Learning to Bid in Proportional Allocation Auctions with Budget Constraints

Younes Ben Mazziane
LIA, Avignon University
Avignon, France
firstname.lastname@univ-avignon.fr

Cleque-marlain Mboulou-Moutoubi
LIA, Avignon University
Avignon, France
firstname.lastname@univ-avignon.fr

Francesco De Pellegrini
LIA, Avignon University
Avignon, France
firstname.lastname@univ-avignon.fr

Eitan Altman
LIA, Avignon University, Avignon, France
INRIA, Valbonne, France
firstname.lastname@inria.fr

***Abstract*—The *Kelly* or *proportional allocation* mechanism is a simple and efficient auction-based decentralized resource allocation scheme that distributes an infinitely divisible resource proportionally to the agents' bids. When agents are aware of the allocation mechanism, their interactions form a game. The properties of its Nash equilibria are well understood under the simplifying assumption of unbounded budgets. In this paper, we analyze the game in a more realistic budget-constrained setting, motivated by its optimality in terms of the liquid price of anarchy (LPoA). Specifically, we establish a sufficient condition for the uniqueness of the Nash equilibrium and design a distributed sequential learning procedure that provably converges to the equilibrium. In particular, our sufficient condition holds when the payoff functions of the agents are of the proportional fair type in the allocated fraction. Finally, extensive numerical experiments shed light on the interplay between the heterogeneity of the payoff functions and the agents' budgets.**

***Index Terms*—Kelly mechanism, decentralized resource allocation, auctions, game theory**

I. Introduction

Decentralized resource allocation in large-scale systems is a fundamental problem extensively studied in network economics [10]. In this context, a resource owner seeks to distribute resources among multiple agents to optimize an objective, such as maximizing *social welfare*, namely, the aggregate net benefit of the agents, or its revenue. It is standard to assume that the resource owner may have partial or lack information about the agents' utilities or preferences. Instead, they depend on signals [13] provided by the agents, such as declared valuations, willingness to pay, or other indirect indicators of agents' preferences. However, agents may often act selfishly and even strategically in order to maximize their benefits. This introduces additional complexity to the decision-making process of the resource owner.

This challenge is actually prevalent in various technological domains, including bandwidth allocation in communication networks [6], [14], task scheduling in cloud computing [18], energy distribution in smart grids [16], and pricing mechanisms in shared transportation systems [20].

This work was supported by the Horizon Europe CoGNETs project (Grant Agreement Number 101135930).

The *Kelly* or *proportional allocation* mechanism stands out among decentralized resource allocation mechanisms for its simplicity and efficiency [10], [12]. In its basic form, agents submit bids to secure shares of a finite, infinitely divisible resource, with allocations distributed proportionally to their bids. This establishes a uniform price per resource unit derived from the total sum of bids and the overall resource supply. In a generalized formulation [13], each user's allocation is determined by a *weighting function* of its bid, enabling diverse allocation strategies such as non-linear weighting of bids or user-specific pricing. In particular, the classic Kelly mechanism arises when this weighting function is simply the identity for all agents.

Originally, the efficiency of the Kelly mechanism has been demonstrated under the assumption that users have unlimited budgets, an assumption that may not align with real-world scenarios. Indeed, in Kelly's seminal paper [12], the mechanism appears as a scheme that maximizes social welfare assuming selfish agents who act as *price takers*, i.e., agents that do not anticipate how their bids affect the market price. However, when agents do anticipate the mechanism behavior, i.e., they are so-called *price-anticipating*, their interaction becomes a continuous-strategy competitive game. In such a setting, existence, uniqueness, and characterization of the Nash equilibrium have been established [8], and the resulting allocation is guaranteed to achieve at least $3/4$ of the optimal social welfare [11]. To further reduce the efficiency loss, modified Kelly-type mechanisms incorporating specific *weighting functions*—such as those analyzed in [13], [19]—can attain optimal social welfare, effectively bridging the efficiency gap in the original mechanism.

The aforementioned works established the Kelly mechanism as a promising approach to decentralized allocation under the simplifying assumption of unlimited budgets. More recently, researchers have investigated the more realistic setting in which price-anticipating users with limited budgets engage in a game induced by the Kelly mechanism [3]–[5], [17]. In this context, a more appropriate goal than conventional social welfare is referred to as *Liquid Social Welfare* (LSW).

ISBN 978-3-903176-73-7

LSW takes into consideration budget constraints by limiting an agent's utility to their maximum budget. The rationale for adopting this metric instead of social welfare is elaborated in [1]. Notably, [4] provides a lower bound on the *Liquid Price of Anarchy* (LPoA), defined as the ratio between the optimal LSW and the worst-case LSW at equilibrium across all possible equilibria. For any decentralized resource allocation mechanism with n agents, the LPoA lower bound is $2 - 1/n$. The Kelly mechanism achieves an LPoA of 2, making its worst case performance asymptotically optimal in the budgeted setting.

The optimality of the Kelly mechanism in budget-constrained settings in terms of LPoA raises several critical questions. Notably, LPoA optimality is achieved at an equilibrium of the game induced among the agents by the Kelly mechanism. However, our understanding of this Nash equilibrium remains limited. Key questions include: Is the equilibrium unique? Does it consist of a finite or infinite set of points? How does it depend on the agents' utility functions and on their budgets? Under what conditions does the system converge to this equilibrium, thereby achieving the optimal LSW?

To our knowledge, these questions have only been partially answered in a particular setting. Namely, for a class of agents whose utilities increase linearly with the fraction of the received resource and where the weighting function is the identity [7]. Specifically, [7] shows that, in this restricted setting, the equilibrium is unique and proposes a revision strategy that enables the system to converge to this unique equilibrium. However, more realistic settings often involve agents with diminishing returns, where the marginal utility of additional resources decreases as the amount received increases. This renders the linear utility model inadequate for many practical scenarios.

Contributions

This paper makes two main contributions. First, we derive a sufficient condition on the utility functions of the game induced by the Kelly mechanism, under general *weighting functions*, that guarantees the uniqueness of the Nash equilibrium. This result is resumed in Theorem 2. Second, we propose a general class of distributed learning algorithms in Alg. 1 and prove that, under a condition similar to, but weaker than the aforementioned sufficient condition, these algorithms drive the system toward the unique Nash equilibrium. This result is resumed in Theorem 3. Interestingly, both sufficient conditions reduce to verifying that the maximum of a specific continuously differentiable function, given in closed form in (14), is strictly negative. In particular, as shown in Lemma 3, both conditions hold when the utility functions grow logarithmically with the fraction of the resource received, and they hold for arbitrary convex, increasing weighting functions.

A central concept common to the proofs of Theorems 2 and 3 is what we term *Strong Diagonal Strict Concavity* (**SDSC**) in Definition 2. **SDSC** requires that the pseudo-Hessian—a matrix of second partial derivatives of the payoff functions—is negative definite for every possible set of agent actions. **SDSC** is a crucial property for establishing both the uniqueness of the Nash equilibrium [2] and convergence toward it [15]. However, directly verifying **SDSC** is challenging because the eigenvalue expressions of the Hessian generally lack closed-form expressions and the condition must be verified over the entire space of possible multi-strategies. Our sufficient conditions ensure that **SDSC** holds, thereby simplifying the verification process to merely checking that the maximum of a closed-form function is strictly negative.

Finally, we complement our results with numerical simulations that provide insights into the interdependence of the Nash equilibrium, e.g., the corresponding LPoA, the players budgets, and their utilities.

Road map

The rest of the paper is organized as follows. In Section II, we formally introduce the Kelly mechanism and its underlying game. Section III explores the connection between the **SDSC** property and both the uniqueness and learnability of the Nash equilibrium, deriving a sufficient condition for it to hold. Section IV identifies a class of utility functions for which the Nash equilibrium is unique and describes a class of distributed learning algorithms able to drive the system toward such equilibrium. Section V reports numerical simulations, and we conclude the paper in Section VI.

II. Problem formulation

We consider the allocation of a unit-sized divisible resource among n users according to the general allocation mechanism proposed in [13], which extends the proportionally fair Kelly mechanism introduced in [12].

Bidding. Each agent i submits a bid b_i that must be at least a fixed positive constant ϵ_i, that is, $b_i \geq \epsilon_i > 0$.

Allocation: Based on these bids, the resource owner allocates fractions $x_i(\mathbf{b})$ of the resource according to

$$x_i(\mathbf{b}) = \begin{cases} \frac{w_i(b_i)}{\sum_{j=1}^{n} w_j(b_j)+\delta}, & \text{if } w_i(b_i) > 0, \\ 0, & \text{otherwise,} \end{cases} \tag{1}$$

where $w_i : \mathbb{R}^+ \to \mathbb{R}^+$ are continuous, increasing functions governing how resources are distributed, and $\delta \geq 0$ is a reservation parameter [13]. If w_i is the identity function, this mechanism reduces to the classic Kelly mechanism.

Following [13], define $z_i = w_i(b_i)$, which can be viewed as the signal submitted by the agents, with the required payment given by the function $p_i : \mathbb{R}^+ \to \mathbb{R}^+$ defined as

$$p_i(z_i) \triangleq w_i^{-1}(z_i), \tag{2}$$

where w_i^{-1} is the inverse of w_i. Under this change of variable, the allocation rule becomes

$$x_i(\mathbf{z}) = \begin{cases} \frac{z_i}{\sum_{j=1}^{n} z_j+\delta} & \text{if } z_i > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

Note that both formulations are equivalent in the sense that the allocated fraction corresponding to a given payment is the same in each setting. For simplicity, we will focus on the second formulation using $\mathbf{z}$ in the remainder of the paper.

Each agent has a valuation function $V_i : [0,1] \to \mathbb{R}^n_{\geq 0}$, where $V_i(x_i)$ quantifies monetary benefit of acquiring a fraction x_i of the resource. The payoff function for agent i, denoted φ_i, is defined as the value the agent derives from the allocated fraction minus the payment, i.e.,

$$\varphi_i(\mathbf{z}) \triangleq V_i(x_i(\mathbf{z})) - p_i(z_i). \tag{4}$$

We study the budgeted Kelly mechanism, where each agent has a budget constraint c_i, i.e., $p_i(z_i) \leq c_i$.

When agents are aware that the Kelly mechanism governs resource allocation, the interaction between them forms a competitive game $\mathcal{G}$ with payoff functions $\boldsymbol{\varphi} = (\varphi_i)_{i\in[n]}$ and an action space constrained by budgets, denoted $\mathcal{R} = \mathcal{R}_1 \times \ldots \times \mathcal{R}_n$, where $\tilde{\epsilon}_i = p_i^{-1}(\epsilon_i)$ and $\tilde{c}_i = p_i^{-1}(c_i)$. We make the following assumptions about the functions V_i and p_i.

Assumption 1. *Over the domain $[0,1]$, V_i is strictly increasing, concave, and twice continuously differentiable ($V_i \in \mathcal{C}^2([0,1])$).*

Assumption 2. *Over the domain $\mathbb{R}^n_{\geq 0}$, p_i is convex, increasing with respect to z_i for any i, and twice continuously differentiable, i.e., $\mathbf{p} \in \mathcal{C}^2(\mathbb{R}^n_{\geq 0})$.*

Note that the above assumptions are standard [10], [13]. Moreover, it is natural for V_i and p_i to be increasing. Agents gain greater utility from receiving a larger share of the resource, and submitting a higher signal leads to a larger payment.

Let $(z_i, \mathbf{z}_{-i})$ denote the vector where agent i submits a bid z_i, while the other agents submit bids $\mathbf{z}_{-i}$.

Definition 1 (Nash Equilibrium)**.** *A strategy profile $\mathbf{z}^* = (z_1^*, z_2^*, \ldots, z_n^*) \in \mathcal{R}$ is a* Nash Equilibrium (NE) *if, for every player $i \in [n]$,*

$$\varphi_i(z_i^*, \mathbf{z}_{-i}^*) \geq \varphi_i(z_i, \mathbf{z}_{-i}^*) \quad \forall z_i \in \mathcal{R}_i \tag{5}$$

As a consequence of Assumptions 1 and 2, the function φ_i is concave in its i-th component and twice continuously differentiable on $\mathbb{R}^n_{\geq 0}$, and the actions set $\mathcal{R}$ is non empty, closed, bounded, and convex. Existence of a Nash Equilibrium (NE) of the game $\mathcal{G}$ follows by [2, Thm. 1].

Theorem 1. *The set of Nash equillibria of $\mathcal{G}$, denoted $NE(\mathcal{G})$ is non-empty, i.e., $NE(\mathcal{G}) \neq \emptyset$.*

A Nash equilibrium point is important as it represents a *stable* state where no player has an incentive to deviate, thereby serving as a reference for analyzing the system's behavior. For example, prior work has investigated the optimality of the system's social welfare, namely $\sum_{i=1}^N V_i(x_i)$, at an equilibrium point in the no-budget setting [11], and the *liquid social welfare*, denoted by LSW, in the budgeted setting. Formally, LSW is defined as,

$$\text{LSW}(\mathbf{x}) \triangleq \sum_{i=1}^n \min\left(V_i(x_i), c_i\right) : \sum_{i=1}^n x_i \leq 1. \tag{6}$$

Notably, if $\mathbf{z}^* \in NE(\mathcal{G})$, then $\text{LSW}(\mathbf{x}(\mathbf{z}^*)) \geq \frac{1}{2}\text{LSW}^*$, where LSW^* is the maximal liquid social welfare across all possible allocations [4, Thm. 5.1]. Moreover, this worst-case lower bound is asymptotically optimal in n for any resource allocation mechanism [4, Thm. 3.1], and it is attained by the Kelly mechanism.

III. Strong Diagonal Strict Concavity

For a vector $\mathbf{r} \in \mathbb{R}^n_{>0}$, **SDSC** is defined in terms of the $n \times n$ matrix $\mathbf{H}_\mathbf{r}(\mathbf{z})$, whose (i,j)-entry is given by

$$(\mathbf{H}_\mathbf{r}(\mathbf{z}))_{i,j} \triangleq r_i\, \partial^2_{i,j}\varphi_i(\mathbf{z}) + r_j\, \partial^2_{j,i}\varphi_j(\mathbf{z}), \tag{7}$$

where the partial derivative is taken with respect to the actions of player i and j, i.e., z_i and z_j, respectively.

Definition 2 (Strong Diagonal Strict Concavity)**.** *The game $\mathcal{G}$ satisfy* Strong Diagonal Strict Concavity *in $\mathbf{r}$ if and only if the matrix $\mathbf{H}_\mathbf{r}(\mathbf{z})$ is negative definite for all $\mathbf{z} \in \mathcal{R}$:*

$$\max_{\mathbf{v}:\|\mathbf{v}\|=1} \{\mathbf{v}^\intercal \mathbf{H}_\mathbf{r}(\mathbf{z})\mathbf{v}\} < 0. \tag{8}$$

*In this case, we write that $\mathcal{G}$ satisfies **SDSC**$(\mathbf{r})$.*

In this section, we derive a sufficient condition that simplifies verifying whether the game $\mathcal{G}$ satisfies **SDSC**$(\mathbf{r})$. Identifying this property reveals important characteristics of $\mathcal{G}$. Specifically, if there exists an $\mathbf{r}$ such that $\mathcal{G}$ satisfies **SDSC**$(\mathbf{r})$, then the game is *diagonally strictly concave* with respect to $\mathbf{r}$ [2, Thm. 6], guaranteeing the uniqueness of the Nash equilibrium [2, Thm. 2]. Furthermore, if $\mathcal{G}$ satisfies **SDSC**$(\mathbf{1})$, then it is *globally variationally stable* [15, Proposition 2.8]. This stability property enables adapting online learning algorithms—originally developed to minimize the *regret metric* [9]—to drive the system toward the unique Nash equilibrium [15, Thm. 4.6]. We leverage these insights to establish uniqueness in Theorem 2, propose a decentralized user bidding strategy in Algorithm 1, and prove its convergence to the Nash equilibrium in Theorem 3.

Given the central role of **SDSC**, we now examine the matrix $\mathbf{H}_\mathbf{r}(\mathbf{z})$ to derive a sufficient condition for the **SDSC** property. The second-order partial derivatives of the payoff functions $\partial^2_{i,j}\varphi_i$ can be written as

$$\partial^2_{i,j}\varphi_i(\mathbf{z}) = \partial^2_{i,j}U_i(\mathbf{z}) - \ddot{p}_i(z_i), \tag{9}$$

where $U_i(\mathbf{z}) = V_i(x_i(\mathbf{z}))$, and $\ddot{p}_i$ is the second derivative of p_i. Note that p_i is convex, which implies it contributes to the matrix $\mathbf{H}_\mathbf{r}(\mathbf{z})$ as a diagonal matrix whose entries are $-2r_i\ddot{p}_i(z_i) \leq 0$ and thus is negative semi-definite. Consequently, any failure of $\mathbf{H}_\mathbf{r}(\mathbf{z})$ to be negative definite would stem from the matrix formed by $\partial^2_{i,j}U_i(\mathbf{z})$. Lemma 1 shows that this matrix has a particular structure.

Lemma 1. *The partial derivatives of U_i verify,*

$$m(z)\partial^2_{i,j}U_i(\mathbf{z}) = \begin{cases} f_i(x_i(\mathbf{z})), & \text{if } i=j, \\ g_i(x_i(\mathbf{z})), & \text{otherwise,} \end{cases} \tag{10}$$

where:

$$f_i(x_i) = (1-x_i)^2\ddot{V}_i(x_i) - 2(1-x_i)\dot{V}_i(x_i), \tag{11}$$

$$g_i(x_i) = -x_i(1-x_i)\ddot{V}_i(x_i) + (2x_i-1)\dot{V}_i(x_i). \tag{12}$$

Here, $\dot{V}_i(\cdot)$ and $\ddot{V}_i(\cdot)$ denote the first and second derivatives of V_i, respectively, and $m(\mathbf{z}) = (\delta + \sum_{i=1}^n z_i)^2$.

The proof is reported in Appendix A. Lemma 1 shows that the matrix of second partial derivatives of the functions U_i can be rewritten using a change of variable from $\mathbf{z} \in \mathcal{R}$ to $\mathbf{x}(\mathbf{z})$, where the set Δ is defined as the image of $\mathcal{R}$ under mapping $\mathbf{z} \mapsto \mathbf{x}(\mathbf{z})$ as in (3). Moreover, this reformulated matrix can be decomposed into the sum of a diagonal matrix, with entries $f_i(x_i) - g_i(x_i)$, and a rank-one matrix, whose entries are given by $g_i(x_i)$. It is easy to show then that Δ verifies,

$$\Delta \subset \left\{ \mathbf{x} \in [0,1]^n : \; x_i \geq \frac{\tilde{\epsilon}_i}{\tilde{\epsilon}_i + C_i + \delta}, \; \sum_{i=1}^n x_i \leq \frac{C}{C+\delta} \right\}, \tag{13}$$

where $C = \sum_{i=1}^n \tilde{c}_i$ and $C_i = \sum_{j\neq i} \tilde{c}_i$. Let $\mathcal{F}_\mathbf{r}$ be the family of *games* $\mathcal{G}$ with payoff functions φ in (4) and action set in $\mathcal{R}$, such that $\forall \mathbf{x} \in \Delta$,

$$\psi(\mathbf{x}) < 1 : \; \psi(\mathbf{x}) \triangleq \left(\sum_{i=1}^n \frac{r_i g_i(x_i)^2}{k_i(x_i)} \right) \left(\sum_{i=1}^n \frac{1}{r_i k_i(x_i)} \right), \tag{14}$$

where $k_i(x_i) = g_i(x_i) - f_i(x_i) + \delta^2 L_i$ and $L_i = \min_{z_i \in \mathcal{R}_i} \ddot{p}_i(z_i)$.

Lemma 2 leverages the particular structure of the matrix in Lemma 1 to derive a sufficient condition for **SDSC**.

Lemma 2. $\mathcal{G} \in \mathcal{F}_\mathbf{r} \implies \mathcal{G}$ *satisfies* ***SDSC***$(\mathbf{r})$.

Proof: We assume that $\mathcal{G} \in \mathcal{F}_\mathbf{r}$ and prove that the matrix $\mathbf{H}_\mathbf{r}(\mathbf{z})$ is negative definite for every $\mathbf{z} \in \mathcal{R}$. To this end, we introduce the following notation. We define the vectors $\mathbf{g} = (g_i(x_i(\mathbf{z})))_{i=1}^n$ and $\mathbf{f} = (f_i(x_i(\mathbf{z})))_{i=1}^n$, and set $\tilde{k}_i(\mathbf{z}) = g_i(x_i(\mathbf{z})) - f_i(x_i(\mathbf{z})) + m(\mathbf{z})\ddot{p}_i(z_i)$, so that $\tilde{\boldsymbol{k}} = (\tilde{k}_i(\mathbf{z}))_{i=1}^n$. Here, $\odot$ denotes the Hadamard (elementwise) product, $\mathbf{1}$ is the vector of ones, and $\Lambda(\mathbf{a})$ represents the diagonal matrix whose diagonal entries are the components of $\mathbf{a}$.

Thanks to Lemma 1, the matrix can be expressed as

$$m(\mathbf{z})(\mathbf{H}_r(\mathbf{z}))_{i,j} = \begin{cases} 2r_i \left(f_i(x_i(\mathbf{z})) - m(\mathbf{z})\, \ddot{p}_i(z_i) \right), & \text{if } i=j, \\ r_i\, g_i(x_i(\mathbf{z})) + r_j\, g_j(x_j(\mathbf{z})), & \text{if } i \neq j. \end{cases} \tag{15}$$

Thus, the matrix $\mathbf{H}_r(\mathbf{z})$ can be written compactly as

$$\mathbf{H}_\mathbf{r}(\mathbf{z}) = -2\,\Lambda\Big(\tilde{\boldsymbol{k}} \odot \mathbf{r}\Big) + \Big(\mathbf{g} \odot \mathbf{r}\Big)\mathbf{1}^\intercal + \mathbf{1}\Big(\mathbf{g} \odot \mathbf{r}\Big)^\intercal. \tag{16}$$

Notice that $\tilde{k}_i(\mathbf{z}) \geq k_i(x_i(\mathbf{z}))$. Moreover, replacing by the expressions of g_i and f_i in (12) and (11), we obtain,

$$k_i(x_i) = (x_i - 1)\,\ddot{V}_i(x_i) + \dot{V}_i(x_i) + \delta^2 \min_{z_i \in \mathcal{R}_i} \ddot{p}_i(z_i). \tag{17}$$

By Assumptions 1 and 2, we have $\dot{V}_i > 0$, $\ddot{V}_i < 0$, and $\ddot{p}_i > 0$. Moreover, since $x_i(\mathbf{z}) < 1$ (as $\mathbf{x} \in \Delta$), it follows that $k_i(x_i(\mathbf{z})) > 0$.

Let $\mathbf{v} \in \mathbb{R}^N$. Using the expression of $\mathbf{H}_r(\mathbf{z})$, we have

$$\begin{aligned} \mathbf{v}^T\mathbf{H}_\mathbf{r}(\mathbf{z})\mathbf{v} &= -2\left(\sum_{i=1}^n v_i^2\, r_i\, \tilde{k}_i(\mathbf{z}) - \langle \mathbf{v},\, \mathbf{g}\odot\mathbf{r}\rangle \langle \mathbf{v},\, \mathbf{1}\rangle \right) \\ &\leq -2\left(\sum_{i=1}^n v_i^2\, r_i\, k_i(x_i(\mathbf{z})) - \langle \mathbf{v},\, \mathbf{g}\odot\mathbf{r}\rangle \langle \mathbf{v},\, \mathbf{1}\rangle \right), \end{aligned} \tag{18}$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product.

We now rewrite the inner products in (18) by noting that, for any vector $\mathbf{a}$, the notation $\sqrt{\mathbf{a}}$ represents the vector obtained by taking the square root of each component of $\mathbf{a}$, and the ratio of two vectors is computed elementwise. Define the vector $\mathbf{k} = (k_i(x_i(\mathbf{z})))_{i=1}^N$. We have that,

$$\langle \mathbf{v},\, \mathbf{g}\odot\mathbf{r}\rangle = \left\langle \mathbf{v} \odot \sqrt{\mathbf{k}\odot\mathbf{r}},\, \frac{\sqrt{\mathbf{r}}\odot\mathbf{g}}{\sqrt{\mathbf{k}}} \right\rangle, \tag{19}$$

$$\langle \mathbf{v},\, \mathbf{1}\rangle = \left\langle \mathbf{v} \odot \sqrt{\mathbf{k}\odot\mathbf{r}},\, \frac{\mathbf{1}}{\sqrt{\mathbf{k}\odot\mathbf{r}}} \right\rangle. \tag{20}$$

Applying the Cauchy-Schwarz inequality to these expressions and substituting into (18) yields

$$\begin{aligned} \mathbf{v}^T\mathbf{H}_\mathbf{r}(\mathbf{z})\mathbf{v} &\leq -2\sum_{i=1}^n v_i^2\, r_i\, k_i(x_i(\mathbf{z})) + 2\left(\sum_{i=1}^n v_i^2 r_i k_i(x_i(\mathbf{z})) \right) \\ &\times \sqrt{\sum_{i=1}^n \frac{r_i\, g_i(x_i(\mathbf{z}))^2}{k_i(x_i(\mathbf{z}))}} \sqrt{\sum_{i=1}^n \frac{1}{r_i k_i(x_i(\mathbf{z}))}} \\ &= \underbrace{2\sum_{i=1}^n v_i^2 r_i k_i(x_i(\mathbf{z}))}_{>0} \Bigg(-1 \\ &\qquad + \sqrt{\sum_{i=1}^n \frac{r_i g_i(x_i(\mathbf{z}))^2}{k_i(x_i(\mathbf{z}))} \sum_{i=1}^n \frac{1}{r_i k_i(x_i(\mathbf{z}))}} \Bigg). \end{aligned} \tag{21}$$

which concludes the proof. ■

Lemma 3. *Consider the game $\mathcal{G}$, with payoff functions φ in (4), $\delta > 0$, and action set $\mathcal{R}$. Let $V_i(x) = a_i \log(x) + d_i$, with $a_i > 0$, then,*

- *$\mathcal{G}$ satisfies **SDSC**$(\mathbf{r})$ for $r_i = a_i^{-1}$.*
- *if $\left(\max_i \frac{a_i^2}{a_i + \delta^2 L_i}\right)\left(\max_i \frac{1}{a_i+\delta^2 L_i}\right) < 1$ and $a_i > 3\delta^2 L_i$, then $\mathcal{G}$ satisfies **SDSC**$(\mathbf{1})$.*

Proof: Assume that $V_i(x) = a_i \log(x) + d_i$, $\delta > 0$, and $r_i = a_i^{-1}$. Straightforward calculations show that $g_i(x) = a_i$ and $k_i(x) = {a_i}/{x^2} + \delta^2 L_i$. Thus $\psi(\mathbf{x})$ verifies,

$$\psi(\mathbf{x}) = \left(\sum_{i=1}^{n} \frac{a_i}{{a_i}/{x_i^2} + \delta^2 L_i}\right)\left(\sum_{i=1}^{n} \frac{1}{{1}/{x_i^2} + \delta^2 a_i^{-1} L_i}\right) \tag{22}$$

$$\leq \left(\sum_{i=1}^{n} x_i^2\right)^2 \leq \left(\sum_{i=1}^{n} x_i\right)^4 \leq \left(\frac{C}{C+\delta}\right)^4 < 1. \tag{23}$$

Now, we assume that $r_i = 1$. In this case,

$$\psi(\mathbf{x}) = \left(\sum_{i=1}^{n} \frac{a_i^2 x_i^2}{a_i + \delta^2 L_i x_i^2}\right)\left(\sum_{i=1}^{n} \frac{x_i^2}{a_i + \delta^2 L_i x_i^2}\right). \tag{24}$$

One can verify that functions of the form ${x^2}/{(a+bx^2)}$ with $a > 0$, $b > 0$, and $a \geq 3b$ are convex over $[0,1]$. Given that, $a_i \geq 3\delta_i^2 L_i$ for any i, we deduce that ψ is the product of two convex functions. The maximum of a convex function over a convex set is attained at an extreme point. Moreover, given that $\mathbf{x} \in \Delta$, and Δ is a subset of the n dimensional simplex, we deduce that,

$$\psi(\mathbf{x}) \leq \left(\max_i \frac{a_i^2}{a_i + \delta^2 L_i}\right)\left(\max_i \frac{1}{a_i + \delta^2 L_i}\right) < 1, \tag{25}$$

which concludes the proof. ■

IV. Nash Equilibrium: Uniqueness and Learning

Lemma 2 guarantees that $\mathcal{G}$ statisfies $\mathbf{SDSC}(\mathbf{r})$ whenever $\mathcal{G} \in \mathcal{F}_{\mathbf{r}}$. This condition, in turn, allows us to apply [2, Thm. 2] and [2, Thm. 6], which together establish the uniqueness of the Nash equilibrium in Theorem 2.

Theorem 2. *If there exists $\mathbf{r} \in \mathbb{R}^n_{>0}$ such that $\mathcal{G} \in \mathcal{F}_{\mathbf{r}}$, then the set of Nash equilibria of $\mathcal{G}$ is a singleton, i.e., $|NE(\mathcal{G})| = 1$.*

We now consider a sequential variant of the game. At time step $t = 0$, all agents submit an initial vector of bids $\mathbf{z}(0)$, which is then broadcast by the resource owner to all agents, or, equivalently, the aggregated demand $m(z)$. At each step $t \geq 1$, the agents update (revise) their actions according to some strategy that might depend on all previous bids $(\mathbf{z}(s))_{s=0}^{t-1}$. Specifically, at step t, each agent i can choose a bid $z_i(t)$ subject to the budget constraint $p_i\big(z_i(t)\big) \leq c_i$ or equivalently, $z_i(t) \leq p_i^{-1}(c_i) = \tilde{c}_i$. The updated bids (or $m(z)$) are then broadcast to all agents again. Algorithm 1 introduces a general class of bidding strategies, parameterized by the partitioning of players into two disjoint groups, I_1 and I_2, where $I_1 \cup I_2 = [n]$ and $I_1 \cap I_2 = \emptyset$. For any player i, $y_i(t)$ denotes the *accumulated discounted gradient* of its utility with respect to its action. Formally, $y_i(0) = 0$ and for $t \geq 1$,

$$y_i(t) = y_i(t-1) + \eta_t \partial_i \varphi_i(\mathbf{z}(t-1)), \tag{26}$$

where $\mathbf{z}(t-1)$ represents the actions of all the players at step $t-1$ and stepsize $\eta_t = t^{-\beta}$ for $\beta \in (0,1]$. The players then convert this gradient information into a bid depending on their group. If $i \in I_1$, then, $z_i(t) = Q_i^{(1)}(y_i(t))$; otherwise, $z_i(t) = Q_i^{(2)}(y_i(t))$, where,

$$Q_i^{(1)}(y) = \max\left(\tilde{\epsilon}_i, \min\left(y, \tilde{c}_i\right)\right), \tag{27}$$

$$Q_i^{(2)}(y) = \max\left(\tilde{\epsilon}_i, \min\left(e^{y-1}, \tilde{c}_i\right)\right). \tag{28}$$

The above update rules bear similarity to the online gradient ascent framework of [21]. However, a key distinction arises in the update mechanism: while the latter defines updates $z_i(t) = Q_i^{(1)}\big(z_i(t-1) + \eta_t \partial_i \varphi_i(\mathbf{z}(t-1))\big)$, the update rule in Alg. 1 decouples the gradient updates. Specifically, gradients are first independently integrated into an auxiliary variable $y_i(t)$, which is subsequently projected as a bid via the transformation $z_i(t) = Q_i^{(1)}(y_i(t))$ for agents in I_1, and similarly $z_i(t) = Q_i^{(2)}(y_i(t))$ for agents in I_2.

Algorithm 1 Bidding strategy

1: **Input:** Initial bids $\{z_i(0)\}_{i=1}^n$, step size parameter β, number of iterations T, groups I_1 and I_2.
2: **Output** bidding strategy at each step: $(\mathbf{z}(t))_{t=1}^T$
3: **for** $t = 1$ to T **do**
4: $\quad \eta_t \leftarrow \frac{1}{t^\beta}$
5: $\quad y_i(t) \leftarrow y_i(t-1) + \eta_t \partial_i \varphi_i(\mathbf{z}(t-1)),\ \forall i \in [n]$
6: $\quad \forall i \in I_1,\ z_i(t) \leftarrow Q_i^{(1)}(y_i(t))$
7: $\quad \forall i \in I_2,\ z_i(t) \leftarrow Q_i^{(2)}(y_i(t))$
8: **end for**

Theorem 3. *If $\mathcal{G} \in \mathcal{F}_{\mathbf{1}}$ (see (14)), then the sequence $\mathbf{z}(t)$ generated by the update rule in Algorithm 1 converges to the unique Nash equilibrium of the game as $t \to \infty$.*

Proof: Let $\mathcal{X} = [\epsilon, c]$ and let $h \colon \mathcal{X} \to \mathbb{R}$ be a regularizer or penalty function. Define the operator Q by

$$Q(y) = \arg\max_{z \in \mathcal{X}} \left(\tilde{h}(y,z)\right) : \ \tilde{h}(y,z) = yz - h(z). \tag{29}$$

First, we will show in Lemma 4 that $Q_i^{(1)}$ and $Q_i^{(2)}$ correspond to this map Q when $h(z) = \frac{1}{2}z^2$ (quadratic regularizer) or $h(z) = z\log z$ (entropic regularizer), respectively, with $c = \tilde{c}_i$ and $\epsilon = \tilde{\epsilon}_i$. Hence, Algorithm 1 belongs to the class of Dual Averaging algorithms introduced in [15]. The convergence of these algorithms to a Nash equilibrium is contingent on two conditions stated in [15, Thm. 4.6].

Condition 1: Define

$$F(\ell, y) = h^*(y) - \tilde{h}(y,\ell),\ h^*(y) = \max_{z \in \mathcal{X}} \left\{\tilde{h}(y,z)\right\}. \tag{30}$$

Here, h^* is the convex conjugate of h. We require that if there is a sequence $\{y_n\} \subset \mathbb{R}$ such that $Q(y_n) \to \ell$, then $F(\ell, y_n) \to 0$. In fact, Lemma 4 (proof in Appendix B) ensures that this property holds for both the quadratic and entropic cases.

Lemma 4. *For the two regularizers below, and for $c = \tilde{c}_i$ and $\epsilon = \tilde{\epsilon}_i$, $Q(y)$ has closed-form expressions:*

- ***Quadratic:*** *if $h(z) = \frac{1}{2}z^2$, then $Q(y) = Q_i^{(1)}(y)$.*

- ***Entropic:*** *if $h(z) = z \log z$ with $h(0) = 0$, then $Q(y) = Q_i^{(2)}(y)$.*

Moreover, if there exists a sequence $\{y_n\} \subset \mathbb{R}$ with $Q(y_n) \to \ell$, then $F(\ell, y_n) \to 0$.

Condition 2: The game $\mathcal{G}$ must be globally variationally stable. A sufficient condition is that $\mathcal{G}$ satisfies $\mathbf{SDSC}(\mathbf{1})$ [15, Prop. 2.8]. Since $\mathcal{G} \in \mathcal{F}_{\mathbf{1}}$ by assumption, Theorem 2 implies $\mathcal{G}$ is indeed globally variationally stable.

As both conditions of [15, Thm. 4.6] are satisfied, the sequence $\mathbf{z}(t)$ converges to the unique Nash equilibrium. ■

Combining Theorems 2, 3, and Lemma 3, we prove Corollary 1.

Corollary 1. *The Nash equilibrium of the game $\mathcal{G}$ is unique for arbitrary weighting functions when $V_i(x) = a_i \log x + d_i$, with $a_i > 0$ and $\delta > 0$. Moreover, if $\left(\max_i \frac{a_i^2}{a_i + \delta^2 L_i}\right)\left(\max_i \frac{1}{a_i + \delta^2 L_i}\right) < 1$ and $a_i > 3\delta^2 L_i$, then Alg. 1 is guaranteed to converge to this unique equilibrium.*

In scenarios where agents only observe a noisy estimate of the aggregate bid $\sum_i z_i$, Alg. 1 may fail to converge to the game's equilibrium without additional assumptions on the noise in the gradient estimates. However, if in addition to the conditions of Theorem 3, the noisy gradients satisfy $\widetilde{\partial_i \phi_i} = \partial_i \phi_i + \zeta_i$ and ζ_i are zero mean random variables with bounded variance, then Alg. 1 converges *almost surely* to the unique equilibrium of the game [15, Thm. 4.7].

V. Numerical Simulations

In this section, we evaluate the efficiency of the Kelly mechanism in terms of the liquid price of anarchy (LPoA). We consider a class of games in $\mathcal{F}_{\mathbf{1}}$. For these games, the Nash equilibrium is unique (Theorem 2), and Algorithm 1 is guaranteed to converge to this equilibrium (Theorem 3). Formally, the LPoA is defined as

$$\mathrm{LPoA} \triangleq \frac{\mathrm{LSW}^*}{\min_{\mathbf{z} \in \mathrm{NE}(\mathcal{G})} \mathrm{LSW}(\mathbf{x}(\mathbf{z}))}, \tag{31}$$

where the optimal liquid social welfare is given by $\mathrm{LSW}^* = \max_{\mathbf{x} \in \Delta} \mathrm{LSW}(\mathbf{x})$, and the liquid social welfare LSW is defined in (6).

In our simulations, we set the number of players to $n = 20$, $\delta = 1$, require an entry bid of at least 1, i.e., $\epsilon_i = 1$, use a constant step size $\eta = 0.01$ in Alg. 1, and consider payoff functions of the form

$$\varphi_i(\mathbf{z}) = a_i \log\big(x_i(\mathbf{z})\big) + d_i - z_i, \tag{32}$$

with $a_i > 0$ and $d_i > 0$. We define d_i^0 as the smallest value of d_i for which $a_i \log\big(x_i(\mathbf{z})\big) + d_i \geq 0$, i.e., $d_i^0 = a_i \log\left(1 + \frac{\delta + \sum_{j \neq i} c_j}{\epsilon_i}\right)$. Indeed, whenever LSW is positive at an equilibrium point of the game, LPoA ≥ 1, and setting $d_i = d_i^0$ ensures that.

For this simulation setup, the computation of the LPoA requires two elements: (i) the equilibrium point of the game, and (ii) the maximizer of $\sum_{i=1}^n \min(a_i \log(x_i) + d_i, c_i)$ over the convex set Δ (see (13)). Observing that this optimization problem has a concave objective, we solve the KKT conditions via a bisection algorithm, thus obtaining LSW^*. We then compute the unique equilibrium of the game $\mathcal{G}$ using Algorithm 1, and finally deduce LPoA.

We consider three settings. First, we examine the impact of different regularizers among players, specifically the choice between a quadratic regularizer and an entropic regularizer. Second, we analyze the impact of heterogeneous utilities among players (i.e., different values of a_i). Finally, we evaluate the effect of varying budgets.

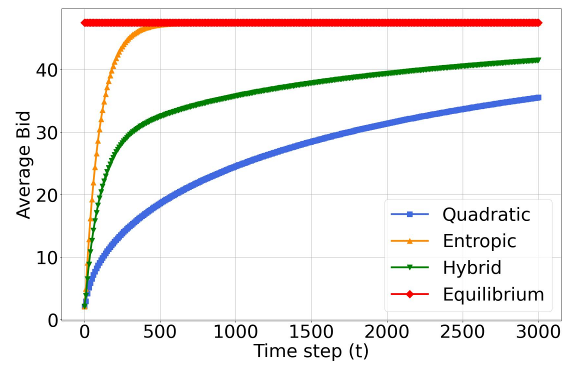

(a) Average Bid

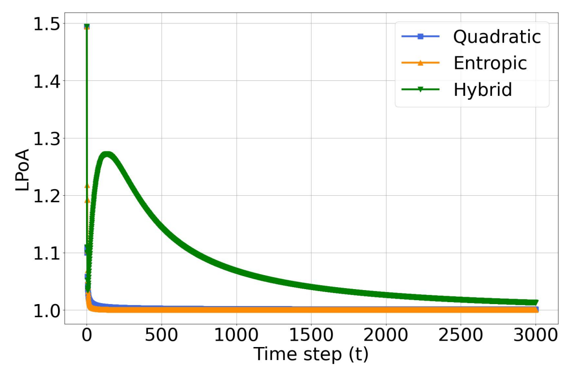

(b) LPoA

Fig. 1: LPoA and bidding dynamics for identical utilities and different regularizers: $c_i = 100$, $a_i = 50$, $d_i = 0.5 d_i^0$.

Impact of Strategy Revision Policies. In this scenario, all players have the same utility function and the same budget. Specifically, we set $a_i = a = 50$ and $c_i = c = 100$ for every player i. By a symmetry argument, all players share the same strategy at an equilibrium point. If $z^* \in (\epsilon_i, c_i)$ is an equilibrium of the game, then it satisfies the following quadratic equation:

$$n(z^*)^2 - z^*(a(n-1) - \delta) - a\delta = 0. \tag{33}$$

In Figure 1, we assess the influence of different regularizers, which correspond to distinct player revision policies: *(i)* when only the regularizer $\mathcal{Q}^{(1)}$ (*Quadratic*) is applied, *(ii)* when only the regularizer $\mathcal{Q}^{(2)}$ (*Entropic*) is applied, and *(iii)* when half of the players uses $\mathcal{Q}^{(1)}$ and the other half uses $\mathcal{Q}^{(2)}$ (*Hybrid*).

Figure 1a illustrates the average bid at each iteration, i.e., $\frac{1}{n}\sum_{i=1}^{n} z_i(t)$. The red curve represents the positive solution to (33). We observe that the *Entropic* regularizer strategy converges the fastest to the Nash equilibrium, followed by the *Hybrid* regularizer, and then by the *Quadratic* regularizer. However, in general, we cannot assert that one strategy is always faster than the others. For instance, for values of η larger than 0.01, the *Quadratic* regularizer converges faster, but the *Entropic* regularizer becomes slower. Figure 1b shows the evolution of the LPoA under different regularizers. While all regularizers attain near-optimal performance in the symmetric case, the *Hybrid*-regularizer strategy requires many iterations to achieve such performance. Indeed, in a symmetric setup, the LSW is maximized when the resource is uniformly allocated among the agents. We ascribe this behavior to the fact that, under the *Hybrid* regularizer, agents are more likely to submit different bids, which delays convergence to the optimal allocation.

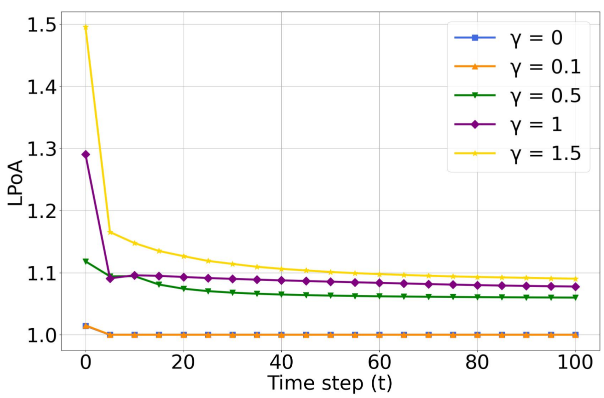

Fig. 2: LPoA with heterogeneous utilities: $d_i = 0.9d_i^0$.

Impact of Heterogeneous Utilities. In the second experiment, shown in Figure 2, players have heterogeneous utility functions defined by $a_i = 50 \cdot i^{-\gamma}$ for $\gamma \in \{0, 0.1, 0.5, 1, 1.5\}$, but they have identical budgets $c_i = 100$. We use the *Entropic* regularizer for all players. As γ increases, the heterogeneity among players grows, and the resulting LPoA decreases. This can be explained when $n = 2$, $\delta = 0$, $\epsilon_i = 0$, and for large enough budgets: the optimal allocation is $(x_1^*, x_2^*) = ({}^{1}/_{\rho+1}, {}^{\rho}/_{\rho+1})$, whereas the allocation at the Nash equilibrium is $(x_1^{\text{eq}}, x_2^{\text{eq}}) = ({}^{1}/_{\sqrt{\rho}+1}, {}^{\sqrt{\rho}}/_{\sqrt{\rho}+1})$, with $\rho = {}^{a_2}/_{a_1}$. One can then show that LPoA is increasing in ρ over $(0,1)$ and decreasing over $(1,+\infty)$ with $\rho^* = 1$ being the maximizer of LPoA.

Impact of Heterogeneous Budgets. Finally in the last experiment reported in Fig. 3, we describe the behavior of LPoA for players with heterogeneous budgets, namely $c_i = 100 \cdot i^{\mu}$ with $\mu \in \{0, 1, 1.25, 1.5\}$, and identical utilities ($a_i = a = 10$). As μ increases, the LPoA becomes worse.

This is expected because when agents have identical utilities, the uniform allocation is optimal and varying budgets across users introduces asymmetry to the allocation at the Nash equilibrium.

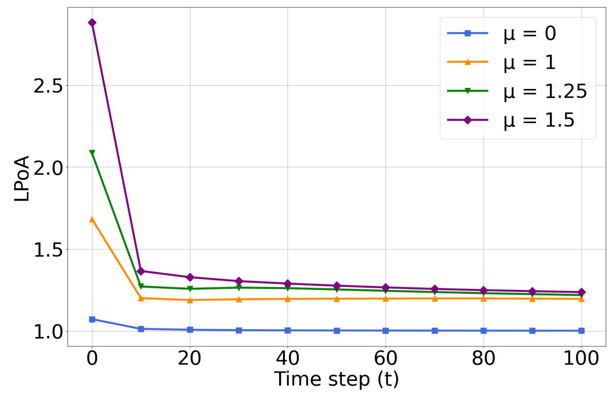

Fig. 3: LPoA with heterogenous budgets: $d_i = 0.9d_i^0$.

VI. Conclusion

In this paper, we studied the game that arises from the competition among budget-limited agents in an auction where an infinitely divisible resource is allocated proportionally to the bids. We established a sufficient condition for the uniqueness of the Nash equilibrium and designed a general class of distributed bidding algorithms that drive the system toward this equilibrium. Moreover, we leveraged these algorithms to shed light on the relationship between the efficiency of the resource allocation mechanism, measured by the liquid price of anarchy, and the heterogeneity of both the payoff functions and the budgets.

In future work, we plan to extend our analysis to more general payoff functions, such as the α-fair utilities. We also intend to explore a sequential setting in which agents do not always participate in the auction, as well as scenarios where agents bid simultaneously for more than one resource.

Appendix

A. Proof of Lemma 1

Proof: The calculation of the derivative of U_i with respect to z_i writes

$$\partial_i U_i(z_i) = \frac{\sum_{j\neq i}^{N} z_j + \delta}{\left(\sum_{l=1}^{N} z_l + \delta\right)^2} \dot{V}_i\left(\frac{z_i}{\sum_{l=1}^{N} z_l + \delta}\right). \tag{34}$$

The elements on the diagonal can be written as

$$\begin{aligned}\partial_{i,i}^2 U_i(\mathbf{z}) &= \frac{-2\sum_{l\neq i}^{N} z_i + \delta}{(\sum_{l=1}^{N} z_l + \delta)^3} \dot{V}_i\left(\frac{z_i}{\sum_{l=1}^{N} z_l + \delta}\right) \\ &+ \left(\frac{\sum_{l\neq i}^{N} z_l + \delta}{(\sum_{l=1}^{N} z_l + \delta)^2}\right)^2 \ddot{V}_i\left(\frac{z_i}{\sum_{l=1}^{N} z_l + \delta}\right) = \frac{f_i(x_i(\mathbf{z}))}{m(z)^2}\end{aligned} \tag{35}$$

In a similar fashion we obtain the off-diagonal entries

$$\partial_{i,j}^2 U_i(\mathbf{z}) = \frac{\sum_{l=1}^N z_l + \delta - 2\left(\sum_{l\neq i}^N z_l + \delta\right)}{\left(\sum_{l=1}^N z_l + \delta\right)^3} \dot{V}_i\left(\frac{z_i}{\sum_{l=1}^N z_l + \delta}\right) - \frac{\sum_{l\neq i}^N z_l + \delta}{\left(\sum_{l=1}^N z_l + \delta\right)^4} \cdot z_i \cdot \ddot{V}_i\left(\frac{z_i}{\sum_{l=1}^N z_l + \delta}\right) = \frac{g_i(x_i(\mathbf{z}))}{m(z)^2} \tag{36}$$

which concludes the calculation. ■

B. Proof of Lemma 4

Proof: Assume that $h(z) = \frac{1}{2}z^2$.

1. Compute $Q(y)$: The maximizer of the function $yz - \frac{1}{2}z^2 = -\frac{1}{2}(y-z)^2 + \frac{1}{2}y^2$ over $z \in [\epsilon, c]$ coincides with the minimizer of $(y-z)^2$, which is the projection of y onto the interval $[\epsilon, c]$ and thus Q coincides with $Q_i^{(1)}$ when $\epsilon = \tilde{\epsilon}_i$ and $c = \tilde{c}_i$.

2. Derive $h^*(y)$: Substitute $Q(y)$ into $yz - \frac{1}{2}z^2$:

$$h^*(y) = \begin{cases} \tilde{h}(y, \epsilon), & y \leq \epsilon, \\ \frac{1}{2}y^2, & 0 < y < c, \\ \tilde{h}(y, c), & y \geq c. \end{cases} \tag{37}$$

3. Prove $Q(y_n) \to l \implies F(l, y_n) \to 0$:

- **Case 1:** $l \in (\epsilon, c)$. Then $y_n \to l$, and $F(l, y_n) = \frac{1}{2}l^2 + \frac{1}{2}y_n^2 - y_n l \to 0$.
- **Case 2:** $l = c$ or $l = \epsilon$. Then, there exists n_0, such that $\forall n \geq n_0$, $y_n \geq c$ (or $y_n \leq \epsilon$ if $l = \epsilon$). Thus for such n, $F(c, y_n) = \frac{1}{2}c^2 + \left(y_n c - \frac{1}{2}c^2\right) - y_n c = 0$ and a similar reasoning holds for $l = \epsilon$.

Assume that $h(z) = z \ln z$.

1. Compute $Q(y)$. The maximizer of $yz - z\log z$ over $z \in [\epsilon, c]$ is found by setting the derivative $y - (\log z + 1)$ to zero, yielding $z = e^{y-1}$. Constraining $z \in [\epsilon, c]$, we obtain $Q(y) = Q_i^{(2)}(y)$ when $\epsilon = \tilde{\epsilon}_i$ and $c = \tilde{c}_i$.

2. Derive $h^*(y)$. Substitute $Q(y)$ into $h^*(y) = \max_{z\in[\epsilon,c]}(yz - z\log z)$:

- If $y \leq 1 + \log\epsilon$: $h^*(y) = \tilde{h}(y, \epsilon)$
- If $y \in (1 + \log\epsilon, 1 + \log c])$: $h^*(y) = e^{y-1}$.
- If $y \geq 1 + \log c$: $h^*(y) = \tilde{h}(y, c)$.

3. Prove $Q(y_n) \to l \implies F(l, y_n) \to 0$. Consider three cases:

- **Case 1:** $0 < l < c$. Here $y_n \to 1 + \log l$, leading to:

$$F(l, y_n) = l\log l + e^{y_n - 1} - y_n l \to= 0.$$

- **Case 2:** $l = \epsilon$ or $l = c$. Then, there exists n_0, such that $\forall n \geq n_0$, $y_n \geq c$ (or $y_n \leq \epsilon$ if $l = \epsilon$). Thus for such n, $F(c, y_n) = c\log c + (y_n c - c\log c) - y_n c = 0$, and a similar reasoning holds for $l = \epsilon$.

In all cases, $F(l, y_n) \to 0$, which concludes the proof. ■

References

[1] Y. Azar, M. Feldman, N. Gravin, and A. Roytman. Liquid price of anarchy. In *SAGT*, pages 3–15, 2017.

[2] J. B. Rosen. Existence and uniqueness of equilibrium points for concave N-person games. *Econometrica*, 33(3), July 1965.

[3] I. Caragiannis and A. A. Voudouris. Welfare guarantees for proportional allocations. *Theory Comput. Syst.*, 59, 2016.

[4] I. Caragiannis and A. A. Voudouris. The efficiency of resource allocation mechanisms for budget-constrained users. In *EC*, pages 681–698. ACM, 2018.

[5] G. Christodoulou, A. Sgouritsa, and B. Tang. On the efficiency of the proportional allocation mechanism for divisible resources. *Theory Comput. Syst.*, 59, 2016.

[6] M. Datar, E. Altman, F. D. Pellegrini, R. E. Azouzi, and C. Touati. A mechanism for price differentiation and slicing in wireless networks. In *WiOPT*, pages 121–128. IEEE, 2020.

[7] S. D'Oro, L. Galluccio, P. Mertikopoulos, G. Morabito, and S. Palazzo. Auction-based resource allocation in openflow multi-tenant networks. *Comput. Networks*, 115:29–41, 2017.

[8] B. Hajek and G. Gopalakrishnan. Do greedy autonomous systems make for a sensible internet? In *Presented at the Conference on Stochastic Networks*, 2002.

[9] E. Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.

[10] R. Johari. *Efficiency Loss in Market Mechanisms for Resource Allocation*. PhD thesis, Department of Electrical Engineering and Computer Science, 2004.

[11] R. Johari and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Math. Oper. Res.*, 29(3):407–435, 2004.

[12] F. Kelly. Charging and rate control for elastic traffic. *Eur. Trans. Telecommun.*, 8(1):33–37, 1997.

[13] R. T. Maheswaran and T. Basar. Efficient signal proportional allocation (ESPA) mechanisms: decentralized social welfare maximization for divisible resources. *IEEE JSAC.*, 24(5), 2006.

[14] P. Maillé and B. Tuffin. Multi-bid auctions for bandwidth allocation in communication networks. In *IEEE INFOCOM*, 2004.

[15] P. Mertikopoulos and Z. Zhou. Learning in games with continuous action sets and unknown payoff functions. *Math. Program.*, 173(1-2):465–507, 2019.

[16] W. Saad, Z. Han, H. V. Poor, and T. Başar. A noncooperative game for double auction-based energy trading between phevs and distribution grids. In *IEEE SmartGridComm*, 2011.

[17] V. Syrgkanis and É. Tardos. Composable and efficient mechanisms. In *STOC*, pages 211–220. ACM, 2013.

[18] X. Wang, Y. Sui, J. Wang, C. Yuen, and W. Wu. A distributed truthful auction mechanism for task allocation in mobile cloud computing. *IEEE Trans. Serv. Comput.*, 14(3):628–638, 2021.

[19] S. Yang and B. E. Hajek. Vcg-kelly mechanisms for allocation of divisible goods: Adapting VCG mechanisms to one-dimensional signals. *IEEE J. Sel. Areas Commun.*, 25(6):1237–1243, 2007.

[20] L. Zheng, P. Cheng, and L. Chen. Auction-based order dispatch and pricing in ridesharing. In *35th IEEE ICDE*. IEEE, 2019.

[21] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *ICML*, pages 928–936, 2003.