bidders for the resource shares they receive, i.e., $\text{SW}(d) = \sum_i v_i(d_i)$. We denote by SW^* the maximum value of the social welfare over all possible allocations. The price of anarchy over pure Nash equilibria is defined as the minimum value of the social welfare among all pure Nash equilibria divided by the optimal social welfare.

The bid strategy of a bidder i can be randomized. In this case, b_i is a random variable and the bidder aims to maximize her expected utility $\mathbb{E}[u_i(\mathbf{b})]$. The bid strategies of different bidders can be independent or correlated. A vector of independent randomized bid strategies is called a mixed Nash equilibrium if it simultaneously maximizes the expected utility of each bidder, given the bid strategies of the other bidders. More generally, coarse-correlated equilibria are solution concepts that capture correlated bid strategies. A vector of (possibly correlated) bid strategies is called

¹Correa et al. [5] studied proportional allocation games in the less standard scenario of non-concave valuation functions.

a coarse-correlated equilibrium if no bidder has any incentive to unilaterally deviate to any deterministic bid strategy in order to improve her expected utility (again, given the strategies of the other bidders). The notion of the price of anarchy naturally extends to these solution concepts as well. For example, the price of anarchy over coarse-correlated equilibria is defined as the minimum value of the expected social welfare among all coarse-correlated equilibria divided by the optimal social welfare.

The above setting is known as the full (or complete) information setting. We consider the incomplete information (or Bayesian) setting as well; in this case, the valuation function v_i of each bidder i is drawn randomly (and independently from the other bidders) from a probability distribution F_i over concave, non-decreasing, and non-negative functions in $[0, 1]$. Again, bidder i aims to maximize her expected utility for each possible valuation function v_i drawn from F_i . In the incomplete information setting, each bidder i bases her decision on her exact valuation v_i and on the probability distributions according to which other bidders draw their valuations (and their corresponding bid strategies); these distributions are common knowledge.

So, the bid strategy of bidder i is a (possibly random) bid function $b_i(v_i)$. A vector with one such strategy per bidder (with independence between bid strategies of different bidders) is called a mixed Bayes-Nash equilibrium if no bidder has any incentive to deviate to some other bid for any valuation function drawn from F_i . In pure Bayes-Nash equilibria, bidders use deterministic bid functions. The price of anarchy over Bayes-Nash equilibria is defined as the minimum value of the expected social welfare among all Bayes-Nash equilibria divided by the expectation of the optimal social welfare. With some abuse in notation, we also use SW^* to denote the expectation of the optimal social welfare in the Bayesian setting.

We also extend the above model by adding budget constraints to the bidders. In this setting, each bidder i has a non-negative budget c_i and she is never allowed to bid above her budget. This restriction can result to equilibria that have extremely low social welfare compared to the optimal one (whose definition does not take budgets into account). Following [23] and [6], we use the effective welfare benchmark in order to assess the quality of equilibria with budget-constrained bidders. The effective welfare of a (deterministic) allocation $d = (d_1, d_2, \dots, d_n)$ is defined as $EW(d) = \sum_i \min\{v_i(d_i), c_i\}$. Note that the definition is similar to the definition of the social welfare; the important difference is that the value of each bidder is capped by her budget. We extend this definition to random allocations d as $EW(d) = \sum_i \min\{\mathbb{E}[v_i(d_i)], c_i\}$. We denote by EW^* the maximum value of the effective welfare over all allocations. The price of anarchy with respect to the effective welfare benchmark (over equilibria in a given class) is the minimum value of the effective welfare (among all allocations induced by equilibria in the class) divided by the optimal effective welfare.

In the Bayesian setting, both the budget c_i of bidder i and her valuation v_i are drawn randomly according to the probability distribution F_i . We refine the effective welfare benchmark in this case as

$$EW(d) = \sum_i \mathbb{E}_{(v_i, c_i) \sim F_i} [\min\{\mathbb{E}_{(v_{-i}, c_{-i}) \sim F_{-i}} [v_i(d_i)], c_i\}],$$

where the inner expectation is taken over the valuation-budget value pairs of the other bidders once the pair for bidder i has been fixed (and over the corresponding bid strategies).

3 Bounding the Social Welfare of Equilibria

In this section, we prove the price of anarchy bounds with respect to the social welfare. We consider both coarse-correlated equilibria in the full information setting as well as pure Bayes-Nash equilibria in the Bayesian setting. Our proofs use the following lemma which bounds the utility of a bidder at a deterministic deviation. We also use this lemma later in Section 4 where we study budget-constrained bidders.

Lemma 3.1 *Consider a bidder with a concave and non-decreasing valuation function $v : [0, 1] \rightarrow \mathbb{R}^+$ and let Γ be the random variable denoting the sum of bids of the other bidders. Then, for every $z \in [0, 1]$ and for every $\mu > 0$, the expected utility the bidder would have by deviating to the deterministic bid $\mu z \mathbb{E}[\Gamma]$ is at least $\frac{3\mu-1}{4\mu} v(z) - \mu z \mathbb{E}\Gamma$.*

Proof It suffices to show that the expected value of the bidder when she deviates to the deterministic bid $y = \mu z \mathbb{E}[\Gamma]$ is at least $\frac{3\mu-1}{4\mu} v(z)$. Define the event $T := \left\{ \Gamma \geq y \left(\frac{1}{z} - 1 \right) \right\}$. When T is false, we have $\frac{y}{y+\Gamma} > z$ and, since v is non-decreasing, we clearly have that

$$v\left(\frac{y}{y+\Gamma}\right) \geq v(z).$$

Otherwise, when T is true, $\frac{y}{y+\Gamma} \in [0, z]$. Since v is concave and non-negative, its value in $[0, z]$ is lower-bounded by the line connecting points $(0, 0)$ and $(z, v(z))$. Hence,

$$v\left(\frac{y}{y+\Gamma}\right) \geq \frac{y}{y+\Gamma} \cdot \frac{v(z)}{z}.$$

So, we can bound the expected value of the bidder when she deviates to the deterministic bid y using the two observations above and linearity of expectation.

$$\begin{aligned} \mathbb{E}\left[v\left(\frac{y}{y+\Gamma}\right)\right] &= \mathbb{E}\left[v\left(\frac{y}{y+\Gamma}\right) \mathbb{1}_{\bar{T}}\right] + \mathbb{E}\left[v\left(\frac{y}{y+\Gamma}\right) \mathbb{1}_T\right] \\ &\geq \mathbb{E}\left[v(z) \mathbb{1}_{\bar{T}}\right] + \mathbb{E}\left[\frac{y}{y+\Gamma} \cdot \frac{v(z)}{z} \mathbb{1}_T\right] \\ &= v(z)(1 - \Pr[T]) + \frac{v(z)}{z} \mathbb{E}\left[\frac{y}{y+\Gamma} \mathbb{1}_T\right]. \end{aligned} \quad (1)$$

Here, we have used the notation $X \mathbb{1}_T$ to denote the random variable that is equal to X if T is true and is zero otherwise.

We will now work with the rightmost term of the above right-hand side expression. Since the function $\frac{y}{y+\Gamma}$ is convex with respect to Γ , we can apply Jensen’s inequality to obtain that

$$\mathbb{E} \left[\frac{y}{y + \Gamma} \mathbb{1}T \right] = \mathbb{E} \left[\frac{y}{y + \Gamma} | T \right] \cdot \Pr[T] \geq \frac{y \Pr[T]}{y + \mathbb{E}[\Gamma|T]}$$

and, since $\mathbb{E}[X|T] \leq \frac{\mathbb{E}[X]}{\Pr[T]}$ for every random variable $X \geq 0$, we have

$$\mathbb{E} \left[\frac{y}{y + \Gamma} \mathbb{1}T \right] \geq \frac{y \Pr[T]^2}{y \Pr[T] + \mathbb{E}[\Gamma]} \geq \frac{y \Pr[T]^2}{y + \mathbb{E}[\Gamma]}.$$

The second inequality follows trivially since $\Pr[T] \leq 1$. Substituting y and using the fact that $z \leq 1$, we get

$$\mathbb{E} \left[\frac{y}{y + \Gamma} \mathbb{1}T \right] \geq \frac{\mu z \Pr[T]^2}{1 + \mu z} \geq \frac{\mu z}{\mu + 1} \Pr[T]^2. \tag{2}$$

Now, using (1), (2), and the fact that $1 - \alpha + \frac{\mu}{\mu+1}\alpha^2 \geq \frac{3\mu-1}{4\mu}$ for every α , we obtain that

$$\mathbb{E} \left[v \left(\frac{y}{y + \Gamma} \right) \right] \geq v(z) \left(1 - \Pr[T] + \frac{\mu}{\mu + 1} \Pr[T]^2 \right) \geq \frac{3\mu - 1}{4\mu} v(z),$$

as desired. □

We are ready to prove our price of anarchy bounds. We begin with the case of coarse-correlated equilibria in the full information setting which is much simpler.

Theorem 3.2 *The price of anarchy of proportional allocation games over coarse-correlated equilibria is at least 1/2.*

Proof Consider a full information proportional allocation game with n bidders in which bidder i has valuation function v_i and denote by x_i the resource fraction bidder i gets in the optimal allocation. Let \mathbf{b} be a coarse-correlated equilibrium that induces a random allocation $d = (d_1, \dots, d_n)$ and let $B = \sum_i b_i$ be the random variable denoting the sum of bids of all bidders, with B_{-i} being the sum of bids of all bidders besides bidder i . Since \mathbf{b} is a coarse-correlated equilibrium, bidder i has no incentive to deviate to any deterministic bid (including the deviating bid $x_i \mathbb{E}[B_{-i}]$). By applying Lemma 3.1 for bidder i with $z = x_i$, $\mu = 1$ and $\Gamma = B_{-i}$, we obtain that

$$\mathbb{E}[u_i(\mathbf{b})] \geq \mathbb{E}[u_i(x_i \mathbb{E}[B_{-i}], \mathbf{b}_{-i})] \geq \frac{1}{2} v_i(x_i) - x_i \mathbb{E}[B_{-i}].$$

Summing over all bidders and using the facts that $\sum_i x_i = 1$ and $B_{-i} \leq B$ for every bidder i , we have

$$\begin{aligned} \sum_i \mathbb{E}[u_i(\mathbf{b})] &\geq \frac{1}{2} \sum_i v_i(x_i) - \sum_i x_i \mathbb{E}[B_{-i}] \\ &\geq \frac{1}{2} \sum_i v_i(x_i) - \sum_i x_i \mathbb{E}[B] \\ &= \frac{1}{2} \text{SW}^* - \mathbb{E}[B]. \end{aligned} \tag{3}$$

The theorem follows by this inequality since the social welfare equals the sum of bidders’ utilities plus their bids, i.e., $\mathbb{E}[\text{SW}(d)] = \sum_i \mathbb{E}[u_i(\mathbf{b})] + \mathbb{E}[B]$. \square

The last step of the proof above begins with inequality (3). Essentially, this inequality has the form

$$\sum_i \mathbb{E}[u_i(\mathbf{b})] \geq \lambda \text{SW}^* - \mu \sum_i x_i \mathbb{E}[B_{-i}].$$

The price of anarchy bound of [23] follows after first proving an inequality of this type and then concluding to a price of anarchy bound of $\frac{\lambda}{\max\{1, \mu\}}$. The smoothness arguments of [23] lead to a version of this inequality with $\lambda = 2 - \sqrt{3}$ and $\mu = 1$. Here, we have been able to improve the parameters to $\lambda = 1/2$ and $\mu = 1$. The next lemma demonstrates that these parameters cannot be improved further.

Lemma 3.3 *For every $\epsilon > 0$, there exists a proportional allocation game such that for every λ, μ satisfying*

$$\sum_i u_i(\mathbf{b}) \geq \lambda \text{SW}^* - \mu \sum_i x_i B_{-i} \tag{4}$$

where x_i is the resource fraction of bidder i in the optimal allocation and B_{-i} is the sum of bids of all bidders besides bidder i at a (pure Nash) equilibrium, it holds that $\frac{\lambda}{\max\{1, \mu\}} \leq \frac{1}{2} + \epsilon$.

Proof Consider the proportional allocation game with $n \geq 2$ bidders in which bidder 1 has valuation $v_1(x) = x$ and bidder i has valuation $v_i(x) = \frac{n-1}{2n-3}x$ for $i \geq 2$. We can show that the bids in the (unique) pure Nash equilibrium are $b_1 = 1/4$ and $b_i = \frac{1}{4(n-1)}$ for $i \geq 2$. Indeed, assuming that this is true for all bidders besides i , it can be verified that the strategy y that maximizes the utility $u_i(y, \mathbf{b}_{-i}) = v_i\left(\frac{y}{y+B_{-i}}\right) - y$ for bidder i satisfies $y = b_i$. I.e., bidder 1 gets half of the resource and the remaining

bidders share the remaining resource equally. Hence, $u_1(\mathbf{b}) = v_1(1/2) - 1/4 = 1/4$ and $\sum_{i \neq 1} u_i(\mathbf{b}) = (n - 1) \left(v_i \left(\frac{1}{2(n-1)} \right) - \frac{1}{4(n-1)} \right) = \frac{1}{4(2n-3)}$.

In the optimal allocation, the whole resource is allocated to bidder 1, i.e., $SW^* = 1, x_1 = 1$ and $x_i = 0$ for $i \geq 2$. Hence, inequality (4) becomes

$$\frac{1}{4} + \frac{1}{4(2n - 3)} \geq \lambda - \frac{\mu}{4}$$

and implies that $\lambda \leq \max\{1, \mu\} \left(\frac{1}{2} + \frac{1}{4(2n-3)} \right)$. The lemma follows by setting n sufficiently large. □

The proof for Bayes-Nash equilibria follows the same general approach with that of Theorem 3.2.

Theorem 3.4 *The price of anarchy of proportional allocation games over pure Bayes-Nash equilibria is at least 1/2.*

Proof Consider an incomplete information proportional allocation game in which the valuation function \mathbf{v}_i of bidder i is drawn from the probability distribution \mathbf{F}_i , independently for each bidder. Let x_i be the random variable denoting the resource fraction bidder i gets in the optimal allocation. Let \mathbf{b} be a pure Bayes-Nash equilibrium and B be the random variable denoting the sum of bids of all bidders; again, B_{-i} denotes the sum of bids of all bidders besides bidder i . Since \mathbf{b} is a pure Bayes-Nash equilibrium, bidder i has no incentive to deviate to any deterministic bid (including the deviating bid $\mathbb{E}[x_i|v_i] \mathbb{E}[B_{-i}|v_i]$) when the valuation drawn from probability distribution \mathbf{F}_i is v_i . So, in all conditional expectations below, we simply write v_i to denote the event that the valuation \mathbf{v}_i drawn from \mathbf{F}_i is v_i . By applying Lemma 3.1 for bidder i with $z = \mathbb{E}[x_i|v_i]$, $\mu = 1$ and $\Gamma = B_{-i}$, we obtain that

$$\begin{aligned} \mathbb{E}[u_i(\mathbf{b})|v_i] &\geq \mathbb{E}\left[u_i\left(\mathbb{E}[x_i|v_i] \mathbb{E}[B_{-i}|v_i], \mathbf{b}_{-i}\right)|v_i\right] \\ &\geq \frac{1}{2} v_i (\mathbb{E}[x_i|v_i] - \mathbb{E}[x_i|v_i] \mathbb{E}[B_{-i}|v_i]) \\ &\geq \frac{1}{2} \mathbb{E}[v_i(x_i)|v_i] - \mathbb{E}[x_i|v_i] \mathbb{E}[B]. \end{aligned}$$

The second inequality follows by Jensen’s inequality since the valuation function v_i is concave and due to the fact that in a pure Bayes-Nash equilibrium, the bid of a bidder different than i does not depend on the exact valuation of bidder i and, hence, $\mathbb{E}[B_{-i}|v_i] = \mathbb{E}[B_{-i}] \leq \mathbb{E}[B]$. Considering all possible valuations for bidder i that are drawn from probability distribution \mathbf{F}_i , we have that her unconditional expected utility is

$$\mathbb{E}[u_i(\mathbf{b})] \geq \frac{1}{2} \mathbb{E}[v_i(x_i)] - \mathbb{E}[x_i] \mathbb{E}[B].$$

Summing over all bidders and using the fact that $\sum_i x_i = 1$, we have

$$\sum_i \mathbb{E}[u_i(\mathbf{b})] \geq \frac{1}{2} \sum_i \mathbb{E}[v_i(x_i)] - \sum_i \mathbb{E}[x_i] \mathbb{E}[B] = \frac{1}{2} \text{SW}^* - \mathbb{E}[B].$$

The theorem follows from this inequality since, again, the social welfare equals the sum of expected bidders' utilities plus the total amount of bids. \square

4 Budget-Constrained Bidders

In this section, we consider budget-constrained bidders and prove a lower bound of approximately 36 % and an upper bound of 50 % on the price of anarchy in terms of the effective welfare benchmark (Theorem 4.1 for coarse-correlated equilibria and Theorem 4.2 for Bayes-Nash equilibria). Furthermore, we present an upper bound (Theorem 4.3) which applies even to pure Nash equilibria.

Before proceeding to the presentation of our bounds for budget-constrained bidders, we remark that minor modifications of the proofs in the previous section can show that the social welfare over equilibria with budget-constrained bidders is at least 1/2 of the optimal effective welfare, improving a corresponding bound of 26.8 % from [23]. The necessary modifications are as follows. First, we need to define the deviating bids in terms of the resource shares in the allocation that maximizes the effective welfare. Then, there is a subtle case where Lemma 3.1 cannot be used, namely when the deviating bid for a bidder exceeds her budget. Fortunately, the inequality provided by Lemma 3.1 follows trivially in this case (actually, we use this argument in the proof below). By repeating the analysis in the proofs of Theorems 3.2 and 3.4, we can conclude that the social welfare at equilibrium is at least 1/2 of the social welfare of the allocation that maximizes the effective welfare. The bound then follows by observing that the effective welfare of this allocation is upper-bounded by its social welfare.

We are now ready to prove our main results for budget-constrained bidders. Again, we begin with the case of coarse-correlated equilibria in the full information setting which is simpler.

Theorem 4.1 *The price of anarchy of proportional allocation games with budget-constrained bidders over coarse-correlated equilibria is at least 0.3596.*

Proof Let $\mu \in (1/3, 1]$ be a parameter whose exact value will be defined later. Consider a full information proportional allocation game with n bidders in which bidder i has valuation function v_i and budget c_i and denote by x_i the resource fraction bidder i gets in the allocation that maximizes the effective welfare. Let \mathbf{b} be a coarse-correlated equilibrium inducing an allocation $d = (d_1, \dots, d_n)$ and let $B = \sum_i b_i$ be the random variable denoting the sum of bids of all bidders, with B_{-i} being the sum of bids of all bidders besides bidder i . Let A be the set of bid-

ders with $\mathbb{E}[v_i(d_i)] \leq c_i$. Clearly, for every bidder not belonging to set A , it holds that

$$\min\{\mathbb{E}[v_i(d_i)], c_i\} \geq \min\{v_i(x_i), c_i\}.$$

Summing over all bidders not belonging to A (and multiplying by $1 - \mu$), we obtain that

$$(1 - \mu) \sum_{i \notin A} \min\{\mathbb{E}[v_i(d_i)], c_i\} \geq (1 - \mu) \sum_{i \notin A} \min\{v_i(x_i), c_i\}. \tag{5}$$

For every bidder $i \in A$, we distinguish between two cases. If $\mu x_i \mathbb{E}[B_{-i}] > c_i$, then clearly

$$\begin{aligned} \mathbb{E}[u_i(\mathbf{b})] &\geq 0 \\ &> c_i - \mu x_i \mathbb{E}[B_{-i}] \\ &\geq \frac{3\mu - 1}{4\mu} \min\{\mathbb{E}[v_i(x_i)], c_i\} - \mu x_i \mathbb{E}[B_{-i}]. \end{aligned}$$

In order to prove the same inequality when $\mu x_i \mathbb{E}[B_{-i}] \leq c_i$, we bound the utility of bidder i by the utility she would have when deviating to bid $\mu x_i \mathbb{E}[B_{-i}]$ (which is within i 's budget c_i). Using Lemma 3.1, we have again

$$\mathbb{E}[u_i(\mathbf{b})] \geq \frac{3\mu - 1}{4\mu} \mathbb{E}[v_i(x_i)] - \mu x_i \mathbb{E}[B_{-i}].$$

Summing this last inequality over all bidders of A , we obtain

$$\begin{aligned} \sum_{i \in A} \mathbb{E}[u_i(\mathbf{b})] &\geq \frac{3\mu - 1}{4\mu} \sum_{i \in A} \min\{\mathbb{E}[v_i(x_i)], c_i\} - \sum_{i \in A} \mu x_i \mathbb{E}[B_{-i}] \\ &\geq \frac{3\mu - 1}{4\mu} \sum_{i \in A} \min\{\mathbb{E}[v_i(x_i)], c_i\} - \mu \mathbb{E}[B] \sum_{i \in A} x_i \\ &\geq \frac{3\mu - 1}{4\mu} \sum_{i \in A} \min\{\mathbb{E}[v_i(x_i)], c_i\} - \mu \mathbb{E}[B]. \end{aligned}$$

Using the equality $B = \sum_i b_i$ and linearity of expectation, this inequality implies that

$$\sum_{i \in A} (\mathbb{E}[u_i(\mathbf{b})] + \mu \mathbb{E}[b_i]) + \mu \sum_{i \notin A} \mathbb{E}[b_i] \geq \frac{3\mu - 1}{4\mu} \sum_{i \in A} \min\{\mathbb{E}[v_i(x_i)], c_i\}.$$

Since $\mathbb{E}[u_i(\mathbf{b})] + \mu \mathbb{E}[b_i] \leq \min\{\mathbb{E}[v_i(d_i)], c_i\}$ for every bidder in A (recall that $\mu \leq 1$ and $v_i(d_i) = u_i(\mathbf{b}) + b_i$) and $\mathbb{E}[b_i] \leq \min\{\mathbb{E}[v_i(d_i)], c_i\}$ for every bidder not belonging to A , the above inequality yields

$$\sum_{i \in A} \min\{\mathbb{E}[v_i(d_i)], c_i\} + \mu \sum_{i \notin A} \min\{\mathbb{E}[v_i(d_i)], c_i\} \geq \frac{3\mu - 1}{4\mu} \sum_{i \in A} \min\{\mathbb{E}[v_i(x_i)], c_i\} \tag{6}$$

By summing (5) and (6), we obtain

$$\begin{aligned}
 \text{EW}(d) &= \sum_{i \in A} \min\{\mathbb{E}[v_i(d_i)], c_i\} + \sum_{i \notin A} \min\{\mathbb{E}[v_i(d_i)], c_i\} \\
 &\geq \frac{3\mu - 1}{4\mu} \sum_{i \in A} \min\{\mathbb{E}[v_i(x_i)], c_i\} + (1 - \mu) \sum_{i \notin A} \min\{\mathbb{E}[v_i(x_i)], c_i\} \\
 &\geq \min\left\{\frac{3\mu - 1}{4\mu}, 1 - \mu\right\} \sum_i \min\{\mathbb{E}[v_i(x_i)], c_i\} \\
 &= \min\left\{\frac{3\mu - 1}{4\mu}, 1 - \mu\right\} \text{EW}^*.
 \end{aligned}$$

Hence, the price of anarchy with respect to the effective welfare benchmark is bounded by the quantity $\min\left\{\frac{3\mu - 1}{4\mu}, 1 - \mu\right\}$ which is maximized to $\frac{7 - \sqrt{17}}{8} \approx 0.3596$ for $\mu = \frac{1 + \sqrt{17}}{8}$. □

The proof for Bayes-Nash equilibria in the incomplete information setting is a bit more complicated. Again, as we did in the proof of Theorem 4.1, we will show that

$$\text{EW}(d) \geq \min\left\{\frac{3\mu - 1}{4\mu}, 1 - \mu\right\} \text{EW}^*.$$

Recall that in the incomplete information setting, both the valuation \mathbf{v}_i and the budget \mathbf{c}_i of bidder i are drawn randomly according to a probability distribution \mathbf{F}_i , and the effective welfare is defined as

$$\text{EW}(d) = \sum_i \mathbb{E}_{(\mathbf{v}_i, \mathbf{c}_i) \sim \mathbf{F}_i} \left[\min\{\mathbb{E}_{(\mathbf{v}_{-i}, \mathbf{c}_{-i}) \sim \mathbf{F}_{-i}} [v_i(d_i)], \mathbf{c}_i\} \right].$$

In order to simplify notation in the proof of Theorem 4.2, we will not explicitly use the subscripts in the expectations.

Theorem 4.2 *The price of anarchy of proportional allocation games with budget-constrained bidders over Bayes-Nash equilibria is at least 0.3596.*

Proof Let $\mu \in (1/3, 1]$ be a parameter whose exact value will be defined later. Consider an incomplete information proportional allocation game with n bidders in which the valuation function \mathbf{v}_i and the budget \mathbf{c}_i of bidder i are drawn from the probability distribution \mathbf{F}_i , independently for each bidder. Let x_i be the random variable denoting the resource fraction bidder i gets in the allocation that maximizes the effective welfare. Let \mathbf{b} be a pure Bayes-Nash equilibrium that induces a random allocation $d = (d_1, \dots, d_n)$ and B be the random variable denoting the sum of bids of all bidders; again, B_{-i} denotes the sum of bids of all bidders besides bidder i . We denote by A_i the set that contains all pairs of a valuation function and a corresponding

budget value (v_i, c_i) that are drawn from the probability distribution \mathbf{F}_i and satisfy $\mathbb{E}[\mathbf{v}_i(d_i)|v_i] \leq c_i$. Consider a bidder i with valuation-budget pair $(v_i, c_i) \notin A_i$. By the definition of A_i , we have

$$\min\{\mathbb{E}[\mathbf{v}_i(d_i)|(v_i, c_i)], c_i\} \geq \min\{\mathbb{E}[\mathbf{v}_i(x_i)|(v_i, c_i)], c_i\}.$$

By considering all valuation-budget pairs not belonging to A_i , we obtain

$$\mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(d_i)], c_i\} \mathbb{1}(v_i, c_i) \notin A_i] \geq \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(x_i)], c_i\} \mathbb{1}(v_i, c_i) \notin A_i],$$

and summing over all bidders, we have

$$\sum_i \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(d_i)], c_i\} \mathbb{1}(v_i, c_i) \notin A_i] \geq \sum_i \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(x_i)], c_i\} \mathbb{1}(v_i, c_i) \notin A_i]. \tag{7}$$

Now consider a valuation-budget pair $(v_i, c_i) \in A_i$ for bidder i that is drawn from \mathbf{F}_i . If $\mu \mathbb{E}[x_i|(v_i, c_i)] \mathbb{E}[B_{-i}|(v_i, c_i)] \leq c_i$, we can bound the expected utility $\mathbb{E}[u_i(\mathbf{b})|(v_i, c_i)]$ by considering the deviation of bidder i to bid $\mu \mathbb{E}[x_i|(v_i, c_i)] \mathbb{E}[B_{-i}|(v_i, c_i)]$ (which is within bidder i 's budget c_i). By Lemma 3.1, we have

$$\begin{aligned} \mathbb{E}[u_i(\mathbf{b})|(v_i, c_i)] &\geq \frac{3\mu - 1}{4\mu} v_i(\mathbb{E}[x_i|(v_i, c_i)] - \mu \mathbb{E}[x_i|v_i]) \mathbb{E}[B_{-i}|(v_i, c_i)] \\ &\geq \frac{3\mu - 1}{4\mu} \mathbb{E}[v_i(x_i)|(v_i, c_i)] - \mu \mathbb{E}[x_i|(v_i, c_i)] \mathbb{E}[B] \\ &\geq \frac{3\mu - 1}{4\mu} \min\{\mathbb{E}[v_i(x_i)|(v_i, c_i)], c_i\} - \mu \mathbb{E}[x_i|(v_i, c_i)] \mathbb{E}[B]. \end{aligned}$$

The second inequality follows by Jensen's inequality and by the fact $\mathbb{E}[B_{-i}|(v_i, c_i)] = \mathbb{E}[B_{-i}]$. Otherwise, if $\mu \mathbb{E}[x_i|(v_i, c_i)] \mathbb{E}[B_{-i}|(v_i, c_i)] > c_i$, the same inequality follows easily since

$$\begin{aligned} \mathbb{E}[u_i(\mathbf{b})|(v_i, c_i)] &\geq 0 \\ &> c_i - \mu \mathbb{E}[x_i|(v_i, c_i)] \mathbb{E}[B_{-i}|(v_i, c_i)] \\ &\geq \frac{3\mu - 1}{4\mu} \min\{\mathbb{E}[v_i(x_i)|(v_i, c_i)], c_i\} - \mu \mathbb{E}[x_i|(v_i, c_i)] \mathbb{E}[B]. \end{aligned}$$

Hence, when $(v_i, c_i) \in A_i$, we have

$$\mathbb{E}[u_i(\mathbf{b})|(v_i, c_i)] + \mu \mathbb{E}[x_i|(v_i, c_i)] \mathbb{E}[B] \geq \frac{3\mu - 1}{4\mu} \min\{\mathbb{E}[v_i(x_i)|(v_i, c_i)], c_i\}.$$

By considering all valuation-budget values belonging to A_i , we have

$$\mathbb{E}[u_i(\mathbf{b}) \mathbb{1}(v_i, c_i) \in A_i] + \mu \mathbb{E}[x_i \mathbb{1}(v_i, c_i) \in A_i] \mathbb{E}[B] \geq \frac{3\mu - 1}{4\mu} \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(x_i)], c_i\} \mathbb{1}(v_i, c_i) \in A_i].$$

Using the obvious fact that $\mathbb{E}[x_i] \geq \mathbb{E}[x_i \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \in A_i]$ and the above inequality, we obtain that

$$\begin{aligned} \mathbb{E}[u_i(\mathbf{b}) \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \in A_i] + \mu \mathbb{E}[x_i] \mathbb{E}[B] &\geq \mathbb{E}[u_i(\mathbf{b}) \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \in A_i] \\ &\quad + \mu \mathbb{E}[x_i \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \in A_i] \mathbb{E}[B] \\ &\geq \frac{3\mu - 1}{4\mu} \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(x_i)], c_i\} \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \in A_i]. \end{aligned} \tag{8}$$

Now, we have

$$\begin{aligned} &\sum_i \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(d_i)], c_i\} \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \in A_i] + \mu \sum_i \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(d_i)], c_i\} \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \notin A_i] \\ &\geq \sum_i \mathbb{E}[u_i(\mathbf{b}) + b_i \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \in A_i] + \mu \sum_i \mathbb{E}[b_i \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \notin A_i] \\ &\geq \sum_i \mathbb{E}[u_i(\mathbf{b}) + \mu b_i \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \in A_i] + \mu \sum_i \mathbb{E}[b_i \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \notin A_i] \\ &= \sum_i \mathbb{E}[u_i(\mathbf{b}) \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \in A_i] + \mu \mathbb{E}[B] \\ &= \sum_i (\mathbb{E}[u_i(\mathbf{b}) \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \in A_i] + \mu \mathbb{E}[x_i] \mathbb{E}[B]) \\ &\geq \frac{3\mu - 1}{4\mu} \sum_i \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(x_i)], c_i\} \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \in A_i]. \end{aligned} \tag{9}$$

The first inequality follows since the quantity $\min\{\mathbb{E}[\mathbf{v}_i(d_i)], c_i\}$ equals $\mathbb{E}[\mathbf{v}_i(d_i)]$ when $(\mathbf{v}_i, \mathbf{c}_i) \in A_i$ and c_i otherwise; in the latter case, the budget is clearly not smaller than the bid of bidder i . The second inequality follows since $\mu \leq 1$, the two equalities are obvious, and the last inequality follows by (8). Now, using (7) and (9), we have

$$\begin{aligned} \text{EW}(d) &= \sum_i \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(d_i)], c_i\}] \\ &= \sum_i \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(d_i)], c_i\} \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \in A_i] \\ &\quad + \mu \sum_i \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(d_i)], c_i\} \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \notin A_i] \\ &\quad + (1 - \mu) \sum_i \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(d_i)], c_i\} \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \notin A_i] \\ &\geq \frac{3\mu - 1}{4\mu} \sum_i \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(x_i)], c_i\} \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \in A_i] \\ &\quad + (1 - \mu) \sum_i \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(x_i)], c_i\} \mathbb{1}(\mathbf{v}_i, \mathbf{c}_i) \notin A_i] \\ &\geq \min \left\{ \frac{3\mu - 1}{4\mu}, 1 - \mu \right\} \sum_i \mathbb{E}[\min\{\mathbb{E}[\mathbf{v}_i(x_i)], c_i\}] \\ &= \min \left\{ \frac{3\mu - 1}{4\mu}, 1 - \mu \right\} \text{EW}^*. \end{aligned}$$

Hence, the price of anarchy with respect to the effective welfare benchmark is bounded by the quantity $\min \left\{ \frac{3\mu-1}{4\mu}, 1-\mu \right\}$ which is maximized to $\frac{7-\sqrt{17}}{8} \approx 0.3596$ for $\mu = \frac{1+\sqrt{17}}{8}$. \square

We conclude this section by presenting our upper bound on the price of anarchy; note that it holds even for pure Nash equilibria.

Theorem 4.3 *For every $\epsilon > 0$, there exists a proportional allocation game among budget-constrained bidders with price of anarchy at most $1/2 + \epsilon$ over pure Nash equilibria, with respect to the effective welfare benchmark.*

Proof Let $\alpha \in (0, 1)$. Consider a proportional allocation game with two bidders. Bidder 1 has valuation $v_1(x) = x$ and budget $c_1 = \frac{\alpha}{(1+\alpha)^2}$. Bidder 2 has valuation $v_2(x) = \alpha x$ and infinite budget. The state in which bidder 1 bids $\frac{\alpha}{(1+\alpha)^2}$ (i.e., her budget) and bidder 2 bids $\frac{\alpha^2}{(1+\alpha)^2}$ is a pure Nash equilibrium, since the derivatives $\frac{b_2}{(b_1+b_2)^2} - 1$ and $\frac{\alpha b_1}{(b_1+b_2)^2} - 1$ of the utilities of the bidders (as functions of their strategies) are equal to zero. Observe that $v_1\left(\frac{1}{1+\alpha}\right)$ significantly exceeds her budget for every value of α . Hence, $\text{EW}(\mathbf{b}) = \frac{\alpha}{(1+\alpha)^2} + \frac{\alpha^2}{1+\alpha} = \frac{\alpha+\alpha^2+\alpha^3}{(1+\alpha)^2}$. The optimal effective welfare is bounded by the welfare at the state when bidder 1 bids her budget $\frac{\alpha}{(1+\alpha)^2}$ and bidder 2 bids $1 - \frac{\alpha}{(1+\alpha)^2}$ so that bidder 1 gets value equal to her budget. Hence, the optimal effective welfare is $\text{EW}^* = \frac{2\alpha+\alpha^2+\alpha^3}{(1+\alpha)^2}$. Clearly, the ratio $\text{EW}(\mathbf{b})/\text{EW}^*$ approaches $1/2$ from above as α approaches 0 and the theorem follows by selecting α to be sufficiently small. \square

5 Discussion and Open Problems

Our work leaves the obvious open problem of computing the tight bound on the price of anarchy over coarse-correlated and Bayes-Nash equilibria. In the full information model, proving bounds specifically for correlated equilibria (a well-known equilibrium class that lies between mixed Nash and coarse-correlated equilibria) is interesting as well. So far, the only upper bound that is known is the counterexample of $3/4$ from [10] for pure Nash equilibria. Is $3/4$ the tight bound for all equilibrium concepts considered in the current paper? Actually, we have not been able to identify any coarse-correlated equilibrium in the full information model that is non-pure. Do such equilibria really exist? Interestingly, we show in Lemma 5.1 that mixed Nash equilibria coincide with pure ones. More generally, this statement applies to mixed Bayes-Nash equilibria in the budget-constrained setting. Does it extend to coarse-correlated ones? We believe that this is an interesting open problem.

Lemma 5.1 *The set of mixed Bayes-Nash equilibria in any proportional allocation game (possibly with budget-constrained bidders) coincides with that of pure Bayes-Nash equilibria.*

Proof Assume that there exists an incomplete information proportional allocation game (possibly with budget-constrained bidders) that has a mixed Bayes-Nash equilibrium \mathbf{b} in which bidder i bids two different values y_1 and y_2 (with $y_1 < y_2$) with non-zero probability when her valuation function is v_i ; both values are within the budget of bidder i (if any). We will show that this is not possible.

By the mixed Bayes-Nash equilibrium condition, both y_1 and y_2 should yield the same maximum expected utility U to bidder i , i.e.,

$$U = \mathbb{E} [u_i(y_1, \mathbf{b}_{-i})|v_i] = \mathbb{E} [u_i(y_2, \mathbf{b}_{-i})|v_i].$$

Let B_{-i} be the random variable denoting the sum of bids of all bidders besides bidder i and let $f(y) = \mathbb{E} \left[v_i \left(\frac{y}{y+B_{-i}} \right) \right]$ be the expected value of bidder i when unilaterally deviating to bid y . Clearly, f is non-decreasing. It is also concave in $[y_1, y_2]$ since it is defined as the linear combination of concave functions: for every value of B_{-i} , $v_i \left(\frac{y}{y+B_{-i}} \right)$ is a concave function with respect to y , and the expectation over B_{-i} is simply a linear combination over such functions. Clearly, $\mathbb{E} [u_i(y, \mathbf{b}_{-i})] = f(y) - y$.

We furthermore claim that f is strictly increasing in $[y_1, y_2]$. If this was not the case, then due to the concavity of f there should exist $y' \in [y_1, y_2]$ such that $f(y') = f(y_2)$ and, hence, bidder i could deviate to bid y' (which is clearly within her budget, if any) for an improved expected utility of $\mathbb{E} [u_i(y', \mathbf{b}_{-i})|v_i] = f(y') - y' > f(y_2) - y_2 = U$. This would contradict the mixed Bayes-Nash equilibrium condition.

The fact that f is strictly increasing clearly implies that $\Pr[B_{-i} > 0] > 0$. But then, for every positive value of B_{-i} , $v_i \left(\frac{y}{y+B_{-i}} \right)$ is a strictly concave function of y and, subsequently, f is also strictly concave as a linear combination of concave functions including strictly concave ones. Hence, there exists $\lambda \in [0, 1]$ such that $f(\lambda y_1 + (1 - \lambda)y_2) > \lambda f(y_1) + (1 - \lambda)f(y_2)$ and bidder i has a profitable deviation to bid $y' = \lambda y_1 + (1 - \lambda)y_2$ (which is again within her budget, if any) since

$$\begin{aligned} \mathbb{E} [u_i(y', \mathbf{b}_{-i})] &= f(y') - y' \\ &> \lambda f(y_1) + (1 - \lambda)f(y_2) - \lambda y_1 - (1 - \lambda)y_2 \\ &= \lambda(f(y_1) - y_1) + (1 - \lambda)(f(y_2) - y_2) \\ &= U. \end{aligned}$$

We conclude that the support of any mixed Bayes-Nash equilibrium cannot contain two different bid values for bidder i when her valuation is v_i and, subsequently (by extending the same argument to all possible valuations of bidder i and to all bidders), it must be a pure Bayes-Nash equilibrium. □

We remark that if the valuation functions are differentiable (we do not make any such assumption in the above proof), a much simpler proof of theorem 5.1 follows by observing that the utility of bidder i , when seen as a function of bidder i 's strategy,

has strictly decreasing derivative. Then, the utility is maximized either by a bid equal to the budget of the bidder (if any) or at the unique bid that nullifies its derivative.

In the Bayesian setting, we have not considered more general equilibrium concepts such as coarse-correlated Bayesian equilibria. A vector of possibly correlated bid functions is called a coarse-correlated Bayesian equilibrium if no bidder i has any incentive to unilaterally deviate to any deterministic bid strategy in order to improve her expected utility (again, given the strategies of the other bidders), for any valuation she draws from her probability distribution \mathbf{F}_i . Coarse-correlated Bayesian equilibria are more general than mixed Bayes-Nash equilibria since the bid functions of different bidders are not restricted to being independent [2]. The main reason for which we have not considered coarse-correlated Bayesian equilibria is that our analysis requires that the expectation of the sum of bids of the other bidders is the same for any possible valuation bidder i can draw from her distribution; this property is not satisfied by more general equilibrium concepts. What is the price of anarchy in this case? Interestingly, the answer cannot be $3/4$ as our next counter-example indicates.

Lemma 5.2 *There exists a proportional allocation game that has price of anarchy at most 0.7154 over coarse-correlated Bayesian equilibria.*

Sketch of Proof Our counter-example has two bidders. Bidder 1 has valuation function x with probability p_1 and $\epsilon_1 x$ with probability $1 - p_1$. Bidder 2 has valuation function αx with probability p_2 and $\epsilon_2 x$ with probability $1 - p_2$. We require $1 \geq \alpha$ and, furthermore, α is significantly larger than ϵ_1 which in turn is significantly larger than ϵ_2 (multiplicatively).

We construct a coarse-correlated Bayesian equilibrium of the following form: When the valuations of the bidders are x and αx , the bid strategies are γ and δ respectively. These values are significant and yield constant resource fractions to both bidders. In all other cases where at least one of the bidders has an almost zero valuation, the bids are extremely close to zero. However, the bidder that has significantly higher valuation than the other submits a significantly higher bid (but still very close to zero) and gets almost 100 % of the resource. In the following, we round negligibly small bids or valuations to 0 and treat an allocation of almost 100 % of the resource to some bidder as exactly 100 %. This rounding does not affect the final result that we can obtain but significantly more detailed (and tedious) calculations are needed for a formal proof. So, we will assume that when the valuations are x or $\epsilon_1 x$ for bidder 1 and $\epsilon_2 x$ for bidder 2, the bid strategies are (almost) 0 but the bid of bidder 1 is significantly higher so that she gets (almost) 100 % of the resource. Similarly, when the valuations are $\epsilon_1 x$ and αx , the bid strategies are (almost) 0 but the bid of bidder 2 is significantly higher so that she gets almost 100 % of the resource.

Notice that the expected utility of bidder 1 when her valuation is x is (approximately) $p_2 \frac{\gamma}{\gamma+\delta} + 1 - p_2 - p_2 \gamma$ and becomes (approximately) $p_2 \frac{y}{y+\delta} + 1 - p_2 - y$ when deviating to a deterministic bid y . We require that the first quantity is higher than the second one so that no such deviation exists, i.e., $p_2 \frac{\gamma}{\gamma+\delta} - p_2 \gamma \geq p_2 \frac{y}{y+\delta} - y$ for every $y \geq 0$. Similarly, we require that $p_1 \frac{\alpha \delta}{\gamma+\delta} - p_1 \delta \geq p_1 \frac{\alpha y}{\gamma+y} - y$. Note that

the right-hand side of the above constraints are maximized to $(\sqrt{p_2} - \sqrt{\delta})^2$ and $(\sqrt{\alpha p_1} - \sqrt{\gamma})^2$, respectively. It remains to compute the exact bid values that satisfy these constraints and minimize the price of anarchy. This is done in the following non-linear mathematical program

$$\begin{aligned} \text{minimize} \quad & \frac{p_1 p_2 \left(\frac{\gamma}{\gamma + \delta} + \alpha \frac{\delta}{\gamma + \delta} \right) + p_1(1 - p_2) + \alpha(1 - p_1)p_2}{p_1 + \alpha(1 - p_1)p_2} \\ \text{subject to:} \quad & p_2 \frac{\gamma}{\gamma + \delta} - p_2 \gamma \geq (\sqrt{p_2} - \sqrt{\delta})^2 \\ & \alpha p_1 \frac{\delta}{\gamma + \delta} - p_1 \delta \geq (\sqrt{\alpha p_1} - \sqrt{\gamma})^2 \\ & \gamma, \delta \geq 0, 0 \leq \alpha \leq 1, 0 \leq p_1, p_2 \leq 1 \end{aligned}$$

which has been solved using Matlab to give an upper bound of 0.7154 for $\alpha = 0.2913$, $\gamma = 0.1071$, $\delta = 0.1510$, $p_1 = 0.6682$, and $p_2 = 0.7616$. \square

Also, recall that we have assumed that bidders have independent valuations. This is a typical assumption in the Bayes-Nash price of anarchy literature [1, 4, 7, 11, 20, 22, 23] with [2] being the only exception we are aware of. Unfortunately, our proof of the pure Bayes-Nash price of anarchy bound does not carry over to the case of correlated valuations either (for the same reason mentioned above). Still, we have not been able to find any counter-example with non-constant price of anarchy in this setting. Again, what is the price of anarchy in this case? These questions are interesting in the budget-constrained setting as well.

Another important issue in the budget-constrained setting is related to our refinement of the effective welfare benchmark for random allocations. Even though our definition naturally extends the one for deterministic allocations, an even more natural definition of the effective welfare for a random allocation d would be $\text{EW}(d) = \sum_i \mathbb{E}[\min\{v_i(d_i), c_i\}]$ (as opposed to the definition $\text{EW}(d) = \sum_i \min\{\mathbb{E}[v_i(d_i)], c_i\}$ that we have adopted); this extends to the Bayesian setting accordingly. Unfortunately, the main obstacle that did not allow us to use this more natural definition is that our techniques cannot establish a relation between the utility of a bidder i and the quantity $\mathbb{E}[\min\{v_i(d_i), c_i\}]$; this is necessary in order to bound the price of anarchy and this has been our approach in the proofs of Theorems 4.1 and 4.2 for the substantially different quantity $\min\{\mathbb{E}[v_i(d_i)], c_i\}$. Interestingly, we have not found any upper bound on the price of anarchy over coarse-correlated or Bayes-Nash equilibria under these alternative definitions for the effective welfare. Is the price of anarchy still bounded by a constant? This question deserves further investigation.

References

1. Bhawalkar, K., Roughgarden, T.: Welfare guarantees for combinatorial auctions with item bidding. In: Proceedings of the 22Nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 700–709 (2011)

2. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M., Lucier, B., Paes Leme, R., Tardos, E.: Bounding the inefficiency of outcomes in generalized second price auctions. *J. Econ. Theory (JET)* **156**, 343–388 (2015)
3. Caragiannis, I., Voudouris, A.A.: Welfare guarantees for proportional allocations. In: *Proceedings of the 7Th International Symposium on Algorithmic Game Theory (SAGT)*, LNCS 8768, Springer, pp. 206–217 (2014)
4. Christodoulou, G., Kovács, A., Schapira, M.: Bayesian combinatorial auctions. In: *Proceedings of the 35Th International Colloquium on Automata, Languages, and Programming (ICALP)*, Part 1, LNCS 5125, Springer, pp. 820–832 (2008)
5. Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: The price of anarchy of the proportional allocation mechanism revisited. In: *Proceedings of the 9Th Conference on Web and Internet Economics (WINE)*, LNCS 8289, Springer, pp. 109–120 (2013)
6. Dobzinski, S., Paes Leme, R.: Efficiency guarantees in auctions with budgets. In: *Proceedings of the 41St International Colloquium on Automata, Languages and Programming (ICALP)*, LNCS 8572, Springer, pp. 392–404 (2014)
7. Feldman, M., Fu, H., Gravin, N., Lucier, B.: Simultaneous auctions are (almost) efficient. In: *Proceedings of the 45Th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 201–210 (2013)
8. Hajek, B., Gopalakrishnan, G.: Do greedy autonomous systems make for a sensible Internet? Unpublished manuscript (2002)
9. Johari, R.: The Price of Anarchy and the Design of Scalable Resource Allocation Algorithms. Chapter 21 in *Algorithmic Game Theory*, pp. 543–568. Cambridge University Press (2007)
10. Johari, R., Tsitsiklis, J.N.: Efficiency loss in a network resource allocation game. *Math. Oper. Res.* **29**(3), 407–435 (2004)
11. de Keijzer, B., Markakis, E., Schaefer, G., Telelis, O.: Inefficiency of standard multi-unit auctions. In: *Proceedings of the 21St Annual European Symposium on Algorithms (ESA)*, LNCS 8125, Springer, pp. 385–396 (2013)
12. Kelly, F.P.: Charging and rate control for elastic traffic. *Eur. Trans. Telecommun.* **8**, 33–37 (1997)
13. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. In: *Proceedings of the 16Th Symposium on Theoretical Aspects of Computer Science (STACS)*, LNCS 1563, Springer, pp. 404–413 (1999)
14. La, R.J., Anantharam, V.: Charge-sensitive TCP and rate control in the internet. In: *Proceedings of the 19Th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pp. 1166–1175 (2000)
15. Maheswaran, R.T., Basar, T.: Nash equilibrium and decentralized negotiation in auctioning divisible resources. *Group Decis. Negot.* **12**(5), 361–395 (2003)
16. Maheswaran, R.T., Basar, T.: Social welfare of selfish agents: motivating efficiency for divisible resources. In: *Proceedings of the 43Rd IEEE Conference on Decision and Control (CDC)*, pp. 1550–1555 (2004)
17. Nguyen, T., Tardos, E.: Approximately maximizing efficiency and revenue in polyhedral environments. In: *Proceedings of the 8Th ACM Conference on Electronic Commerce (EC)*, pp. 11–20 (2007)
18. Nguyen, T., Vojnovic, M.: Weighted proportional allocation. In: *Proceedings of the 2011 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pp. 173–184 (2011)
19. Roughgarden, T.: Potential functions and the inefficiency of equilibria. In: *Proceedings of the International Congress of Mathematicians*, vol. III, pp. 1071–1094 (2006)
20. Roughgarden, T.: The price of anarchy in games of incomplete information. In: *Proceedings of the 13Th ACM Conference on Electronic Commerce (EC)*, pp. 862–879 (2012)
21. Sanghavi, S., Hajek, B.: Optimal allocation of a divisible good to strategic buyers. In: *Proceedings of the 43Rd IEEE Conference on Decision and Control (CDC)*, pp. 2748–2753 (2004)
22. Syrgkanis, V.: Bayesian games and the smoothness framework (2012). arXiv:[1203.5155](https://arxiv.org/abs/1203.5155)
23. Syrgkanis, V., Tardos, E.: Composable and efficient mechanisms. In: *Proceedings of the 45Th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 211–220 (2013)
24. Syrgkanis, V., Tardos, E.: Bayesian sequential auctions. In: *Proceedings of the 13Th ACM Conference on Electronic Commerce (EC)*, pp. 929–944 (2012)