

Profit Sharing and Efficiency in Utility Games*

Sreenivas Gollapudi¹, Kostas Kollias², Debmalya Panigrahi³, and Venetia Pliatsika⁴

1 Google Research, Mountain View, USA

2 Google Research, Mountain View, USA

3 Duke University, Durham, USA

4 University of Pennsylvania, Philadelphia, USA

Abstract

We study utility games (Vetta, FOCS 2002) where a set of players join teams to produce social utility, and receive individual utility in the form of payments in return. These games have many natural applications in competitive settings such as labor markets, crowdsourcing, etc. The efficiency of such a game depends on the profit sharing mechanism – the rule that maps utility produced by the players to their individual payments. We study three natural and widely used profit sharing mechanisms – egalitarian or equal sharing, marginal gain or value addition when a player joins, and marginal loss or value depletion when a player leaves. For these settings, we give tight bounds on the price of anarchy, thereby allowing comparison between these popular mechanisms from a (worst case) social welfare perspective.

1998 ACM Subject Classification J.4 Social and Behavioral Sciences: Economics

Keywords and phrases Price of anarchy, submodular maximization, coverage functions

Digital Object Identifier 10.4230/LIPIcs.ESA.2017.43

1 Introduction

In *utility games* (introduced by Vetta [20], see also [14]), individual agents (e.g., employees) offer their services to entities (e.g., employers) to create social utility, and receive individual utility in the form of payments in return. It is natural to expect the agents to behave strategically, i.e., offer their services to the entity giving them the highest payment. This represents a game where each agent (called a *player*) selects one of the available entities (called *teams*) to maximize their individual payments (called *payoffs*), but the overall *social welfare* is the total utility cumulatively produced by all the teams. A stable outcome, called a *Nash equilibrium* or NE, is achieved when no player can unilaterally change her team and increase her payoff. The goal of this paper is to study the (in)efficiency of such stable outcomes – the maximum (worst-case) ratio between the social welfare produced by an optimal allocation and that in an NE, called the Price of Anarchy or POA of the game. Clearly, the social welfare produced at equilibrium depends on the *profit sharing* mechanism in use, i.e., the payoff of the players as a function of the utility produced by them. We consider three natural and widely used profit sharing rules:

- *egalitarian*: any unit of utility produced by a team is divided equally among the team members who helped produce it.

* A full version of this paper is available at http://theory.stanford.edu/~kkollias/profit_sharing.pdf



43:2 Profit Sharing and Efficiency in Utility Games

- *marginal gain*: the payoff of a player in a team is the utility that the team gained when she joined.
- *marginal loss*: the payoff of a player in a team is the utility that the team would lose if she were to leave.

The main motivation for this work is to compare these popular profit sharing mechanisms in terms of their impact on (worst case) social utility.

Formally, there is a set of *players* N and a set of *teams* T that they can join. The utility produced by a team is given by a *weighted coverage function* over the players in the team. To interpret this, consider a set of tasks S with respective utilities v_s for $s \in S$. If each player $i \in N$ can perform a subset of tasks $S_i \subseteq S$, then the tasks completed by a team $t \in T$ is given by $S_t = \cup_{i \in t} S_i$, and the utility produced is $U_t = \sum_{s \in S_t} v_s$. This is precisely a weighted coverage function over the team members; we say that team t *covers* task s if $s \in S_t$.

From a social perspective, the goal is to maximize the total utility produced by all teams,

$$U = \sum_{t \in T} U_t = \sum_{t \in T} \sum_{s \in S_t} v_s.$$

We will call U the *social welfare* or *objective value*.

Profit Sharing Mechanisms

Each player i is interested in maximizing her own payoff, denoted u_i , which depends on the profit sharing mechanism. We consider the following popular profit sharing mechanisms.

Egalitarian Profit Sharing. For every task that is covered in a team, the utility of the task is equally shared among the members of the team who perform the task. We may note that the egalitarian model is an instantiation of the *Shapley value* utility sharing method [19]. This is defined as the expected contribution of a player to her team's utility, assuming players are sequentially added to the team using a uniformly random ordering.

For egalitarian sharing, we exactly determine the POA to be 1.6 by proving matching upper and lower bounds on the POA. The upper bound employs the smoothness framework due to Roughgarden [18]. However, unlike the standard approach of applying the smoothness inequality for every resource, (in our setting, for every team), we apply the smoothness inequality across all teams and players simultaneously for a single task. The matching lower bound of 1.6, on the other hand, uses a careful combinatorial construction of the worst case POA instance, using symmetrization techniques to argue stability of the solution. Both these results appear in Section 2, thereby proving the following theorem.

► **Theorem 1.** *The POA of egalitarian profit sharing is 1.6.*

Marginal Gain Profit Sharing. In this model, players have an order of arrival and each player's utility is the value added to the team when that player joins. The marginal gain method is an instantiation of an *ordinal Shapley value*, which is a variation on the Shapley value discussed above where player ordering is not random but predefined.

For marginal gain profit sharing, we show an upper bound of 1.71 and a lower bound of 1.58 on the POA. The lower bound can be established by hardness of approximation results (see [11]) assuming $\mathbf{P} \neq \mathbf{NP}$. We give an alternative proof of this result based on an explicit construction that does not rely on complexity theoretic assumptions in our full paper. On the other hand, our upper bound, which appears in Section 3, is based on a charging argument, which carefully matches tasks that are not covered in the NE to covered tasks.

► **Theorem 2.** *The POA of marginal gain profit sharing is at most $1 + \frac{1}{\sqrt{2}} \simeq 1.71$ and at least $\frac{e}{e-1} \simeq 1.58$.*

Marginal Loss Profit Sharing. In this model, the utility of a player is the value lost if she were to leave the team. The marginal loss model is an instantiation of the *marginal contribution* method, where each player is rewarded with her marginal contribution to the utility of her team [19, 10]. One interesting point of difference between this and the previous two models is that the sum of individual payoffs in this case may be strictly smaller than the overall social utility, whereas this sum was exactly equal to the social utility in the previous cases. This is because the only tasks whose utility is awarded as payoff in this model are those that are uniquely performed by a single team member.

For marginal loss profit sharing, we prove the POA is exactly 2. An upper bound of 2 follows from the work of Vetta [20]. We show a matching lower bound via an explicit construction with ideas that are similar to the lower bound construction in Theorem 1. Due to space constraints, we present this lower bound construction in our full paper, thereby proving the next theorem.

► **Theorem 3.** *The POA of marginal loss profit sharing is 2.*

Existence of NE. Omitting further details, we briefly mention why existence of a NE is guaranteed in all three profit sharing models. Egalitarian profit sharing induces a congestion game [17, 15], which implies existence of a NE is guaranteed. In marginal gain profit sharing, the property follows by equivalence to the process of having players appear online and letting each player select the team yielding the higher profit at the time of her arrival. In marginal loss profit sharing, the optimal solution is always a NE since any beneficial deviation by a player by definition increases the objective value.

Extensions

Submodular Utilities. While we primarily consider utility functions that are weighted coverage functions of the players in a team, Vetta [20] has originally proposed the utility game framework for more general *submodular* utility functions. The marginal gain and marginal loss profit sharing models naturally extend to this setting with the same definitions. In both cases, we show that the POA is 2, i.e., that Vetta's upper bound is tight. For the marginal loss model, this follows as a corollary of Theorem 3 since weighted coverage functions are a special case of submodular functions. For the marginal gain mode, however, this establishes a separation between the POA for general submodular functions and the special case of weighted coverage functions. Our lower bound construction for general submodular functions makes use of a knapsack welfare function, and appears in our full paper.

► **Theorem 4.** *The POA of marginal gain and marginal loss profit sharing with submodular utilities is 2.*

One can also extend the egalitarian profit sharing model to the case of general submodular functions by using the analogy with Shapley profit sharing. In particular, for submodular functions, the payoff of a player in the egalitarian case is her expected contribution to the utility of the team, assuming a uniform random order of arrival of players. Determining the POA in this case is an interesting problem for future work.

Asymmetric Utilities. We can also consider an *asymmetric* setting, where the utility functions of teams are not necessarily identical. This is the case, e.g., if a task produces different utility to different teams. In this case, even for weighted coverage utility functions (and therefore, also for submodular utilities), we show that the POA is 2, i.e., Vetta’s upper bound is tight. For the marginal loss model, this follows from Theorem 3 since symmetric utilities are a special case, but for the egalitarian and symmetric gain models, we require new lower bound constructions that are given in our full paper.

► **Theorem 5.** *The POA in the egalitarian, marginal gain, and marginal loss models for asymmetric teams is 2.*

Related Work

Utility games were introduced by Vetta [20] to model strategic agents who produce submodular social welfare in a team, and seek to maximize their individual utility or payoff in return. For this general setting, Vetta showed an upper bound of 2 on the POA, subject to some mild conditions on the agent payoff functions that are satisfied in all our models above. The profit sharing models that we study in this paper are inspired by standard cost sharing rules from the economics literature such as Shapley and marginal contribution costs as mentioned earlier. The POA of these cost sharing models has been studied in several resource selection problems with negative externalities among players. Specifically, Marden and Wierman [14] studied utility sharing methods in a general distributed utility maximization model, and Harks and Miller [9] studied cost sharing methods in networking applications. Bachrach *et al.* [1] considered the effect of *positive* externalities among players working on multiple projects simultaneously, but restricted by an effort budget. In [13], the authors study a special case of utility games where the welfare produced by a resource is a function of the number of players on it, and prove that under certain conditions (such as symmetric players), the POA drops below 2.

Utility games with coverage utility functions are also related to congestion games [17, 15]; in fact, utility games in the egalitarian profit sharing model *are* congestion games. The POA of cost sharing methods in generalizations of congestion games has been extensively studied [3, 12, 8]. Gairing [7] studied a congestion game with a coverage utility function and showed how to modify the payoffs of the tasks so that better-response dynamics reaches an equilibrium with an inefficiency of $1 - \frac{1}{e}$ or better in polynomial time.

Finally, we mention known results for the corresponding optimization problem – make an assignment of players to teams to maximize overall social utility, where the utility on every team is given by a weighted coverage function. This problem is called *submodular welfare maximization with coverage functions*. The best approximation ratio in the general case is $\frac{e}{e-1} \simeq 1.58$ [2, 5, 6]. Algorithmic results on combinatorial auctions, which are similar to our setting (teams are bidders and players are items) include a $1 - \frac{1}{e}$ approximation algorithm for submodular valuations [4], a proof of optimality of the greedy algorithm in various online and offline settings [16], and a (matching) hardness of approximation result [11].

2 Egalitarian Payoffs

In this section, we prove Theorem 1, i.e., show that the POA for the egalitarian profit sharing model is exactly 1.6. This comprises two parts: a lower bound of 1.6 (in Section 2.1) and a matching upper bound of 1.6 (in Section 2.2).

2.1 Lower Bound for Egalitarian Payoffs

► **Lemma 6.** *The POA of egalitarian profit sharing is at least 1.6.*

We construct an instance of the egalitarian model and an assignment of players to teams that is an NE and whose social welfare is a $\frac{5}{8}$ fraction of the optimal solution. We begin with an overview of this construction, and then give details of each step. First, we create a simple instance parameterized by integers x and y (we will precisely define these integers later), and an assignment of players to teams with utility $\frac{x+y}{2x+y}$ times the optimal. Our assignment in this preliminary game will *not* be an NE. We then modify the instance in two stages, where we preserve the ratio $\frac{x+y}{2x+y}$ w.r.t. the optimal solution, while creating sufficient structure to argue that the final assignment is an NE for appropriate values of x and y . The worst case among these equilibrium-inducing (x, y) values will yield the POA lower bound of 1.6.

Checking whether the final assignment is an NE can be a complicated task in general, since there will eventually be a large number of players and possible deviations in the game. Our two-stage transformation will ensure, however, that this task reduces to verifying a single inequality. This will be achieved by imposing symmetry across players (first transformation) and symmetry across possible deviations of a player (second transformation). We now present the four stages of our proof (initialization, imposing player symmetry, imposing deviation symmetry, and picking the values of x, y) in detail.

Stage 1: Initialization. Our preliminary game uses the parameter $k = 2x + y$. There are k tasks s_1, s_2, \dots, s_k , and k types of players where a player of type i can *only* perform task s_i . There are k players for each type, i.e., a total of k^2 players. The number of teams is also k . The utility produced by covering any single task in a team is 1.


We will crucially maintain two properties of the assignment. The first property imposes symmetry over how players are divided among teams.

► **Property 7.** *Note that there are $k = 2x + y$ players who can perform a task. Our assignment will ensure that every task is covered by 2 players in x teams, by 1 player in y teams, and remains uncovered in x teams. We will also ensure that every team has $k = 2x + y$ players. These k players in any team will cover tasks as follows: x tasks will be covered by 2 players, y tasks will be covered by 1 player, and x tasks will remain uncovered.*


Note that the above property ensures that every team only covers $x + y$ tasks out of the total of $k = 2x + y$ tasks. Similarly, every task is covered in only $x + y$ teams out of the total of $k = 2x + y$ teams.

The second property relates our assignment to an optimal assignment (call it OPT). To encode OPT, let us use k colors c_1, c_2, \dots, c_k , where all players assigned to team i by OPT are said to have color c_i .

► **Property 8.** *OPT will satisfy the property that there is exactly one player with color c_i who can perform a specific task s_j , for any i and j . In other words, the $k = 2x + y$ players who can perform any specific task will be divided among the $k = 2x + y$ teams, thereby ensuring that all tasks in all teams are covered. Contrast this to our assignment that only covers $x + y$ tasks in every team, and $x + y$ teams cover every task, according to Property 7. Finally, in our assignment, there will be exactly one player of each color c_i in every team t . In other words, the overlap between any team in our assignment and any team in OPT will be exactly one player.*



■ **Figure 1** The preliminary assignment for $x = y = 1$. Teams are rows, tasks are columns. Cell (i, j) corresponds to team i , task s_j . OPT has all white players in team 1, all gray players in team 2, and all black players in team 3. The occupancy is symmetric across rows and columns and each color appears once per row and column.



■ **Figure 2** The intermediate assignment for $x = y = 1$. The first copy is the original preliminary assignment. In the second one, white becomes gray, gray becomes black, black becomes white. In the third we shift the colors again. All three together form the intermediate assignment.

As noted above, Property 7 implies that the coverage of our assignment is $\frac{x+y}{2x+y}$ times the total number of tasks, while Property 8 ensures that the optimal solution covers every task. However, it is not a priori clear that these properties can be satisfied by an assignment: the next lemma asserts this.


► **Lemma 9.** *Given k teams, k tasks, and, for each task, k players who can perform only that task, there is an assignment of the players to the teams and a coloring that satisfies Properties 7 and 8.*

Proof. (See Fig. 1 for an illustration of the $x = y = 1$ case.) The first team’s structure is as follows: tasks s_1, s_2, \dots, s_x are covered by two players, tasks $s_{x+1}, s_{x+2}, \dots, s_{x+y}$ are covered by one player each, and tasks $s_{x+y+1}, s_{x+y+2}, \dots, s_k$ are left uncovered. For this first team, we use any coloring that has a different color for each of the k players. The structure and coloring of the second team is obtained by performing a left circular shift to the first team’s structure, i.e., $s_k, s_1, s_2, \dots, s_{x-1}$, are covered by 2 players, $s_x, s_{x+1}, \dots, s_{x+y-1}$ are covered by 1 player, and $s_{x+y}, s_{x+y+1}, \dots, s_{k-1}$, are left uncovered. Colors are also shifted, i.e., the color(s) of the player(s) covering s_i in the first team is applied to the player(s) covering s_{i-1} (s_k , if $i = 1$) in the second team. We continue with similar left circular shifts to define the remaining teams. This assignment and coloring satisfies Properties 7 and 8. ◀

We have now completed the first stage; we will call this the *preliminary assignment*. By Property 8, the optimal assignment covers all tasks; hence, the ratio of the coverage of this preliminary assignment to the optimum is $\frac{x+y}{2x+y}$. However, this assignment is not an NE, since players sharing a task have unilateral incentive to deviate to a team where the corresponding task is not covered. We now proceed to the next stages, which will modify this assignment to an NE.

Stage 2: Imposing player symmetry. During this stage, we will augment the game by adding new tasks. In our preliminary assignment, not all players have the same payoff since some of them share a task with a teammate while others do not. In this stage, we impose symmetry across players: every player will share exactly $2x$ tasks with another player and will cover exactly y tasks by herself. To do this, we create k copies of our preliminary assignment, and exchange roles between players in the different copies in a way that they all end up being symmetric. We will call this the *intermediate assignment*.

The first copy is identical to the preliminary assignment. In the second copy, we take the preliminary assignment and perform a circular shift on the colors, i.e., we change color c_i to



■ **Figure 3** All 6 versions of the intermediate assignment for $x = y = 1$. The first one is the original intermediate assignment. The rest are all possible permutations of the team structures (i.e., rows). All 6 together form the final assignment with 9 players: one white, one gray, and one black player in each team.

color c_{i+1} (color c_k changes to c_1). Next, we rename the tasks so that they are distinct from those in the first copy. We continue this process of doing a circular shift on the colors and renaming the tasks in each subsequent copy until we have k copies in total. (See Fig. 2 for all the copies of the $x = y = 1$ case.) The intermediate assignment is constructed by appending all k copies (recall that the tasks are distinct in the copies), and merging all players in the same team with the same color into a single player.

Note that properties 7 and 8 continue to hold; in particular, this implies that the intermediate assignment covers a $\frac{x+y}{2x+y}$ fraction of tasks in each team, while OPT covers every task in every team. Moreover, since every color assumes the role of every other color in the preliminary assignment in one of the copies, it follows that every player covers $2x$ tasks with another player and y tasks by herself in the intermediate assignment. This implies that the players are symmetric in their coverage and payoff in their current team. However, the possible deviations of a player to another team are not symmetric, i.e., the payoff of a player depends on the team that the player moves to. In the next stage, we impose symmetry on the deviations of players, thereby reducing the equilibrium condition to a single inequality.

Stage 3: Imposing Deviation Symmetry. In this stage, we repeatedly perform an operation that we call *team structure switch*. Switching the structure of team t to that of t' involves taking each player in t , stripping her of her existing tasks, and granting her the tasks of the player in t' with the same color. By Property 8, this player in t' is uniquely defined given a specific player in t . A *team structure permutation* is said to be performed when we switch the structure of every team t to the structure of team $\pi(t)$, where π is a permutation on the teams T .

For every possible permutation π , we generate a copy of the intermediate assignment and perform a team structure permutation based on π . As we did in the previous stage, we rename tasks so that they are different for each permuted copy and incorporate all $k!$ copies into our game by merging players in the same team with the same color into a single player with $k \cdot k!$ tasks. This generates our *final assignment*. (See Fig. 3 for the copies corresponding to the six permutations for the $x = y = 1$ case.)

Again, note that Properties 7 and 8 continue to hold; as a consequence, each team only covers a $\frac{x+y}{2x+y}$ of the tasks in the final assignment whereas OPT covers every task in every

team. Additionally, every deviation of a player to another team now result in exactly the same utility; therefore, not only are the players symmetric in their current team, but their deviation to any other team is also symmetric.

► **Lemma 10.** *In the final assignment, the utility of any player i who deviates to a team t' that is not her assigned team t is given by*

$$\left[\left(\frac{x-1}{3} + \frac{y}{2} + x \right) 2x + \left(\frac{x}{3} + \frac{y-1}{2} + x \right) y \right] k(k-2)!.$$

Proof. Consider a player i , and call her assigned team t . Fix some structure for t in an intermediate assignment, and focus on all versions of the intermediate assignment in which team t has that structure. There will be $(k-1)!$ such versions. Consider any task s that is covered by i and another player in this structure. Our first goal is to determine the coverage of task s in any other team t' in each of the $(k-1)!$ versions of the intermediate assignment that we are considering.

The copies of t' in the $(k-1)!$ versions we are considering can assume one of $k-1$ possible structures, excluding the structure that we have fixed for t . Each one of these $k-1$ possible structures for t' appears an equal number of times, i.e., $(k-2)!$ times. By Property 7, $x-1$ of these $k-1$ structures have 2 players covering task s . (Note that t itself has 2 players covering s , hence the number is $x-1$ and not x .) Similarly, in y of these structures, s is covered by a single player, and it is not covered at all in x structures. This implies that the payoff of i due to s , if she deviates to t' , will be $\frac{x-1}{3} + \frac{y}{2} + x$ when we sum across one copy of each structure of t' . For the overall payoff of i after deviation to t' due to tasks shared with another player in t , we need to multiply this expression by:

- $2x$, which represents the number of different tasks s that are covered by i and another player in t ,
- $(k-2)!$, which represents the number of copies with the same structure of t' , given a fixed structure of t , and
- k , which is the number of different structures of t .

This yields a payoff of:

$$\left(\frac{x-1}{3} + \frac{y}{2} + x \right) 2xk \cdot (k-2)! \tag{1}$$

In a similar manner, we can calculate the payoff that i would get by deviating to t' due to the tasks she uniquely covers in t . This comes out to:

$$\left(\frac{x}{3} + \frac{y-1}{2} + x \right) yk \cdot (k-2)! \tag{2}$$

The total payoff after deviation for player i is then given by:

$$\left[\left(\frac{x-1}{3} + \frac{y}{2} + x \right) 2x + \left(\frac{x}{3} + \frac{y-1}{2} + x \right) y \right] k \cdot (k-2)!, \tag{3}$$

which is independent of i , t and t' . This completes the proof of the lemma. ◀

Stage 4: Choice of the parameters x and y . Note that the payoff of a player in the final assignment is $(x+y)k!$, since the payoff in every copy of the intermediate assignment is $x+y$ and there are $k!$ copies. Therefore, by Lemma 10, the equilibrium condition is:

$$\left[\left(\frac{x-1}{3} + \frac{y}{2} + x \right) 2x + \left(\frac{x}{3} + \frac{y-1}{2} + x \right) y \right] k(k-2)! \leq (x+y)k!.$$

Since $k = 2x + y$, this simplifies to:

$$(x + y)(2x + y - 1) \geq \left(\frac{x-1}{3} + \frac{y}{2} + x \right) 2x + \left(\frac{x}{3} + \frac{y-1}{2} + x \right) y.$$

We can verify that this equilibrium condition holds if we set $y = \frac{2+\epsilon}{3}x$, with $\epsilon > 0$ arbitrarily small, and let $x \rightarrow \infty$. We then get a $\frac{2x+y}{x+y}$ lower bound on the POA, which is arbitrarily close to 1.6. This completes the proof of Lemma 6.

2.2 Upper Bound for Egalitarian Payoffs

► **Lemma 11.** *The POA of egalitarian profit sharing is at most 1.6.*

We will apply the (λ, μ) smoothness framework of Roughgarden [18]. For the purposes of this proof, we extend the game by introducing a new strategy for each player, which is to split herself into $|T|$ fractions of $\frac{1}{|T|}$ each, and assign each fraction to a different team; call this the *fractional* strategy. We define the payoff of a $\frac{1}{|T|}$ -sized fractional player sharing a task s with another n (integral) players in a team as $\frac{1}{|T|} \cdot \frac{v_s}{n+1}$. If every player plays her fractional strategy, then we denote the outcome OPT-FR. Let N_s be the set of players who can perform task s . We define the utility of a task s in team t in the outcome OPT-FR as $v_s \cdot \min \left\{ \frac{|N_s|}{|T|}, 1 \right\}$.

We prove two important properties of this augmented game. The first property is that OPT-FR represents an optimal fractional solution to the optimization problem maximizing the total utility; therefore, its utility is at least that of OPT, which is the optimal integral solution to the same problem.

► **Property 12.** *The total utility (social welfare) of OPT-FR is at least that of OPT.*

This property allows to compare the utility in any NE with that in OPT-FR instead of OPT in order to obtain an upper bound on the POA. Since OPT-FR is highly symmetric, this is a simpler comparison that does not require delving into the structure of OPT.

The second property establishes that any NE in the original game is also an NE in the augmented game, i.e., no player has an incentive to deviate to her fractional strategy. This property holds because a deviation to the fractional strategy would produce payoff for the player that is a convex combination of her current payoff and the payoffs produced by deviating (integrally) to the other teams. Since none of these integral deviations produces a higher payoff, neither does the deviation to the fractional strategy.

► **Property 13.** *If NASH is an NE in the original game, then no player has an incentive to deviate to her fractional strategy.*

Let u_i^d be the payoff of player i if she unilaterally deviates from her team in NASH to her fractional strategy. Also, let U be the total utility (social welfare) in NASH and U^* be the total utility in OPT-FR. Our goal will be to identify positive parameters λ and μ such that for any equilibrium NASH,

$$\sum_{i \in N} u_i^d \geq \lambda U^* - \mu U. \quad (4)$$

Using Property 13, we have: $U = \sum_{i \in N} u_i \geq \sum_{i \in N} u_i^d \geq \lambda U^* - \mu U$. By rearranging the terms, we get $\frac{U^*}{U} \leq \frac{\mu+1}{\lambda}$. A POA bound of $\frac{\mu+1}{\lambda}$ now follows from Property 12. We will show Eq. (4) with $\lambda = \frac{5}{6}$ and $\mu = \frac{1}{3}$, which will then give us the desired upper bound of 1.6.

43:10 Profit Sharing and Efficiency in Utility Games

Our task, therefore, is to prove Eq. (4) with $\lambda = \frac{5}{6}$ and $\mu = \frac{1}{3}$. Let us initially focus on a single task s . Let n be the number of players who can perform task s , k be the number of teams, and h be the total utility produced by task s across all the teams in OPT-FR. To compare this utility with that in the equilibrium NASH, we use γ to denote the ratio of utilities for task s in the two assignment OPT-FR and NASH. In other words, γh denotes the total utility produced by task s in NASH. We examine two cases: $\gamma < \frac{1}{2}$ and $\gamma \geq \frac{1}{2}$. (Since we argue Eq. (4) for each task separately, we assume wlog that $v_s = 1$.)

Case 1: $\gamma < \frac{1}{2}$. When a player $i \in N_s$ deviates to her fractional strategy unilaterally, her payoff from task s is the sum of payoffs from the $k - \gamma h$ teams that do not cover task s and the γh teams that already cover s . This is given by:

$$\frac{1}{k}(k - \gamma h) + \frac{1}{k} \sum_{t: n_{s,t} > 0} \frac{1}{n_{s,t} + 1} \geq \frac{1}{k}(k - \gamma h) + \frac{1}{k} \cdot \gamma h \cdot \frac{1}{\frac{n}{\gamma h} + 1} \quad (\text{by convexity}). \quad (5)$$

Summing over the n players who can perform task s gives:

$$\frac{n}{k}(k - \gamma h) + \frac{n}{k} \cdot \frac{(\gamma h)^2}{n + \gamma h} \geq (1 - \gamma)h + \frac{(\gamma h)^2}{h + \gamma h}. \quad (6)$$

The inequality follows by replacing n and k with their smaller or equal number h , since the left-hand side is increasing as a function of n and k . We can then verify that for our values of $\lambda = \frac{5}{6}$ and $\mu = \frac{1}{3}$ and for any $\gamma < \frac{1}{2}$, the last expression from (6) satisfies,

$$(1 - \gamma)h + \frac{(\gamma h)^2}{h + \gamma h} \geq \lambda h - \mu \gamma h. \quad (7)$$

Case 2: $\gamma \geq \frac{1}{2}$. Similar to Case 1, the sum of payoffs for deviating from NASH to the fractional strategies is at least

$$\frac{n}{k}(k - \gamma h) + \frac{n}{k} \sum_{t: n_{s,t} > 0} \frac{1}{n_{s,t} + 1}. \quad (8)$$

Note that the sum of all $n_{s,t}$ values must be equal to n . Also, note that an adversary minimizing $\sum_{t: n_{s,t} > 0} \frac{1}{n_{s,t} + 1}$ sets all $n_{s,t}$ values equal and, if it turns out to be non-integral, then rounds some of them up and some down to keep their sum at n . Now consider the expression $n \sum_{t: n_{s,t} > 0} \frac{1}{n_{s,t} + 1}$ and suppose that the $n_{s,t}$ values have been picked by the adversary as above. Consider increasing n by one. Then the adversary will also increase exactly one of the $n_{s,t}$ values to restore the property that they sum to n . This will clearly increase the value of the expression $n \sum_{t: n_{s,t} > 0} \frac{1}{n_{s,t} + 1}$. Hence, we again get a lower bound on (8) by substituting n and k with their smaller number h , and letting the adversary pick $n_{s,t}$ values summing to h .

At this point, we know that the sum of $n_{s,t}$ values will be equal to h and that the number of $n_{s,t}$ variables is γh , with $\gamma \geq \frac{1}{2}$. Therefore, each $n_{s,t}$ value chosen by the adversary will be either 1 or 2. Since the average of the $n_{s,t}$ values must be equal to $\frac{1}{\gamma}$, there is also the constraint that $2\beta + 1(1 - \beta) = \frac{1}{\gamma}$, where β is the fraction of $n_{s,t}$ variables with value 2. After solving, we get $\beta = \frac{1}{\gamma} - 1$. Then, the inequality corresponding to (7) in Case 1 becomes

$$(1 - \gamma)h + \left(\frac{1}{\gamma} - 1\right) \gamma h \frac{1}{3} + \left(2 - \frac{1}{\gamma}\right) \gamma h \frac{1}{2} \geq \lambda h - \mu \gamma h, \quad (9)$$

which is always true for $\lambda = \frac{5}{6}$ and $\mu = \frac{1}{3}$.

Combining the above two cases, and summing over all tasks s , we can conclude that (4) holds as desired.

3 Marginal Gain Payoffs

In this section, we prove our upper bound from Theorem 2, i.e., show that the POA for marginal gain is at most $1 + \frac{1}{\sqrt{2}}$. We note that the lower bound of Theorem 2 can be established by hardness of approximation results (see [11]) assuming $\mathbf{P} \neq \mathbf{NP}$ but we also provide an explicit construction in our full paper.

► **Lemma 14.** *The POA of marginal gain is at most $1 + \frac{1}{\sqrt{2}}$.*

We will assume wlog that all tasks have unit utility, since any task can be decomposed into multiple unit-utility tasks. We say that player i has a *hit* on a task that she can perform if she receives payoff for it, i.e., is the first person in her team to perform the task; if not, we call it a *waste* of the task. We will write OPT for the optimal outcome and NASH for some given NE. Consider some task $s \in S$ and denote the number of hits on this task in OPT (resp., NASH) by h_s^* (resp., h_s) and the number of wastes by w_s^* (resp., w_s). Consider the quantity $w_s^+ = w_s - w_s^*$. For tasks that have $w_s^+ > 0$, we will arbitrarily select w_s^+ of the wastes in NASH and label them as the *additional* wastes of NASH against OPT. Note that since $h_s^* + w_s^* = h_s + w_s$, w_s^+ is precisely the difference in social utilities of NASH and OPT due to task s . For ease of exposition we make the following modification to the values of h_s^* and w_s^* : for tasks with $w_s^+ < 0$, we raise h_s^* (and accordingly lower w_s^*) until $w_s^+ = 0$. These changes improve the situation for OPT, and so an upper bound after the modification also holds for the original scenario.

In what follows, let $k = |T|$ be the number of teams and $m = |S|$ be the number of tasks. Now focus on any of the additional wastes w , which was a waste of s by player i in team t . We can charge this waste to k hits as follows:

1. Task s was already covered in team t when i appeared; hence, we can infer that a hit occurred for task s in team t in a previous arrival. We charge to that hit.
2. For every team $t' \neq t$:
 - a. either has s covered (a hit from a previous arrival),
 - b. or has some other task s covered, for which i received payoff in t (again a hit from a previous arrival). If not, player i would have chosen t' over t .

We charge to these hits in teams $t' \neq t$.

We now need to bound the maximum number of times that a hit can be charged in the above scheme. Whenever some hit h for some task s is charged for an additional waste w , one of the following is true:

- Another hit h' , on the same task s as h , is happening at the same time as w . This is true for charging arguments of the form (2b).
- w is a waste of the same task s that is a hit in h . This is true for charging arguments of the form (1) and (2a).

Hence, a hit h on task s may be charged in the above scheme only if, at the time of the charging, there is a hit h' on the same task s or a waste w of the same task s . We also note that if a player i incurs multiple wastes in her selected team t , then for each team $t' \neq t$ and for each of these wastes (that are labeled as additional), we can find a distinct hit to charge with an argument of the form (2a) or (2b).

It follows that the first hit on s can be charged at most $h_s - 1 + w_s^+ = h_s^* - 1$ times, the second hit on s can be charged at most $h_s^* - 2$ times, and so on. Recall that the total number of hits on task s in NASH is h_s . Therefore, the total number of times that a hit on s can be

43:12 Profit Sharing and Efficiency in Utility Games

charged, denoted χ_s , is upper bounded as

$$\chi_s \leq (h_s^* - 1) + (h_s^* - 2) + (h_s^* - 3) + \dots + (h_s^* - h_s) = \sum_{j=1}^{h_s} (h_s^* - j).$$

Then, it follows that the total number of times all hits are charged is upper bounded as follows, with U (resp., U^*) denoting the total utility of NASH (resp., OPT).

$$\begin{aligned} \sum_{s \in S} \chi_s &\leq \sum_{s \in S} \sum_{j=1}^{h_s} (h_s^* - j) = \sum_{s \in S} h_s h_s^* - \sum_{s \in S} \left(\frac{1}{2} h_s (h_s + 1) \right) = \sum_{s \in S} h_s h_s^* - \frac{1}{2} \sum_{s \in S} h_s^2 - \frac{U}{2} \\ &\leq \sqrt{\sum_{s \in S} h_s^2} \sqrt{\sum_{s \in S} h_s^{*2}} - \frac{1}{2} \sum_{s \in S} h_s^2 - \frac{U}{2} \quad (\text{Cauchy-Schwarz inequality}) \\ &\leq \sqrt{m} \frac{U}{m} \cdot \sqrt{m} \frac{U^*}{m} - \frac{1}{2} m \cdot \frac{U^2}{m^2} - \frac{U}{2} = \frac{U \cdot U^*}{m} - \frac{U^2}{2m} - \frac{U}{2}. \end{aligned} \quad (10)$$

Eqn. (10) follows from the following facts: (a) the sum of squares of m nonnegative numbers with a given sum (here the sum of all h_s is U and the sum of all h_s^* is U^*) is minimized when they are all equal, and (b) the expression is decreasing as a function of the sum of all h_s and as a function of the sum of all h_s^* .

We also know that the total number of times a hit is charged is k times the number of additional wastes. Hence,

$$\sum_{s \in S} \chi_s = k \sum_{s \in S} w_s^+ = k \sum_{s \in S} (w_s - w_s^*) = k \sum_{s \in S} (h_s^* - h_s) = k(U^* - U). \quad (11)$$

From (10) and (11) we get that:

$$\frac{U \cdot U^*}{m} - \frac{U^2}{2m} - \frac{U}{2} \geq k(U^* - U) \quad (12)$$

Now let $\gamma = \frac{U}{U^*}$. Note that upper bounding $\frac{1}{\gamma}$ gives an upper bound on the POA. Substituting in (12) and using the fact that $U^* \leq mk$, we get

$$-\frac{k}{2} \gamma^2 + \frac{4k-1}{2} \gamma - k \geq 0$$

Since, by definition, $\gamma \in [0, 1]$, the expression is increasing in γ and, hence, for the inequality to hold, it must be the case that γ is greater than or equal to the unique root in $[0, 1]$. This gives the following upper bound for the POA:

$$\frac{U^*}{U} \leq \frac{2k}{4k-1 - \sqrt{8k^2 - 8k + 1}}.$$

This is increasing in k and as k goes to ∞ , the limit is $1 + \frac{1}{\sqrt{2}}$. This completes the proof.

4 Price of Stability

Studying the efficiency of the best NE in our setting is an interesting direction. The more optimistic metric that corresponds to the price of anarchy in this framework is the *price of stability*, i.e., the worst case ratio of the efficiency in OPT over the efficiency in the best NE. We conclude the paper with a brief discussion on the topic. For marginal loss profit sharing, we observe that any beneficial unilateral deviation also improves the social objective,

hence, OPT is also a NE and the price of stability is 1. For marginal gain profit sharing, it is possible to take any given NE, NASH, and modify the instance so that NASH becomes the unique NE in the modified instance. The modification is performed by means of making a very large number of copies of each task and introducing new unit tasks for tie-breaking purposes. Then, we get that the price of stability for marginal loss profit sharing is equal to the price of anarchy. We omit the exact details of this modification process. In contrast to the two previous models, determining the price of stability for egalitarian profit sharing appears to be a challenging question that invites future research.

References

- 1 Y. Bachrach, V. Syrgkanis, and M. Vojnovic. Incentives and efficiency in uncertain collaborative environments. In *WINE*, pages 26–39, 2013.
- 2 G. Calinescu, C. Chekuri, M. Pal, and J. Vondrak. Maximizing a submodular set function subject to a matroid constraint (extended abstract). In *IPCO*, pages 182–196, 2007.
- 3 H. Chen, T. Roughgarden, and G. Valiant. Designing network protocols for good equilibria. *SIAM Journal on Computing*, 39(5):1799–1832, 2010.
- 4 S. Dobzinski and M. Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *SODA*, pages 1064–1073, 2006.
- 5 Y. Filmus and J. Ward. The Power of Local Search: Maximum Coverage over a Matroid. In *STACS*, pages 601–612, 2012.
- 6 Y. Filmus and J. Ward. Monotone submodular maximization over a matroid via non-oblivious local search. *SIAM Journal on Computing*, 43(2):514–542, 2014.
- 7 M. Gairing. Covering games: Approximation through non-cooperation. In *WINE*, pages 184–195, 2009.
- 8 V. Gkatzelis, K. Kollias, and T. Roughgarden. Optimal cost-sharing in weighted congestion games. In *WINE*, 2014.
- 9 T. Harks and K. Miller. The worst-case efficiency of cost sharing methods in resource allocation games. *Operations Research*, 59(6):1491–1503, 2011.
- 10 E. Kalai and D. Samet. On weighted Shapley values. *International Journal of Game Theory*, 16(3):205–222, 1987.
- 11 S. Khot, R. J. Lipton, E. Markakis, and A. Mehta. Inapproximability results for combinatorial auctions with submodular utility functions. *Algorithmica*, 52(1):3–18, 2008.
- 12 K. Kollias and T. Roughgarden. Restoring pure equilibria to weighted congestion games. In *ICALP*, 2011.
- 13 J. R. Marden and T. Roughgarden. Generalized efficiency bounds in distributed resource allocation. In *CDC*, pages 2233–2238. IEEE, 2010.
- 14 J. R. Marden and A. Wierman. Distributed welfare games. *Operations Research*, 61(1):155–168, 2013.
- 15 D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- 16 I. Post, M. Kapralov, and J. Vondrak. Online submodular welfare maximization: Greedy is optimal. In *SODA*, pages 1216–1225, 2013.
- 17 R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- 18 T. Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC*, pages 513–522, 2009.
- 19 L. S. Shapley. *Additive and Non-Additive Set Functions*. PhD thesis, Department of Mathematics, Princeton University, 1953.

43:14 Profit Sharing and Efficiency in Utility Games

- 20 A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *FOCS*, 2002.