Nash Social Welfare Approximation for Strategic Agents*

Simina Brânzei
Hebrew U. of Jerusalem
simina.branzei@gmail.com

Vasilis Gkatzelis
U.C. Berkeley
gkatz@cs.berkeley.edu

Ruta Mehta
U.I. Urbana-Champaign
mehta.ruta@gmail.com

Abstract

The fair division of resources among strategic agents is an important age-old problem that has led to a rich body of literature. At the center of this literature lies the question of whether there exist mechanisms that can implement fair outcomes, despite the agents' strategic behavior. A fundamental objective function used for measuring fair outcomes is the *Nash social welfare* (NSW), mathematically defined as the geometric mean of the agents' values in a given allocation. This objective function is maximized by widely known solution concepts such as Nash bargaining and the competitive equilibrium with equal incomes.

In this work we focus on the question of (approximately) implementing this objective. We analyze two classical mechanisms, the Fisher market and the Trading Post mechanism, and provide bounds on the quality of their allocations as measured by the NSW objective. We consider two extreme classes of valuations functions, namely perfect substitutes and perfect complements. For perfect substitutes we show that both mechanisms have a price of anarchy (PoA) of at most 2 and at least $e^{\frac{1}{e}}$ (≈ 1.44). For perfect complements we find that the Fisher market mechanism has arbitrarily bad PoA, but Trading Post performs surprisingly well, achieving a PoA of $(1+\epsilon)$ for every $\epsilon > 0$. In fact, we find that all the equilibria of the Trading Post mechanism are pure, so this bound extends beyond the pure PoA.

*This work was done in part while the authors were visiting the Simons Institute for the Theory of Computing. Simina was also supported by ISF grant 1435/14 administered by the Israeli Academy of Sciences and Israel-USA Bi-national Science Foundation (BSF) grant 2014389, as well as the I-CORE Program of the Planning and Budgeting Committee and The Israel Science Foundation. Vasilis was additionally supported by grants NSF-1408635, NSF-1216073, and NSF-1161813.

1 Introduction

The problem of allocating resources among multiple participants in a fair manner is as old as human society itself, with some of the earliest recorded instances of the problem dating back to more than 2500 years ago[1]. The mathematical study of fair division began with the work of Steinhaus during the second world war, which led to an extensive and growing body of work on fair division protocols within economics and political science, e.g., [4, 7, 50, 61, 66], as well as computer science in recent years, e.g., [2, 8, 16, 46, 49, 58, 59, 62]. Fair division problems have also been surfacing in the realm of allocating computational resources—such as CPU, memory, and bandwidth—among the users of a computing system. Such scenarios are captured through the model of allocating multiple divisible goods, which is an important special case of cake cutting that has drawn considerable interest in computer science [10, 12, 18, 20, 55, 56, 60].

One of the most basic questions underlying the fair division problem is that of *defining fairness* to begin with. A large body of work in economics, especially social choice theory, is concerned with this very question, and numerous solution concepts have been proposed in response. Our fairness concept of choice herein is the *Nash social welfare* (NSW), which dates back to the fifties [43, 52] and has been proposed by Nash as a solution for bargaining problems, using an axiomatic approach. This objective aims to choose an outcome $\mathbf{x}$ maximizing the geometric mean of the utilities ($u_i(\mathbf{x})$) of the n participating agents and, like other standard welfare objectives, it is captured by the following family of functions known as generalized (power) means:

$$M_p(\mathbf{x}) \;=\; \left(\frac{1}{n}\sum_i [u_i(\mathbf{x})]^p\right)^{1/p}.$$

In particular, the NSW corresponds to $M_0(\mathbf{x})$, i.e., the limit of $M_p(\mathbf{x})$ as p goes to zero, i.e., $(\Pi_i u_i(\mathbf{x}))^{\frac{1}{n}}$.

While an extended treatment of the NSW can be found, for example, in Moulin [50], we highlight a fundamental property of the NSW objective, namely that it achieves *a natural compromise between individual fairness and efficiency.* Two other well-studied functions captured by $M_p(\mathbf{x})$ are the (i) *egalitarian* (max-min) objective attained as $p \to -\infty$, and (ii) *utilitarian* (average) objective attained at $p = 1$, which correspond to extreme fairness and extreme efficiency, respectively. However, the former may cause vast inefficiencies, while the latter can completely neglect how unhappy some agents might be. The NSW objective lies between these two extremes and strikes a natural balance between them, since maximizing the geometric mean leads to more balanced valuations, but without neglecting efficiency.

The highly desired fairness and efficiency trade-off that the NSW objective provides can be verified via its close connection with market equilibrium outcomes in the *Fisher market* model—one of the fundamental resource allocation models in mathematical economics. This model was developed by Fisher [6] and studied in an extensive body of literature [13, 17, 23, 28, 29, 32–34, 42, 54]. The basic setting involves a seller who brings multiple divisible goods to the market and a set of buyers equipped with monetary endowments, or budgets. The seller's goal is to set prices that extract all the money from the buyers, each of which aims to spend all its budget in acquiring the best possible bundle at the given prices. A *market equilibrium* is an outcome where supply meets demand, and it has been shown to exist under very general conditions. When every buyer has the same budget, this outcome is also known as the *competitive equilibrium from equal incomes* (CEEI) [40, 65]. For a broad family of valuations, including the ones considered in this paper, the market equilibrium allocation is known to maximize the NSW objective when the budgets are equal. In other words, the seller's goal can be achieved by computing the allocation of the goods that maximizes the NSW which, in turn, implies the desired price for each good.

However, in order to be able to maximize the NSW, the seller needs to know the valuations of the participants. When these valuations are private information of the buyers, a natural solution would be to use the *Fisher market mechanism*: ask the buyers to report their valuations, and then compute the NSW maximizing allocation based on the reports. Unfortunately, it is well known that buyers can feign different interests and eventually get better allocations [1, 3, 9, 14, 15]. This strategic behavior can result in the mechanism computing a market equilibrium with respect to utility functions which may have little to do with reality, leading to unfair allocations. In this work we address the following basic question:

[1] See, e.g., Hesiod's Theogony, where a protocol known as *Cut and Choose* is mentioned.

How well does the Fisher market mechanism optimize its own objective—the Nash social welfare—when the participants are strategic? Are there widely used mechanisms with better guarantees?

This question falls under the general umbrella of implementation theory [21], and particularly, of implementing markets [5, 26, 37, 51, 57]. In this literature, the goal is to identify mechanisms (game forms) for which the set of Nash equilibria coincides with the set of market equilibrium allocations for every possible state of the world [21]. In general this can be achieved only in the limit, as the number of players goes to infinity and every player is infinitesimal compared to the entire economy [27]. In this work we show that, even for small markets, there exist classical mechanisms achieving outcomes that closely approximate the optimal NSW on every instance. We measure the quality of a mechanism using the *price of anarchy* [53], defined as the ratio between the optimal NSW and the NSW of the worst Nash equilibrium outcome obtained by the mechanism.

1.1 Our Results

We study the question of approximately implementing the NSW objective for two extensively studied classes of valuations, namely linear (or additive) and Leontief [17, 32, 53]; two extremes. Linear valuations capture goods that are *perfect substitutes*, i.e., that can replace each other in consumption, such as Pepsi and Coca-Cola. Leontief valuations capture *perfect complements*, i.e., goods that have no value without each other, such as left and right shoes.

Our first set of results concerns the Fisher market mechanism, which collects the bidders valuations in the form of bids and then computes the market equilibrium based on those bids:

The Nash equilibria of the Fisher market mechanism approximate the NSW objective within a factor of 2 for linear valuations. For Leontief valuations, the approximation degrades linearly with the number of players.

These bounds reveal significant differences between the quality of the Nash equilibria of the Fisher market mechanism for complements and substitutes.

Much more strikingly, we find that a classic mechanism known as *Trading Post*, originally introduced by Shapley and Shubik [63] and studied in a long line of work in different scenarios [27, 41, 44, 47], offers yet stronger guarantees. At a high level, rather than collecting the players' preferences and computing the market equilibrium, the Trading Post mechanism gives each participant direct control over how to spend its budget. Once the agents choose how to distribute their budget over the available goods, they receive a fraction from each good that is proportional to the amount they spent on it. Our main results for the Trading Post mechanism are:

The Nash equilibria of the Trading Post mechanism approximate the NSW objective within a factor of 2 for linear valuations. But, for Leontief valuations, the Trading Post mechanism achieves an approximation of $1 + \epsilon$ *for any* $\epsilon > 0$.

In other words, not only does the Trading Post mechanism achieve the same approximation as the Fisher market mechanism for linear valuations, but it also almost perfectly implements the market equilibrium outcome for Leontief valuations! We view this as an important result that testifies to the usefulness and robustness of the Trading Post mechanism. An interpretation of the result is that the Trading Post mechanism limits the extent to which an agent can affect the outcome, thus also limiting the extent to which things can go awry. Specifically, when an agent deviates in the Trading Post mechanism, this deviation has no effect on the way that the other agents are spending their money. On the other hand, an agent's deviation in the Fisher market mechanism can lead to a market equilibrium where the other agents' spending and allocation has changed significantly. In addition to this, in the Fisher market mechanism an agent can affect the price of an item even if the agent does not end up spending on that item in the final outcome. This is in contrast to the Trading Post mechanism where an agent can affect only the prices of the items that this agent is spending on, so the agents are forced to "put their money where their mouth is".

Finally, we prove that the set of mixed Nash equilibria of the Trading Post mechanism coincides with the set of pure Nash equilibria, which extends our PoA bounds for this mechanism to mixed PoA as well.

1.2 Related Work

The paper most closely related to our work is that of Cole et al. [20] which proposes *truthful* mechanisms for approximately maximizing the Nash social welfare objective. One of the truthful mechanisms that they propose, the Partial Allocation mechanism, guarantees a 2.718 approximation of the optimal NSW for both linear and Leontief valuations. In fact, the Partial Allocation mechanism guarantees that *every agent* receives a 2.718 approximation of the value that it would receive in the market equilibrium. But, in order to ensure truthfulness, this mechanism is forced to keep some of the goods unallocated, which makes it inapplicable for many real world settings. Complementing this mechanism, our work analyzes simple and well-studied mechanisms that allocate all the goods fully.

Most of the literature on fair division starting from the 1940's deals with the cake-cutting problem, which models the allocation of a divisible heterogeneous resource such as land, time, and mineral deposits, among agents with different preferences [4, 7, 50, 61, 66]. Recent algorithmic developments include a 4-player bounded and envy-free protocol due to Aziz and Mackenzie [2] and an envy-free and approximately proportional protocol for any number of players due to Segal-Halevi et al. [62]. Truthful mechanisms for cake cutting were studied in several recent works [8, 16, 46, 49, 67]. In particular, Chen et al. [16] study truthful cake-cutting with agents having piecewise uniform valuations and provide a polynomial-time mechanism that is truthful, proportional, and envy-free, while Mossel and Tamuz [49] shows that for general valuations there exists a protocol that is truthful in expectation, envy-free, and proportional for any number of players. The work of Maya and Nisan [46] shows that truthfulness comes at a significant cost in terms of efficiency for direct revelation mechanisms, while Brânzei and Miltersen [8] show that the only strategyproof mechanisms in the standard query model for cake cutting are dictatorships (even for two players; a similar impossiblity holds for $n > 2$). The standard cake cutting model assumes additive valuations, and so it does not capture resources with Leontief valuations, which we also analyze in this paper.

The resource allocation literature has seen a resurgence of work studying fair and efficient allocation for Leontief valuations [25, 36, 39, 56]. These valuations exhibit perfect complements and they are considered to be natural valuation abstractions for computing settings where jobs need resources in fixed ratios. Ghodsi et al. [36] defined the notion of Dominant Resource Fairness (DRF), which is a generalization of the egalitarian social welfare to multiple types of resources. This solution has the advantage that it can be implemented truthfully for this specific class of valuations. Parkes et al. [56] assessed DRF in terms of the resulting efficiency, showing that it performs poorly. Dolev et al. [25] proposed an alternate fairness criterion called Bottleneck Based Fairness, which Gutman and Nisan [39] subsequently showed is satisfied by the proportionally fair allocation. Gutman and Nisan [39] also posed the study of incentives related to this latter notion as an interesting open problem. It is worth noting that Ghodsi et al. [36] acknowledge that the CEEI, i.e., the NSW maximizing allocation would actually be the preferred fair division mechanism in their setting, and that the main drawback of this solution is the fact that it cannot be implemented truthfully. Our results show that the Trading Post mechanism can, in fact, approximate the CEEI outcome arbitrarily well, thus shedding new light on this setting.

The Trading Post mechanism, also known as the Shapley-Shubik game [63], has been studied in an extensive body of literature over the years, sometimes under very different names, such as Chinese auctions [45], proportional sharing mechanism [31], and the Tullock contest in rent seeking [30, 48, 64], the latter being a variant of the game with a different success probability for items that nobody bid on. Trading Post can also be interpreted as a congestion game (see, e.g., [35]), or an all-pay auction when the budgets are intrinsically valuable to the players.

The fact that, facing the Fisher mechanism, the agents may gain by bidding strategically is well known. Adsul et. al. [1] studied the agents' incentives and proved existence and structural properties of Nash equilibria for this mechanism. Extending this work, Chen et al. [14, 15] proved bounds on the extent to which an agent can gain by misreporting for various classes of valuation functions, including additive and Leontief. Finally, Brânzei et al. [9] showed bounds for the price of anarchy of this mechanism with respect to the social welfare objective, and Cole and Tao [19] studied large markets under mild randomness and showed that this price of anarchy converges to one.

2 Preliminaries

Let $N = \{1, \ldots, n\}$ be a set of players (agents) and $M = \{1, \ldots, m\}$ a set of divisible goods. Player i's utility for a bundle of goods is represented by a non-decreasing valuation function $u_i : [0,1]^m \to \mathbb{R}_+$. An allocation $\mathbf{x}$ is a partition of the goods to the players such that $x_{i,j}$ represents the amount of good j received by player i. Our goal will be to allocate all the resources fully; it is without loss of generality to assume that a single unit of each good is available, thus the set of feasible allocations is $\mathcal{F} = \{\mathbf{x} \mid x_{i,j} \geq 0 \text{ and } \sum_{i=1}^{n} x_{i,j} = 1\}$.

Our measure for assessing the quality of an allocation is its Nash social welfare. At a given allocation $\mathbf{x}$ it is defined as follows

$$\mathrm{NSW}(\mathbf{x}) = \left(\prod_{i=1}^{n} u_i(\mathbf{x}_i) \right)^{\frac{1}{n}}.$$

In order to also capture situations where the agents may have different importance or priority, such as clout in bargaining scenarios, we also consider the weighted version of the Nash social welfare objective. Note this is the objective maximized by the Fisher market equilibrium solution when the buyers have different budgets. We slightly abuse notation and refer to the weighted objective as the Nash social welfare (NSW) as well. If $B_i \geq 1$ is the budget of agent i and $\mathcal{B} = \sum_{i=1}^{n} B_i$ is the total budget, the induced market equilibrium in Fisher's model maximizes the following objective:

$$\mathrm{NSW}(\mathbf{x}) = \left(\prod_{i=1}^{n} u_i(\mathbf{x}_i)^{B_i} \right)^{\frac{1}{\mathcal{B}}}.$$

Note that we get back the original definition when all players have the same budget. We would like to find mechanisms that maximize the NSW objective in the presence of strategic agents whose goal is to maximize their own utility. We measure the quality of the mechanisms using the price of anarchy [53] with respect to the NSW objective. Given a problem instance $\mathcal{I}$ and some mechanism $\mathcal{M}$ that yields a set of pure Nash equilibria E, the price of anarchy (PoA) of $\mathcal{M}$ for $\mathcal{I}$ is the maximum ratio between the optimal NSW—obtained at some allocation $\mathbf{x}^*$—and the NSW at an allocation $\mathbf{x} \in E$:

$$\mathrm{PoA}(\mathcal{M}, \mathcal{I}) \;=\; \max_{\mathbf{x} \in E} \left\{ \frac{\mathrm{NSW}(\mathbf{x}^*)}{\mathrm{NSW}(\mathbf{x})} \right\}$$

The price of anarchy of $\mathcal{M}$ is the maximum of this value over all possible instances, i.e., $\max_{\mathcal{I}}\{\mathrm{PoA}(\mathcal{M}, \mathcal{I})\}$.

Valuation Functions. We focus on two very common and extensively studied valuation functions that lie at the two extremes of *perfect substitutes* and *perfect complements*. For both, let $\boldsymbol{v}_i = (v_{i,1}, \ldots, v_{i,m}) \in \mathbb{R}_+^m$ be a vector of valuations for agent i, where $v_{i,j}$ captures the liking of agent i for good j.

Perfect substitutes, defined mathematically through *additive valuations*, represent goods that can replace each other in consumption, such as Pepsi and Coca-Cola. In the additive model, the utility of a player i for bundle $\mathbf{x}_i$ is $u_i(\mathbf{x}_i) = \sum_{j=1}^{m} v_{i,j} \cdot x_{i,j}$.

On the other hand, *perfect complements*, represented by *Leontief utilities*, capture scenarios where one good may have no value without the other, such as a left and a right shoe, or the CPU time and computer memory required for the completion of a computing task. In the Leontief model, the utility of a player i for a bundle $\mathbf{x}_i$ is $u_i(\mathbf{x}_i) = \min_{j=1}^{m} \{x_{i,j}/v_{i,j}\}$; that is, player i desires the items in the ratio $v_{i,1} : v_{i,2} : \ldots : v_{i,m}$. In the Leontief model the coefficients can be rescaled freely, and so we assume w.l.o.g. that $v_{i,j} \geq 1$.

3 The Fisher Market Mechanism

In the Fisher market model, given prices $\boldsymbol{p} = (p_1, \ldots, p_m)$, each buyer i demands a bundle $\boldsymbol{x}_i$ that maximizes her utility subject to her budget constraints; we call this an optimal bundle of buyer i at prices $\boldsymbol{p}$. Prices $\boldsymbol{p}$ induce a market equilibrium if buyers get their optimal bundle and market clears. Formally, prices $\boldsymbol{p}$ and allocation $\boldsymbol{x}$ constitute a market equilibrium if

1. Optimal bundle: $\forall i \in N$ and $\forall \boldsymbol{y} : \boldsymbol{y} \cdot \boldsymbol{p} \leq B_i,\; u_i(\mathbf{x}_i) \geq u_i(\boldsymbol{y})$

2. Market clearing: Each good is fully sold or has price zero, i.e., $\forall j \in M$, $\sum_{j=1}^{m} x_{i,j} \leq 1$, and equality holds if $p_j > 0$. Each buyer exhausts all its budget, i.e., $\forall i \in N$, $\sum_{j=1}^{m} x_{i,j} p_j = B_i$.

For linear and Leontief valuations, the market equilibria can be computed using the Eisenberg-Gale (EG) convex program formulations that follow.

$$
\begin{array}{llcll}
 & \textbf{Linear} & & & \textbf{Leontief} \\
\max & \sum_{i=1}^{n} B_i \cdot \log u_i & & \max & \sum_{i=1}^{n} B_i \cdot \log u_i \\
\text{s.t.} & u_i = \sum_{j=1}^{m} v_{i,j} x_{i,j}, \forall i \in N & & \text{s.t.} & u_i \leq \frac{x_{i,j}}{v_{i,j}},\ \forall i \in N, j \in M \\
 & \sum_{i=1}^{n} x_{i,j} \leq 1, \quad \forall\, j \in M & & & \sum_{i=1}^{n} x_{i,j} \leq 1, \quad \forall\, j \in M \\
 & x_{i,j} \geq 0, \quad \forall\, i \in N, j \in M & & & x_{i,j} \geq 0, \quad \forall\, i \in N, j \in M
\end{array}
\tag{1}
$$

Let p_j be the dual variable of the second inequality (for good j) in both cases, which corresponds to price of good j. Since due to strong duality Karush-Kuhn-Tucker (KKT) conditions capture solutions of the formulations, we get the following characterization for market equilibria.

For linear valuations, an outcome $(\boldsymbol{x}, \boldsymbol{p})$ is a market equilibrium if and only if,

Lin$_1$ $\forall i \in N$ and $\forall j \in M$, $x_{i,j} > 0 \Rightarrow \frac{v_{i,j}}{p_j} = \max_{k \in M} \frac{v_{i,k}}{p_k}$.

Lin$_2$ $\forall i \in N$, $\sum_{j \in M} x_{i,j} p_j = B_i$. $\forall j \in M$, either $\sum_{i \in N} x_{i,j} = 1$ or $p_j = 0$.

For Leontief valuations, an outcome $(\boldsymbol{x}, \boldsymbol{p})$ is a market equilibrium if and only if.

Leo$_1$ $\forall i \in N$ and $\forall j \in M$, $v_{i,j} > 0 \Rightarrow u_i = \frac{x_{i,j}}{v_{i,j}} = \frac{B_i}{\sum_{j \in M} v_{i,j} p_j}$. Note that $\sum_{j \in M} v_{i,j} p_j$ is the amount buyer i has to spend to get unit utility.

Leo$_2$ $\forall i \in N$, $\sum_{j \in M} x_{i,j} p_j = B_i$. $\forall j \in M$, either $\sum_{i \in N} x_{i,j} = 1$ or $p_j = 0$.

The Fisher market mechanism asks that the agents report their valuations and then it computes the market equilibrium allocation with respect to the reported valuations using the EG formulations.

Definition 1 (Fisher Market Mechanism)**.** *The* Fisher Market Mechanism *is such that:*

- *The strategy space of each agent i consists of all possible valuations the agent may pose: $S_i = \{\mathbf{s}_i \mid \mathbf{s}_i \in \mathbb{R}^m_{\geq 0}\}$. We refer to an agent's strategy as a* report*.*
- *Given a strategy profile $\mathbf{s} = (\mathbf{s}_i)_{i=1}^n$, the outcome of the game is a market equilibrium of the Fisher market given by $\langle B_i, \mathbf{s}_i \rangle$, after removing the items j for which $\sum_{i \in N} s_i(j) = 0$. If there exists a market equilibrium $\mathcal{E}$ preferred by all the agents (with respect to their true valuations), then $\mathcal{E}$ is the outcome of the game on $\langle B_i, \mathbf{s}_i \rangle$. Otherwise, the outcome is any fixed market equilibrium.*

It is well known that the buyers have incentives to hide their true valuations in a Fisher market mechanism. We illustrate this phenomenon through an example.

Example 1. *Consider a Fisher market with players $N = \{1, 2\}$, items $M = \{1, 2\}$, additive valuations $v_{1,1} = 1$, $v_{1,2} = 0$, $v_{2,1} = v_{2,2} = 0.5$, and budgets equal to 1. If the players are truthful, the market equilibrium allocation is $\mathbf{x}_1 = (1, 0)$, $\mathbf{x}_2 = (0, 1)$. However, if player 2 pretended that its value for item 2 is a very small $v'_{2,2} = \epsilon > 0$, then player 2 would not only get item 2, but also a fraction of item 1.*

In the rest of this section we study the performance of the Fisher mechanism when the items are substitutes and when they are complements. For both of these cases we provide upper and lower bounds for the price of anarchy.

3.1 Fisher Market: Perfect Substitutes

In this section we study efficiency loss in the Fisher market due to strategic agents with additive valuations. We note that pure Nash equilibria in the induced game are known to always exist due to Adsul et al. [1].[2] Our first main result states that the Fisher market approximates the Nash Social Welfare within a small constant factor, even when the players are strategic.

Theorem 1. *The Fisher Market Mechanism with additive valuations has price of anarchy at most 2.*

We first prove a useful lemma, which bounds the change in prices due to a unilateral deviation in the Fisher market.

Lemma 1. *Let* $\mathbf{p}$ *be the prices in a Fisher market equilibrium. Suppose that buyer* i *unilaterally changes its reported values to* $\mathbf{v}'_i$*, where* $M_i = \{j \in M \mid v'_{i,j} > 0\}$ *are the goods with strictly positive value for* i *under* $\mathbf{v}'_i$*. Then the new market equilibrium prices,* $\mathbf{p}'$*, satisfy the inequality:* $\sum_{j\in M_i} p'_j \leq B_i + \sum_{j \in M_i} p_j$.

Sketch. The most extreme change occurs when buyer i initially reports a value of zero for every item in M_i and subsequently changes its report to zero on everything *except* these items. Since the spending of the other players on these items will not increase, the maximum change is due to the budget of i, i.e., at most B_i. □

In addition to this, we will be using the weighted arithmetic and geometric mean inequality, which upper bounds the weighted geometric mean using the weighted arithmetic mean.

Lemma 2. *For any nonnegative numbers* $\rho_1, \rho_2, \ldots, \rho_n$ *and* $w_1, w_2, \ldots, w_n$ *such that* $W = \sum_{i=1}^n w_i$,

$$\left(\prod_{i=1}^{n} \rho_i^{w_i}\right)^{1/W} \leq \frac{\sum_{i=1}^n \rho_i w_i}{W}.$$

We can now prove the main theorem.

Proof. (of Theorem 1) Given a problem instance with additive valuations, let $\mathbf{x}^*$ be the allocation that maximizes the Nash social welfare—i.e., the market equilibrium allocation with respect to the true valuations—and $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{p}}$ the allocation and prices respectively obtained under some Nash equilibrium of the market, where the players report fake valuations $\tilde{\mathbf{v}}$. Additionally, for each player i, let $\mathbf{x}^i$ be the allocation that would arise if every player $k \neq i$ reported $\tilde{\mathbf{v}}_k$ while i reported its true value v_{ij} for every item j with $x^*_{ij} > 0$, and zero elsewhere.

Since $\tilde{\mathbf{v}}$ is an equilibrium, this unilateral deviation of i cannot increase its utility:

$$u_i(\tilde{\mathbf{x}}_i) \;\geq\; u_i(\mathbf{x}^i_i) \;=\; \sum_{j=1}^{m} x^i_{ij} v_{ij}. \tag{2}$$

Since $\mathbf{x}^i$ is a market equilibrium with respect to the reported values then, according to the KKT condition ($\mathbf{Lin}_1$), $x^i_{ij} > 0$ implies $v_{ij}/p^i_j \geq v_{ik}/p^i_k$ for any other item k, where p^i_j is the price of item j in $\mathbf{x}^i$. Also, the KKT condition ($\mathbf{Lin}_2$) imply that $\sum_j x^i_{ij} p^i_j = B_i$ for every bidder i. Therefore, $\mathbf{x}^i_i$ provides i at least as much value as any other bundle that i can afford facing prices p^i, i.e., any bundle that costs at most B_i. In particular, consider the allocation $\mathbf{x}'$ such that

$$x'_{ij} = \frac{B_i x^*_{ij}}{\sum_{k=1}^m x^*_{ik} p^i_k}.$$

For this allocation we have:

$$\sum_{j=1}^{m} x'_{ij} p^i_j \;=\; \sum_{j=1}^{m} \left(\frac{B_i x^*_{ij}}{\sum_{k=1}^m x^*_{ik} p^i_k}\right) \cdot p^i_j \;=\; \frac{B_i \sum_{j=1}^m x^*_{ij} p^i_j}{\sum_{k=1}^m x^*_{ik} p^i_k} \;=\; B_i.$$

[2] The existence is shown under a *conflict-free* tie breaking rule, which tries to allocate best bundle to as many agents as possible when there is a choice.

Therefore, player i can afford this bundle of items using the budget B_i, which implies that $\mathbf{x}_i^i$ provides i at least as much value, i.e.,

$$u_i(\mathbf{x}_i^i) \;\geq\; u_i(\mathbf{x}'_i) \;=\; \sum_{j=1}^{m} v_{i,j} x'_{i,j} = \frac{B_i u_i(\mathbf{x}_i^*)}{\sum_{j=1}^m x_{i,j}^* p_j^i}.$$

Let ρ_i denote the ratio $\frac{u_i(\mathbf{x}^*)}{u_i(\tilde{\mathbf{x}}_i)}$. Using the Nash equilibrium inequality (2), we get:

$$u_i(\tilde{\mathbf{x}}_i) \geq \frac{B_i u_i(\mathbf{x}_i^*)}{\sum_{j=1}^m x_{i,j}^* p_j^i} \;\Rightarrow\; \frac{u_i(\mathbf{x}^*)}{u_i(\tilde{\mathbf{x}}_i)} \leq \frac{\sum_{j=1}^m x_{i,j}^* p_j^i}{B_i} \;\Rightarrow\; \rho_i B_i \leq \sum_{j=1}^m x_{i,j}^* p_j^i.$$

Summing over all players and using Lemma 1, we get:

$$\sum_{i=1}^n \rho_i B_i \;\leq\; \sum_{i=1}^n \sum_{j=1}^m x_{i,j}^* p_j^i \;\leq\; \sum_{i=1}^n \left(B_i + \sum_{j=1}^m x_{i,j}^* \tilde{p}_j \right) \;\leq\; 2\mathcal{B}. \tag{3}$$

Substituting B_i for w_i in Lemma 2 and using Inequality (3) yields:

$$\left(\prod_{i=1}^n \rho_i^{B_i} \right)^{1/\mathcal{B}} \;\leq\; \frac{\sum_{i=1}^n \rho_i B_i}{\mathcal{B}} \;\leq\; 2.$$

□

Next we show a lower bound for the price of anarchy of the Fisher mechanism. We construct a collection of problem instances whose PoA goes to $e^{1/e}$ as number of players grows.

Theorem 2. *The Fisher mechanism has a price of anarchy no better than $e^{1/e} \approx 1.445$.*

Given some value of n, we construct a market with $n+2$ agents and $n+1$ goods. Fix an integer $k \leq n$; we will set its value later. Each player $i \leq k$ likes only good i, *i.e.*, $v_{i,i} = 1$ and all other $v_{i,j} = 0$. On the other hand, every agent $i \in [k+1, n]$, apart from having $v_{i,i} = 1$, also has some small, but positive, value $v_{i,j} = \epsilon$ for all items $j \leq k$. The rest of that agent's $v_{i,j}$ values are zero. Agent $n+1$ has a small but positive value ϵ' for goods $j \in [k+1, n]$ and value 2 for good $n+1$. Finally, agent $n+2$ values only good $n+1$ at value 2. Here $\epsilon << \epsilon'$, and we will set their values later. In the allocation where every agent $i \in [1, n]$ gets all of good i, while agents $n+1$ and $n+2$ share good $n+1$ equally, the NSW is equal to 1. Next, we construct a Nash equilibrium strategy profile $\boldsymbol{s}$ of the above market where the NSW approaches $(1/e)^{1/e}$ as $n \to \infty$.

We define a strategy profile $\boldsymbol{s}$ where the first k agents and the last agent, i.e., agent $n+2$, bid truthfully, while every bidder $i \in [k+1, n]$ misreports in a way such that it ends up spending some small amount $\delta = 2\epsilon'$ on item i and the rest of its budget, namely $(1-\delta)$, is equally divided on items in $[1, k]$. We later set a value for δ such that agent $n+1$ would want to buy only good $n+1$.

We now show that the above profile is a Nash equilibrium for carefully chosen values of ϵ, and ϵ'. Note that, since bidders $i \in [1, k]$ and bidder $n+2$ bids truthfully and values just one item, each one of these bidders will spend all of its budget on the corresponding items, no matter what the remaining bidders report. Therefore, the price of items $j \in [1, k]$ and item $n+1$ will be exactly 1 if we exclude the spending of bidders $i \in [k+1, n+1]$, no matter what these bidders report. Also, if the price of items $j \in [k+1, n]$ is equal to δ in profile $\boldsymbol{s}$, then this price will not drop below δ, irrespective of what bidder $n+1$ reports. In the next two lemmas, we show that a deviation from $\boldsymbol{s}$ does not help the bidders, even when the prices are held constant for the goods where an agent starts spending more money.

Lemma 3. *If $\delta = 2\epsilon'$ then, even if the prices of goods $j \in [k+1, n]$ are fixed at δ, the $(n+1)^{th}$ bidder has no incentive to spend money on any good other than good $n+1$.*

Proof. If agent $n+1$ spends all of its budget on item $n+1$, then its price becomes 2 and the agent receives half of that item, i.e., a utility of 1. On the other hand, if the agent spends at total of $\gamma > 0$ on goods $j \in [k+1, n]$ then, even if the price of these items remains δ, her utility would be $\epsilon' \frac{\gamma}{\delta} + 2\frac{1-\gamma}{2-\gamma}$. For this to be strictly greater than 1 we have to have $\delta < 2\epsilon' - \gamma\epsilon'$, which contradicts our assumption that $\delta = 2\epsilon'$. □

Furthermore, using a similar analysis as that of Lemma 3, it follows that if agent $i \in [k+1, n]$ deviates in a way that decreases its spending on good i to $2\epsilon' - \tau$ for $\tau > 0$, then the Fisher market mechanism outcome will have the $(n+1)^{th}$ agent spending at least $\frac{\tau}{1+\epsilon'}$ on that item. Thus, such a deviation would cause agent i to lose its monopoly on good i. Next, we show that this is not advantageous for agent i if we set $\epsilon = \frac{1}{n^4}$ and $\epsilon' = \frac{1}{n}$, which implies $\delta = \frac{2}{n}$.

Lemma 4. *For any $\tau \geq 0$, if agent i is spending $\delta - \tau$ on good i and agent $(n+1)$ is spending $\frac{\tau}{1+\epsilon'}$, when others are bidding according to $\boldsymbol{s}$, then agent i's utility is maximized at $\tau = 0$.*

Proof. Note that at $\tau = 0$, agent $i \in [k+1, n]$ is spending $\delta = 2\epsilon'$ on good i and, according to Lemma 3, agent $n+1$ would not be interested in spending on good i. Thus, agent i is buying good i exclusively and spends $\frac{1-\delta}{k}$ on each of the first k goods. Thus the price each one of the first k goods is $1 + (n-k)\frac{1-\delta}{k}$. Agent i's utility at this allocation is

$$1 + \frac{\frac{1-\delta}{k}}{1 + (n-k)\frac{1-\delta}{k}} k\epsilon \;=\; 1 + \frac{(1-\delta)k\epsilon}{k + (n-k)(1-\delta)}$$

As mentioned above, if agent i reduces spending on good i by τ, then the $(n+1)^{th}$ agent will end up spending at least $\frac{\tau}{1+\epsilon'}$ on this item. This may lead to increased prices for the first k goods, but we show that this deviation would not benefit i even if these prices remained the same. The utility of agent i after such a deviation would be

$$\frac{\delta - \tau}{\delta - \tau + \frac{\tau}{1+\epsilon'}} + \frac{(1-\delta+\tau)k\epsilon}{k + (n-k)(1-\delta)}$$

The latter utility is greater than the former only when

$$\frac{\tau k \epsilon}{k + (n-k)(1-\delta)} \;>\; \frac{\frac{\tau}{1+\epsilon'}}{\delta - \tau + \frac{\tau}{1+\epsilon'}} \qquad \Rightarrow \qquad \epsilon \;>\; \frac{k + (n-k)(1-\delta)}{k(\delta + \delta\epsilon' - \tau\epsilon')} \;>\; 1,$$

which contradicts the fact that $\epsilon = \frac{1}{n^4}$. □

Lemmas 3 and 4 imply that $\boldsymbol{s}$ is a Nash equilibrium. The price of the first k goods at this Nash equilibrium is $1 + \frac{(n-k)(1-\delta)}{k} = \frac{k+(n-k)(1-\delta)}{k}$. Thus, the utility of buyer $i \in [1, k]$, who spends all of its \$1 on good i, is $u_i = \frac{k}{k+(n-k)(1-\delta)}$. On the other hand, the utility of buyer $i \in [k+1, n]$, who gets all of good i and some of the first k goods, is $u_i = 1 + \frac{(1-\delta)k\epsilon}{k+(n-k)(1-\delta)}$. Agents $(n+1)$ and $(n+2)$ get half of good $n+1$ and thereby get utility of 1 each. Since $\epsilon = \frac{1}{n^4}$ and $\epsilon' = \frac{1}{n}$, then $\forall i \in [1, k]$, $\lim_{n\to\infty} u_i = \frac{k}{n}$ and $\forall i \in [k+1, n] \lim_{n\to\infty} u_i = 1$. Thus, NSW at this bid profile as $n \to \infty$ is $(\frac{k}{n})^{\frac{k}{n}}$, and thereby PoA is at least $(\frac{n}{k})^{k/n}$. Letting $k = \frac{n}{e}$, this becomes $e^{1/e}$, which concludes the proof of Theorem 2.

3.2 Fisher Market: Perfect Complements

On the other hand, the Fisher market with perfect complements has a price of anarchy that grows linearly with the number of players. The existence of pure Nash equilibria in Fisher markets with Leontief valuations was established by Brânzei *et al.* [9].

Theorem 3. *The Fisher Market Mechanism with Leontief valuations has a price of anarchy of n and the bound is tight.*

Proof. Our tool is the following theorem, which states that Fisher markets with Leontief utilities have Nash equilibria where players copy each others strategies.

Lemma 5 ([9]). *The Fisher Market mechanism with Leontief preferences always has a Nash equilibrium where every buyer reports the uniform valuation $(1/m, \ldots, 1/m)$.*

For completeness, we include the worst case example. Consider an instance with n players of equal budgets ($B_i = 1$) and n items, where each player i likes item i and nothing else; that is, $v_{i,i} = 1$, for all $i \in N$ and $v_{i,j} = 0, \forall i \in N, \forall j \neq i$. Then the optimal Nash Social Welfare is obtained in the Fisher market equilibrium, where the price of each item j is $p_j = 1$ and the allocation is $x_{i,i} = 1, \forall i \in N$ and $x_{i,j} = 0, \forall i \in N, \forall j \neq i$.

However, the strategy profile $\mathbf{y} = (\mathbf{y_1}, \ldots, \mathbf{y}_n)$, where $\mathbf{y}_i = (1/n, \ldots, 1/n)$ is a Nash equilibrium in which each player i gets a fraction of $1/n$ from every item, yielding utility $1/n$ for every player. It follows that the price of anarchy is n.

For a general upper bound, we note that any Nash equilibrium must be proportional, i.e. each player i gets a fraction of at least $B_i/\mathcal{B}$ of its best possible utility, OPT_i. Let $(\mathbf{p}, \mathbf{x})$ be a Nash equilibrium of the market, achieved under some reports $\mathbf{v}'$. Suppose for a contradiction that there exists player i with $u_i(\mathbf{x}_i) < B_i/\mathcal{B} \cdot \mathrm{OPT}_i$. Then if i reported instead its true valuation $\mathbf{v}_i$, the new market equilibrium, $(\mathbf{p}', \mathbf{x}')$, achieved under valuations $(\mathbf{v}_i, \mathbf{v}'_{-i})$, should satisfy the inequality $u_i(\mathbf{x}'_i) \geq B_i/\mathcal{B} \cdot \mathrm{OPT}_i$. If this were not the case, the outcome would not be a market $(\mathbf{v}_i, \mathbf{v}'_{-i})$, since i can afford the bundle $\mathbf{y} = (B_1/\mathcal{B}, \ldots, B_m/\mathcal{B})$:

$$\mathbf{p}(\mathbf{y}) = \sum_{j=1}^{m} p_j \cdot \frac{B_i}{\mathcal{B}} = \frac{B_i}{\mathcal{B}} \cdot \sum_{j=1}^{m} p_j = B_i, \tag{4}$$

Moreover, $u_i(\mathbf{y}) = B_i/\mathcal{B} \cdot \mathrm{OPT}_i$, which together with Identity 4 contradicts the market equilibrium property of $(\mathbf{p}', \mathbf{x}')$. Thus in any Nash equilibrium $(\mathbf{p}, \mathbf{x})$ we have that $u_i(\mathbf{x}_i) \geq B_i/\mathcal{B} \cdot \mathrm{OPT}_i$, and so the Nash Social Welfare is $\mathrm{NSW}(\mathbf{x}) \geq \prod_{i=1}^{n} \left(\frac{B_i}{\mathcal{B}} \cdot \mathrm{OPT}_i\right)^{\frac{B_i}{\mathcal{B}}}$. Then the price of anarchy can be bounded as follows:

$$\mathrm{PoA} \leq \prod_{i=1}^{n} \frac{\mathrm{OPT}_i^{B_i/\mathcal{B}}}{\left(\frac{B_i}{\mathcal{B}} \cdot \mathrm{OPT}_i\right)^{B_i/\mathcal{B}}} \leq \prod_{i=1}^{n} \left(\frac{\mathcal{B}}{B_i}^{B_i/\mathcal{B}}\right) \leq \frac{\sum_{i=1}^{n} B_i \cdot \left(\frac{\mathcal{B}}{B_i}\right)}{\mathcal{B}} = n,$$

where for the last inequality we used the fact that the weighted geometric mean is bounded by the weighted arithmetic mean (Lemma 2).This completes the proof of the theorem. □

4 The Trading Post Mechanism

The important difference between the Trading Post mechanism and the Fisher Market mechanism is the strategy space of the agents. More precisely, unlike the Fisher market mechanism, where the agents' are asked to report their valuations, the Trading Post mechanism instead asks the agents to directly choose how to distribute their budgets. Once the agents have chosen how much of their budget to spend on each of the goods, the total spending on each good j is treated as its price, and each agent i is allocated a fraction of good j proportional to the amount that i is spending on j. Therefore, the strategy set of each player i is $\left\{\mathbf{b}_i \in [0,1]^m \mid \sum_{j=1}^{m} b_{i,j} = B_i\right\}$. A bid profile $\mathbf{b} = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ yields the allocation:

$$x_{i,j} = \frac{b_{i,j}}{\sum_{k=1}^{n} b_{k,j}}$$

For the Trading Post mechanism instances we assume that every good is desired by at least two players, i.e., for each good j, there exist buyers i, k such that $v_{i,j} \neq 0, v_{k,j} \neq 0$. If this were not the case, then we could either discard that good or give it away for free. This is a commonly assumed property called *perfect competition*.

Our main results in this section are that the Nash equilibria of Trading Post approximate the NSW objective within a factor of $(2e)^{1/e} \approx 1.864$ for linear valuations (the bound becomes 2 for arbitrary budgets), and within a factor of $1 + \epsilon$ for every $\epsilon > 0$ for Leontief valuations. Moreover, in both cases, all the Nash equilibria of the game are pure.

4.1 Trading Post: Perfect Substitutes

The existence of pure Nash equilibria in the Trading Post mechanism for additive valuations was established by Feldman, Lai, and Zhang [31] under perfect competition. Without competition, even very simple games may not have pure Nash equilibria. To see this, consider for instance a game with two players, two items, and additive valuations $v_{1,1} = 1$, $v_{1,2} = 0$, $v_{2,1} = v_{2,2} = 0.5$. Through a case analysis it can be seen that both players will compete for item 1, while player 2 is the only one that wants item 2, reason for which this player will successively reduce its bid for 2 to get a higher fraction from item 1. However, in the limit of its bid for the second item going to zero, player 2 loses the item.

Our first result for Trading Post quantifies the degradation in the Nash Social welfare value for additive valuations.

Theorem 4. *The Trading Post Mechanism with additive valuations has price of anarchy at most 2.*

Proof. Given a problem instance with additive valuations, let $\mathbf{x}^*$ be the allocation that maximizes the Nash social welfare and p^* be the corresponding market equilibrium prices. Also, let $\tilde{\mathbf{x}}$ be the allocation obtained under a Nash equilibrium where each player i bids $\tilde{\boldsymbol{b}}_i$ and the price of each item j is $\tilde{p}_j = \sum_i \tilde{b}_{i,j}$.

For some player i, let $\mathbf{x}'$ be the allocation that arises if every player $k \neq i$ bids $\tilde{\boldsymbol{b}}_k$ while agent i unilaterally deviates to $b'_{i,j} = \tilde{b}_{i,j} + \delta_{i,j}$ for each item j and some $\delta_{i,j} \in \mathbb{R}$ such that $\sum_j \delta_{i,j} = 0$. Let $\delta_{i,j}$ be such that for some $\beta_i > 0$ and every item j

$$\frac{\tilde{b}_{i,j} + \delta_{i,j}}{\tilde{p}_j + \delta_{i,j}} \;=\; \frac{x^*_{i,j}}{\beta_i}. \tag{5}$$

This bid $\mathbf{b}'_i$ is implied by the solution of the following program with variables $\delta_{i,j}$ for each $j \in M$:

$$\begin{aligned} \text{minimize:} \quad & \beta_i \\ \text{subject to:} \quad & \beta_i = \frac{x^*_{i,j}(\tilde{p}_j + \delta_{i,j})}{\tilde{b}_{i,j} + \delta_{i,j}}, \quad \forall j \in M \\ & \sum_{j \in M} \delta_{i,j} = 0, \end{aligned}$$

The allocation induced by this unilateral deviation of i is

$$x'_{i,j} \;=\; \frac{b'_{i,j}}{\tilde{p}_j + \delta_{i,j}} \;=\; \frac{\tilde{b}_{i,j} + \delta_{i,j}}{\tilde{p}_j + \delta_{i,j}} \;=\; \frac{x^*_{i,j}}{\beta_i}.$$

Therefore, the utility of player i after this deviation is $u_i(\mathbf{x}^*)/\beta_i$. But the outcome $\tilde{\mathbf{x}}$ is a Nash equilibrium, so this deviation cannot yield a higher utility for i, which implies that $u_i(\tilde{\mathbf{x}}) \geq u_i(\mathbf{x}^*)/\beta_i$. By definition of b', we get $\sum_j b'_{i,j} = \sum_j x'_{i,j}(\tilde{p}_j + \delta_{i,j}) = B_i$. Therefore, replacing for $x'_{i,j} = x^*_{i,j}/\beta_i$, we get

$$B_i \beta_i = \sum_j x^*_{i,j}(\tilde{p}_j + \delta_{i,j}) \tag{6}$$

Since $x*_{i,j} \geq 0$, Equation (5) implies that $\delta_{i,j} \geq -\tilde{b}_{i,j}$ for every item j. Therefore, the sum of the negative $\delta_{i,j}$ values is no less than $-B_i$. This implies that $\delta_{i,j} \leq B_i$, since $\sum_j \delta_{i,j} = 0$, so for every agent i we have

$$\sum_{i=1}^{n} x^*_{i,j}(\tilde{p}_j + \delta_{i,j}) \;\leq\; B_i + \sum_{i=1}^{n} x^*_{i,j}\tilde{p}_j. \tag{7}$$

Using Inequalities (6) and (7), and summing over all players gives

$$\sum_{i=1}^{n} B_i \frac{u_i(\mathbf{x}^*_i)}{u_i(\tilde{\mathbf{x}}_i)} \;\leq\; \sum_{i=1}^{n} B_i \beta_i \;\leq\; \sum_{i=1}^{n} \left(B_i + \sum_{i=1}^{n} x^*_{i,j} p_j \right) \;\leq\; 2\mathcal{B}. \tag{8}$$

Note that Inequality (8) is the same as inequality (3) in the proof of Theorem 1. Thus, following the same arguments as those in that proof proves the theorem. □

The next theorem complements the upper bound on the price of anarchy with a lower bound of approximately 1.445.

Theorem 5. *The Trading Post mechanism has a price of anarchy no better than* $e^{1/e} \approx 1.445$.

Proof. In order to prove this theorem we can use the same family of problem instances used in the proof of Theorem 2. In fact, verifying that the market equilibrium outcome induced by the strategy profile $\boldsymbol{s}$ in that construction is a Nash equilibrium for the Trading Post mechanism as well is much more straightforward since an agent's deviation may affect only the way that particular agent ends up spending its budget. □

4.2 Trading Post: Perfect Complements

We begin by characterizing the precise conditions under which the Trading Post mechanism has exact pure Nash equilibria for Leontief utilities; the proof is included in the appendix.

Theorem 6. *The Trading Post mechanism with Leontief utilities has pure Nash equilibria if and only if the corresponding market equilibrium prices are all strictly positive. When this happens, the Nash equilibrium utilities in Trading Post are unique and the price of anarchy is* 1.

This theorem establishes a correspondence between the pure Nash equilibria of Trading Post and the corresponding market equilibria with respect to the agents' (true) valuations. We observe however that existence of pure Nash equilibria in Trading Post is not guaranteed for Leontief utilities.

Example 2. *Consider a game with two players and two items, where player 1 has values $v_{1,1} = v_{1,2} = 0.5$ and player 2 has $v_{2,1} = 0.9$, $v_{2,2} = 0.1$. Assume there is a pure Nash equilibrium profile $\mathbf{b}$. Since both players require a non-zero amount from every item for their utility to be positive, we have that $b_{i,j} > 0$ for all $i, j \in \{1, 2\}$. Denote $b_1 = b_{1,1}$ and $b_2 = b_{2,1}$; then $b_{1,2} = 1 - b_1$ and $b_{2,1} = 1 - b_2$. Note that each player must receive the two items in the same ratio relative to its valuation; that is:*

$$u_i(\mathbf{b}) = \left(\frac{b_i}{b_1 + b_2}\right)\frac{1}{v_{i,1}} = \left(\frac{1 - b_i}{b_1 + b_2}\right)\frac{1}{v_{i,2}} \tag{9}$$

Otherwise, if the two ratios were not equal, then a player could transfer weight among the items to improve the smaller fraction. Then the requirement in 9 are equivalent to the following equations:

$$\left(\frac{b_1}{b_1 + b_2}\right)\frac{1}{0.5} = \left(\frac{1 - b_1}{b_1 + b_2}\right)\frac{1}{0.5} \iff b_1 = b_2 \tag{10}$$

and

$$\left(\frac{b_2}{b_1 + b_2}\right)\frac{1}{0.9} = \left(\frac{1 - b_2}{b_1 + b_2}\right)\frac{1}{0.1} \iff 8b_2^2 + 8b_1b_2 = 9b_1 + 7b_2 \tag{11}$$

Combining equations 10 and 11, we get that $b_1 = 1$ and $b_2 = 1$, which contradicts the requirement that $b_1, b_2 \in (0, 1)$. Thus the equilibrium profile $\mathbf{b}$ cannot exist.

The issue illustrated by this example is that the Trading Post cannot implement market outcomes when there exist items priced at zero in the corresponding market equilibrium. This motivates us to introduce an entrance fee in the Trading Post mechanism, denoted by a parameter $\Delta > 0$, which is the minimum amount that an agent needs to spend on an item in order to receive any of it. We denote the corresponding mechanism by $\mathcal{TP}(\Delta)$. The value of Δ can be arbitrarily small, so its impact on the outcome of the game is insignificant. Formally, given a bid profile $\mathbf{b} = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$, let $\overline{\mathbf{b}} = (\overline{\mathbf{b}}_1, \ldots, \overline{\mathbf{b}}_n)$ be the "effective" bid profile which, for every $i \in N$ and $j \in M$, satisfies $\overline{b}_{i,j} = b_{i,j}$ if $b_{i,j} \geq \Delta$, and $\overline{b}_{i,j} = 0$ otherwise. Then, the bid profile $\mathbf{b}$ yields the allocation:

$$x_{i,j} = \frac{\overline{b}_{i,j}}{\sum_{k=1}^{n} \overline{b}_{k,j}}$$

Clearly, in any Nash equilibrium, we have that for every player i and item j we have that $b_{i,j} \geq \Delta$ or $b_{i,j} = 0$(the latter identity holds for those goods j that are outside of player i's demand). The main result of this section is that, for every $\epsilon \in (0, 1/m)$, the Trading Post mechanism with $\Delta \leq \epsilon/m^2$ has a price of anarchy of at most $1 + \epsilon$.

We first show that Trading Post has a pure Nash equilibrium for Leontief utilities for every strictly positive entrance fee. The proof uses an application of Glicksberg's theorem for continuous games and is included in the appendix.

Theorem 7. *The parameterized Trading Post mechanism $\mathcal{TP}(\Delta)$ is guaranteed to have a pure Nash equilibrium for every $\Delta > 0$.*

We start by defining a notion of approximate market equilibrium that will be useful

Definition 2 (ϵ-market equilibrium)**.** *Given a problem instance and some $\epsilon > 0$, an outcome $(\mathbf{p}, \mathbf{x})$ is an ϵ-market equilibrium if and only if:*

- *All the goods with a positive price are completely sold.*
- *All the buyers exhaust their budget.*
- *Each buyer gets an ϵ-optimal bundle at prices $\mathbf{p}$; that is, for every bundle $\mathbf{y} \in [0,1]^m$ that i could afford at these prices ($\mathbf{p} \cdot \mathbf{y} \leq B_i$), we have $u_i(\mathbf{y}) \leq u_i(\mathbf{x}_i)(1+\epsilon)$.*

The following theorem states that for every small enough $\epsilon > 0$, all the PNE of the Trading Post game with a small enough entrance fee correspond to ϵ-market equilibrium outcomes.

Theorem 8. *Let $\epsilon > 0$. Then for every $0 < \Delta < \min\left\{\frac{\epsilon}{m^2}, \frac{1}{m}\right\}$, every pure Nash equilibrium of the mechanism $\mathcal{TP}(\Delta)$ with Leontief valuations corresponds to an ϵ-market equilibrium.*

Proof. Let $\tilde{\boldsymbol{b}}$ be a pure Nash equilibrium of $\mathcal{TP}(\Delta)$ and $\mathbf{x}$ the induced allocation. For each player i, let $D_i = \{j \in M \mid v_{i,j} > 0\}$ be the set of items that i requires, and let $m_i = |D_i|$. We also override notation and refer to $u_i(\mathbf{b})$ as the utility of player i when the strategy profiles are $\mathbf{b}$.

First note that $\tilde{\boldsymbol{b}}_{i,j} > 0$ for each player i and item $j \in D_i$. If this were not the case, then player i would get zero utility at strategy profile $\tilde{\boldsymbol{b}}$; this is worse than playing the uniform strategy $\mathbf{z}_i = (B_i/m, \ldots, B_i/m)$, which guarantees i a positive value regardless of the strategies of the other players $\tilde{\boldsymbol{b}}_{-i}$, namely:

$$u_i(\mathbf{z}_i, \tilde{\boldsymbol{b}}_{-i}) \;=\; \min_{j \in D_i} \left\{ \frac{z_{i,j}}{z_{i,j} + \sum_{k \neq i} \tilde{b}_{k,j}} \cdot \frac{1}{v_{i,j}} \right\} \;\geq\; \min_{j \in D_i} \left\{ \frac{B_i/m}{B_i/m + \sum_{k \neq i} B_k} \cdot \frac{1}{v_{i,j}} \right\} \;>\; 0$$

For each player i and item $j \in D_i$, denote the fraction of utility that i derives from j by:

$$\phi_{i,j} = \frac{\tilde{b}_{i,j}}{\sum_{k=1}^{n} \tilde{b}_{k,j}} \cdot \frac{1}{v_{i,j}}$$

Then $u_i(\tilde{\boldsymbol{b}}) = \min_{j \in D_i} \phi_{i,j}$. Sort the items in D_i increasingly by their contribution to i's utility: $\phi_{i,i_1} \leq \phi_{i,i_2} \leq \ldots \leq \phi_{i,i_{m_i}}$; it follows that $u_i(\tilde{\boldsymbol{b}}) = \phi_{i,i_1}$. Let $S_i = \{j \in D_i \mid \phi_{i,j} = \phi_{i,i_1}\}$ be the items received in the smallest fraction (equal to i's utility). If $S_i = M$, then the analysis is similar to the exact equilibrium case, where the prices are strictly positive. The difficult case is when $S_i \neq M$. Then player i is getting a higher than necessary fraction from some resource $j \in M \setminus S_i$. Thus i would improve by shifting some of the mass from item j to the items in S. Since $\tilde{\boldsymbol{b}}$ is an equilibrium, no such deviation is possible. Then it must be the case that $\tilde{b}_{i,j} = \Delta$ for all $j \in D_i \setminus S_i$.

Now interpret the bids and allocation as a market equilibrium with Leontief utilities $\mathbf{v}$ and budgets B_i, by setting the prices to $\mathbf{p} = (p_1, \ldots, p_m)$, where $p_j = \sum_{i=1}^{n} \tilde{b}_{i,j}$ for all $j \in M$, and the allocation to $\mathbf{x}$, the same as the one induced by the bids $\tilde{\boldsymbol{b}}$ in the trading post game. We argue that $(\mathbf{p}, \mathbf{x})$ is an ϵ-market equilibrium. Clearly at the outcome $(\mathbf{p}, \mathbf{x})$ all the goods are sold and each buyer exhausts their budget. Moreover observe that all the prices are strictly positive. We must additionally show that each player gets an ϵ-optimal bundle at $(\mathbf{p}, \mathbf{x})$.

Fix an arbitrary buyer i. Let $\mathbf{y}_i$ be an optimal bundle for i given prices $\mathbf{p}$, and let $q_{i,j}$ be the amount of money spent by i to purchase $y_{i,j}$ units of good j at these prices. An upper bound on the optimal value $u_i(\mathbf{y}_i)$ is attained when buyer i shifts *all* the money spent on purchasing items outside S_i to purchase instead higher fractions from the items in S_i. Since the strategy profile $\mathbf{b}$ is an exact equilibrium in the game $\mathcal{TP}(\Delta)$, the amount of money spent by player i on items outside S_i is at most $(m-1)\Delta$; thus i spends at most $B_i - (m-1)\Delta$ on the remaining items in S_i.

By an averaging argument, there exists a good $j \in S_i$ on which i spends the greatest amount of its money, i.e.

$$\tilde{b}_{i,j} \geq \frac{B_i - (m-1)\Delta}{|S_i|}.$$

This will be the item for which the gain brought by the deviation in spending is modest. Formally, the maximum fraction of utility that i can get from item j—without decreasing the ratios at which the other items in S_i are received—is:

$$\begin{aligned}
\phi'_{i,j} &= \frac{q_{i,j}}{p_j \cdot v_{i,j}} \leq \frac{\tilde{b}_{i,j} + (m-1)\Delta}{\left(\sum_{k=1}^{n} \tilde{b}_{k,j}\right) \cdot v_{i,j}} \\
&= \phi_{i,j} + \frac{(m-1)\Delta}{\left(\sum_{k=1}^{n} \tilde{b}_{k,j}\right) \cdot v_{i,j}} \leq \phi_{i,j} + \frac{\tilde{b}_{i,j} \cdot \epsilon}{\left(\sum_{k=1}^{n} \tilde{b}_{k,j}\right) \cdot v_{i,j}} \\
&= \phi_{i,j} \cdot (1+\epsilon) = u_i(\mathbf{x}_i)(1+\epsilon)
\end{aligned}$$

where in the inequalities we additionally used that $\Delta < \epsilon^2/m$, $B_i \geq 1\ \forall i \in N$, and $S_i \leq m-1$. The identities hold since item j is in the tight set S_i. Then $u_i(\mathbf{y}_i) \leq \phi'_{i,j} \leq u_i(\mathbf{x}_i)(1+\epsilon)$. Thus each player gets an ϵ-optimal bundle, and so $(\mathbf{p}, \mathbf{x})$ is an ϵ-market equilibrium. □

The following theorem, which we believe is of independent interest, states that in Fisher markets with Leontief utilities, approximate market equilibria are close to exact equilibria in terms of their Nash Social Welfare.

Theorem 9. *The Nash social welfare at an ϵ-market equilibrium for Leontief utilities is at least a $\frac{1}{(1+\epsilon)}$ factor of the optimal Nash social welfare.*

Proof. For any given problem instance, let $(\boldsymbol{p}', \boldsymbol{x}')$ be an ϵ-market equilibrium and let $(\boldsymbol{p}^*, \boldsymbol{x}^*)$ be exact market equilibrium prices and allocation. By abuse of notation let $u_i(\boldsymbol{p}')$ denote the optimal utility player i can obtain at prices $\boldsymbol{p}'$, *i.e.*, $u_i(\boldsymbol{p}') = \max\{u_i(\boldsymbol{y}) \mid \boldsymbol{y} \geq 0;\ \ \boldsymbol{p}' \cdot \boldsymbol{y} \leq B_i\}$.

For Leontief utility functions, convex formulation of (1) captures the market equilibrium allocation. Note that, in order to get a utility of 1 at prices $\boldsymbol{p}$, agent i would need to spend a total amount of money equal to $\phi_i(\boldsymbol{p}) = \sum_j v_{ij} p_j$. Devanur [24] derived the following dual of this convex program:

$$\begin{array}{ll}
\min: & \sum_j p_j - \sum_i B_i \log(\phi(\boldsymbol{p})) + \sum_i B_i \log(B_i) - \sum_i B_i \\
s.t. & \forall j : p_j \geq 0
\end{array}$$

Note that the term $(\sum_i B_i \log(B_i) - \sum_i B_i)$ is a constant for a given market since B_is are constants, and hence is omitted in [24]. Since $(\boldsymbol{p}^*, \boldsymbol{x}^*)$ is a market equilibrium, using strong duality and the fact that agents spend all their money at equilibrium, i.e., $\sum_j p^*_j = \sum_i B_i$:

$$\sum_i B_i \log(u_i(\boldsymbol{x}^*_i)) \ = \ -\sum_i B_i \log(\phi(\boldsymbol{p}^*)) + \sum_i B_i \log(B_i) \tag{12}$$

Furthermore, at the ϵ-market equilibrium $(\boldsymbol{x}', \boldsymbol{p}')$ all the agents spend all their money, implying $\sum_j p'_j = \sum_i B_i$. Since $\boldsymbol{p}'$ is a feasible dual solution,

$$-\sum_i B_i \log(\phi(\boldsymbol{p}^*)) + \sum_i B_i \log(B_i) \ \leq \ -\sum_i B_i \log(\phi(\boldsymbol{p}')) + \sum_i B_i \log(B_i).$$

Substituting the left hand side using Equation (12), and taking an antilogarithm on both sides yields

$$\prod_i u_i(\boldsymbol{x}^*_i)^{B_i} \ \leq \ \prod_i \left(\frac{B_i}{\phi(\boldsymbol{p}')}\right)^{B_i}. \tag{13}$$

Since the optimal utility that agent i gets at prices $\boldsymbol{p}'$ is $u_i(\boldsymbol{p}')$, which she derives using B_i money, and while for unit utility she needs $\phi(\boldsymbol{p}')$ money, we get

$$\forall i: \quad u_i(\boldsymbol{p}') = \frac{B_i}{\phi(\boldsymbol{p}')} \tag{14}$$

Since $(\boldsymbol{x}', \boldsymbol{p}')$ is an ϵ-market equilibrium, each agent gets an ϵ-optimal bundle, so $u_i(\boldsymbol{p}') \leq u_i(\boldsymbol{x}'_i)(1+\epsilon)$. According to (14), this implies that $\frac{B_i}{\phi(\boldsymbol{p}')} \leq u_i(\boldsymbol{x}'_i)(1+\epsilon)$, which we combine with (13) to get:

$$\prod_i u_i(\boldsymbol{x}^*_i)^{B_i} \leq \prod_i \left(\frac{B_i}{\phi(\boldsymbol{p}')}\right)^{m_i} \leq (1+\epsilon)^{\mathcal{B}} \prod_i u_i(\boldsymbol{x}'_i)^{B_i}$$

Since the Nash social welfare at $\boldsymbol{x}$ is $(\Pi_i u_i(\boldsymbol{x}_i)^{B_i})^{\frac{1}{\mathcal{B}}}$, the result follows. $\square$

Finally, we can state the main result of this section.

Theorem 10. *For every $\epsilon > 0$, the Trading Post game $\mathcal{TP}(\Delta)$ with entrance fee $0 < \Delta < \min\left\{\frac{\epsilon^2}{m}, \frac{1}{m}\right\}$ has a price of anarchy of $1+\epsilon$, even for arbitrary budgets.*

Proof. By Theorem 8, every PNE of $\mathcal{TP}(\Delta)$ corresponds to an ϵ-market equilibrium. By Theorem 9, every ϵ-market equilibrium attains at least a fraction $\frac{1}{1+\epsilon}$ of the optimal Nash social welfare. Thus, the price of anarchy of $\mathcal{TP}(\Delta)$ is $1+\epsilon$, which completes the proof. $\square$

4.3 Trading Post: Beyond Pure Nash Equilibria

So far we have bound Price of Anarchy with respect to pure Nash equilibria in case of both the mechanisms. A natural question is what if NSW at mixed Nash equilibria are bad. In case of Trading Post mechanism we rule out this case by showing that all its Nash equilibria are pure, be it for linear or Leontief valuations. Details of this result may be found in Appendix B, where we show the following theorem.

Theorem 11.
- *For market with linear utilities, every Nash equilibrium of the corresponding trading post game is pure (Theorem 12).*
- *For market with Leontief utilities, every Nash equilibrium of the corresponding Δ trading post game $TP(\Delta)$, for $\Delta > 0$, is pure (Theorem 13).*

Finally, given such an efficiency achieved by Trading Post mechanism one may wonder if these mechanisms in-fact have unique equilibrium. We rule this out through Examples 3 and 4 for linear and Leontief valuations respectively given in Appendix B.

5 Discussion

In this paper we analyzed two well-known mechanisms, namely the Fisher market mechanism and the Trading Post mechanism, in terms of their price of anarchy with respect to the Nash social welfare objective. We showed that both mechanisms manage to obtain a small price of anarchy in the linear valuations case but, when it comes to Leontief valuations, the two mechanisms perform very differently: the Fisher market mechanism has a very high price of anarchy, while the Trading Post mechanism is essentially optimal. As an interpretation of this result we observe that, although both of these mechanisms are closely connected to the market equilibrium in the Fisher model, the Fisher market mechanism allows a unilateral deviation of one agent to affect the way the other agents end up spending their budgets, while the Trading Post mechanism does not provide the agent with the same power. It therefore appears that this limitation is helpful in avoiding undesired Nash equilibria. Furthermore, we show that the bounds for the Trading Post mechanism go beyond pure Nash equilibria, as the set of mixed Nash equilibria is equal to the set of pure Nash equilibria for this mechanism. A natural extension is to see if these nice properties extend to general CES valuation functions. Note that CES valuations are parameterized using some $\rho \in (-\infty, 1]$, and linear and Leontief correspond to $\rho = 1$ and $\rho = -\infty$ respectively. Another important question is to achieve good approximation of the NSW using truthful non-wasteful mechanisms. There has been some work along this line, but with other objectives as their focus.

References

[1] Bharat Adsul, Ch. Sobhan Babu, Jugal Garg, Ruta Mehta, and Milind Sohoni. Nash equilibria in Fisher market. In *SAGT*, pages 30–41, 2010.

[2] H. Aziz and S. Mackenzie. A discrete and bounded envy-free cake cutting protocol for four agents. In *Proceedings of STOC*, 2016.

[3] Moshe Babaioff, Brendan Lucier, Noam Nisan, and Renato Paes Leme. On the efficiency of the walrasian mechanism. In *EC*, pages 783–800, 2014.

[4] J.B. Barbanel. *The Geometry of Efficient Fair Division.* Cambridge Univ. Press, 2004.

[5] C. Bevia, L. Corchón, and S. Wilkie. Implementation of the Walrasian correspondence by market games. *Review of Economic Design*, 7:429–442, 2003.

[6] W. C. Brainard and H. E. Scarf. How to compute equilibrium prices in 1891. *Cowles Foundation Discussion Paper*, 1270, 2000.

[7] S. Brams and A. Taylor. *Fair Division: from cake cutting to dispute resolution.* Cambridge University Press, Cambridge, 1996.

[8] S. Brânzei and P. B. Miltersen. A dictatorship theorem for cake cutting. In *Proceedings of IJCAI*, pages 482–488, 2015.

[9] S. Brânzei, Y. Chen, X. Deng, A. Filos-Ratsikas, S. Frederiksen, and J. Zhang. The Fisher market game: Equilibrium and welfare. In *AAAI*, pages 587–593, 2014.

[10] S. Brânzei, I. Caragiannis, D. Kurokawa, and A. D. Procaccia. An Algorithmic Framework for Strategic Fair Division. In *Proceedings of AAAI*, 2016.

[11] A. Cambini and L. Martein. *Generalized Convexity and Optimization: Theory and Applications.* Springer, 2009.

[12] Deeparnab Chakrabarty, Julia Chuzhoy, and Sanjeev Khanna. On allocating goods to maximize fairness. In *FOCS*, pages 107–116, 2009.

[13] N. Chen, Xiatie Deng, Xiaoming Sun, and Andrew Yao. Fisher Equilibrium Price with a class of Concave Utility Functions. In *ESA*, pages 169–179, 2004.

[14] Ning Chen, Xiaotie Deng, and Jie Zhang. How profitable are strategic behaviors in a market? In *ESA*, 2011.

[15] Ning Chen, Xiaotie Deng, Hongyang Zhang, and Jie Zhang. Incentive ratios of Fisher markets. In *ICALP*, 2012.

[16] Yiling Chen, John K. Lai, David C. Parkes, and Ariel D. Procaccia. Truth, justice, and cake cutting. volume 77, pages 284–297, 2013.

[17] B. Codenotti and K. Vardarajan. Equilbrium for Markets with with Leontief Utilities. In *ICALP*, 2004.

[18] Richard Cole and Vasilis Gkatzelis. Approximating the Nash Social Welfare with Indivisible Items. In *STOC*, pages 371–380, 2015.

[19] Richard Cole and Yixin Tao. The price of anarchy of large Walrasian auctions. arXiv:1508.07370, 2015.

[20] Richard Cole, Vasilis Gkatzelis, and Gagan Goel. Mechanism design for fair division: allocating divisible items without payments. In *EC*, pages 251–268, 2013.

[21] P. Dasgupta, P. Hammond, and E. Maskin. The implementation of social choice rules: Some general results on incentive compatibility. *The Review of Economic Studies*, 46(2):185–216, 1979.

[22] G. Debreu. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, 38:886–893, 1952.

[23] N. Devanur, C. H. Papadimitriou, A. Saberi, and V. V. Vazirani. Market equilibrium via a primal-dual algorithm for a convex program. *JACM*, 55(5), 2008.

[24] Nikhil R. Devanur. Fisher markets and convex programs. Manuscript, 2009.

[25] D. Dolev, D.G. Feitelson, J.Y. Halpern, R. Kupferman, and N. Linial. No justified complaints: on fair sharing of multiple resources. In *ITCS*, pages 68–75, 2012.

[26] P. Dubey and M. Shubik. The noncooperative equilibria of a closed trading economy with market supply and bidding strategies. *JET*, 17:1–20, 1978.

[27] Pradeep Dubey and John Geanakoplos. From Nash to Walras via Shapley-Shubik. *Journal of Mathematical Economics*, 39(5):391–400, 2003.

[28] B. C. Eaves. A finite algorithm for the linear exchange model. *J. Math. Econ.*, 3:197–203, 1976.

[29] E. Eisenberg and D. Gale. Consensus of subjective probabilities: the Pari-Mutuel method. *The Annals of Mathematical Statistics*, 30:165–168, 1959.

[30] H. Fang. Lottery versus all-pay auction models of lobbying. *Public Choice*, 112(3-4):351–71, 2002.

[31] M. Feldman, K. Lai, and L. Zhang. The proportional-share allocation market for computational resources. *ITPDS*, 20(8), 2009.

[32] D. Gale. *Theory of Linear Economic Models*. McGraw Hill, N.Y., 1960.

[33] D. Gale. The linear exchange model. *J. Math. Econ.*, 3:205–209, 1976.

[34] R. Garg, S. Kapoor, and V. V. Vazirani. An auction-based market equilbrium algorithm for the separable gross substitutability case. In *APPROX*, 2004.

[35] C. Georgiou, T. Pavlides, and A. Philippou. Network uncertainty in selfish routing. In *International Parallel and Distributed Processing Symposium*, 2006.

[36] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. Dominant resource fairness: fair allocation of multiple resource types. In *NSDI*, 2011.

[37] G. Giraud. Strategic market games: an introduction. *Journal of Mathematical Economics*, 39:355–375, 2003.

[38] I. L. Glicksberg. A further generalization of the kakutani fixed-point theorem. *Proceedings of the American Mathematical Society*, 3:170–174, 1952.

[39] A. Gutman and N. Nisan. Fair allocation without trade. In *AAMAS*, pages 719–728, 2012.

[40] Aanund Hylland and Richard Zeckhauser. The Efficient Allocation of Individuals to Positions. *Journal of Political Economy*, 87(2):293–314, April 1979.

[41] M. O. Jackson and J. Peck. Asymmetric information in a competitive market game: Reexamining the implications of rational expectations. *Econ. Theory*, 13:603–628, 1999.

[42] K. Jain, V. V. Vazirani, and Y. Ye. Market equilibrium for homothetic, quasi-concave utilities and economies of scale in production. In *SODA*, 2005.

[43] Mamoru Kaneko and Kenjiro Nakamura. The Nash Social Welfare Function. *Econometrica*, 47(2): 423–435, 1979.

[44] C. G. Korpeoglu and S. E. Spear. The market game with production: Coordination equilibrium and price stickiness. CMU, Tepper School of Business, 2015.

[45] A. Matros. Chinese auctions, 2007. Mimeo, University of Pittsburgh.

[46] A. Maya and N. Nisan. Incentive compatible two player cake cutting. In *WINE*, 2012.

[47] Jean-Franois Mertens and Sylvain Sorin (eds). *Game-theoretic methods in general equilibrium analysis.* Springer Science & Business Media, 2013.

[48] B. Moldovanu and A. Sela. The optimal allocation of prizes in contests. *American Economic Review*, 91(3):542–558, 2001.

[49] E. Mossel and O. Tamuz. Truthful fair division. In *SAGT*, pages 288–299, 2010.

[50] H. Moulin. *Fair Division and Collective Welfare.* The MIT Press, 2003.

[51] Shinsuke Nakamura. A feasible Nash implementation of Walrasian equilibria in the two-agent economy. *Economics Letters*, 34(1):5–9, 1990.

[52] J. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, April 1950.

[53] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani. *Algorithmic Game Theory.* Cambridge University Press, Cambridge, Massachusetts, 2007.

[54] J. B. Orlin. Improved algorithms for computing Fisher's market clearing prices. In *STOC*, pages 291–300, 2010.

[55] A. Othman, T. Sandholm, and E. Budish. Finding approximate competitive equilibria: efficient and fair course allocation. AAMAS '10, pages 873–880, Richland, SC, 2010. ISBN 978-0-9826571-1-9.

[56] D.C. Parkes, A.D. Procaccia, and N. Shah. Beyond dominant resource fairness: extensions, limitations, and indivisibilities. In *EC*, pages 808–825, 2012.

[57] Andrew Postlewaite and David Schmeidler. Implementation in differential information economies. *J. of Econ. Theory*, 39:14–33, 1986.

[58] A. D. Procaccia. Cake cutting: Not just child's play. *Communications of the ACM*, 56(7):78–87, 2013.

[59] Ariel D. Procaccia. Cake cutting algorithms, 2014. Draft of Chapter 13 of the Handbook of Computational Social Choice.

[60] Ariel D. Procaccia and Junxing Wang. Fair enough: guaranteeing approximate maximin shares. In *EC*, pages 675–692, 2014.

[61] J.M. Robertson and W.A. Webb. *Cake-cutting algorithms - be fair if you can.* A K Peters, 1998. ISBN 978-1-56881-076-8.

[62] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. Waste makes haste: Bounded time protocols for envy-free cake cutting with free disposal. In *Proceedings of AAMAS*, pages 901–908, 2015.

[63] L. Shapley and M. Shubik. Trade using one commodity as a means of payment. *Journal of Political Economy*, 85(5):937–968, 1977.

[64] G. Tullock. Efficient rent-seeking, in: J.m. buchanan, r.d. tollison, g. tullock (eds.), toward a theory of the rent-seeking society (texas a. & m, university press, college station, tx). pages 97–112, 1980.

[65] H.R. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974.

[66] H.P. Young. *Equity.* Princeton University Press, 1995.

[67] R. Zivan, M. Dudík, S. Okamoto, and K.P. Sycara. Reducing untruthful manipulation in envy-free pareto optimal resource allocation. In *IAT*, pages 391–398, 2010.

A Trading Post: Perfect Complements

Theorem 6 (restated) : The Trading Post mechanism with Leontief utilities has exact pure Nash equilibria if and only if the corresponding Fisher market has market equilibrium prices that are strictly positive everywhere. When this happens, the Nash equilibrium utilities in Trading Post are unique and the price of anarchy is 1.

Proof. Suppose $(\mathbf{v}, \mathbf{B})$ are valuations and budgets for which the Fisher market equilibrium prices are strictly positive everywhere (recall also that the valuations satisfy perfect competition). Then we show that some pure Nash equilibrium exists in the Trading Post game with the same valuations and budgets, and which induces the same allocation.

Let $(\mathbf{p}, \mathbf{x})$ be the Fisher market equilibrium prices and allocation. Define matrix of bids $\mathbf{b}$ by $b_{i,j} = p_j \cdot x_{i,j}$, for all $i \in N, j \in M$. We claim that $\mathbf{b}$ is a pure Nash equilibrium in the Trading Post game. From the conditions in the theorem statement, we have that for each item j, there are two players $i \neq i'$ such that $b_{i,j} \cdot b_{i',j} > 0$. Also the utility of each player i in the market equilibrium is the same as that of strategy profile $\mathbf{b}$ and can be written as follows:

$$u_i(\mathbf{x}_i) = \min_{j \in M : v_{i,j} > 0} \left\{ \frac{x_{i,j}}{v_{i,j}} \right\}$$

By definition of the market equilibrium, each player i gets each item in its demand set in the same fraction f_i, that is

$$f_i = \frac{x_{i,j}}{v_{i,j}}, \text{for } j \in M \text{ with } v_{i,j} > 0$$

Then player i can only improve its utility by taking weight from some item(s) and shifting it towards others in its demand. However, since all the items are received in the same fraction f_i, it follows that player i can only decrease its utility by such deviations. Thus the profile $\mathbf{b}$ is a pure Nash equilibrium of Trading Post.

For the other direction, if a bid profile $\mathbf{b}$ is a pure Nash equilibrium in the Trading Post game, then consider the market allocation and prices $(\mathbf{x}, \mathbf{p})$, where $p_j = \sum_{k=1}^{n} b_{k,j}$ and the induced allocation $\mathbf{x}$ be given by $x_{i,j} = \frac{b_{i,j}}{\sum_{k=1}^{n} b_{k,j}}$. From the perfect competition requirement, clearly $x_{i,j}$ is always well defined. Then at $(\mathbf{x}, \mathbf{p})$ all the goods are allocated, all the money is spent, and each player gets an optimal bundle from its desired goods. To see the latter, note again that a player receives all the items in the same fractions at $(\mathbf{x}, \mathbf{p})$ and since all the prices are positive (since on each item there are at least two non-zero bids), then a player cannot decrease its spending on any item(s). Thus $(\mathbf{x}, \mathbf{p})$ is a market equilibrium and since the market equilibrium utilities are unique, it follows that the same is true for the PNE of Trading Post.

From the correspondence between the Nash equilibria of Trading Post and the market equilibria of the Fisher mechanism, we obtain that on such instances the price of anarchy is 1. □

Theorem 7 (restated) : The parameterized Trading Post mechanism $\mathcal{TP}(\Delta)$ is guaranteed to have a pure Nash equilibrium for every $\Delta > 0$.

Proof. Let $\mathcal{TP}(\Delta)$ be the Trading Post game with minimum fee Δ. We first show that a variant of the game, $\mathcal{TP}'(\Delta)$, where the strategy space of each player i is restricted as follows, must have a pure Nash equilibrium.

- i is forced to bid at least Δ on every item j with the property that $v_{i,j} > 0$
- i must bid zero on every item j for which $v_{i,j} = 0$.

Clearly, the strategy space S_i of each player i is a nonempty compact convex subset of a Euclidean space. Moreover, the utility function of each player is continuous in $\mathbf{x}$ and quasi-concave in the player's own strategy (see, e.g. [11], chapter 2).

We use the following theorem due to Debreu, Glicksberg, and Fan.

Lemma 6. *(Debreu 1952; Glicksberg 1952; Fan 1952) Consider a strategic form game whose strategy spaces S_i are nonempty compact convex subsets of a Euclidean space. If the payoff functions u_i are continuous in s and quasi-concave in s_i, then there exists a pure strategy Nash equilibrium.*

The conditions of Lemma 6 apply and so $\mathcal{TP}'(\Delta)$ has a pure Nash equilibrium $\mathbf{b}^*$. Consider now the Trading Post game with minimum bid Δ, $\mathcal{TP}(\Delta)$. Note the strategy profile $\mathbf{b}^*$ dominates every other strategy in $\mathcal{TP}(\Delta)$, including those that allow the players to bid zero on items of interest to them, since such strategies can only decrease utility. Thus $\mathbf{b}^*$ is also a PNE in $\mathcal{TP}(\Delta)$, which completes the proof. $\square$

B Beyond Pure Nash Equilibria in Trading Post

In this section we show that all Nash equilibria of the trading post game are pure, be it for linear utilities or Leontief utilities. For this we will show that no matter what other players play, there is a *unique pure* best response strategy for player i. At Nash equilibrium, since strategy of a player is a probability distribution over pure strategies that are in best response, this will imply that every Nash equilibrium has to be pure.

To show uniqueness of best response we will show that fixing mixed strategy/bid profile for all other players, player i's utility is a strictly concave function of its own bid profile. Let $\sigma_{-i} = (\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n)$ be a mixed-strategy of all the agents except i, and let agent i's payoff function w.r.t. σ_{-i} be denoted by $u_i^{\sigma_{-i}} : S_i \to \mathbb{R}$, where $S_i = \{(s_{i1}, \dots, s_{im}) \geq 0 \mid \sum_j s_{ij} = B_i\}$ is the set of pure strategies (bids) of player i, i.e., set of all possible ways in which she can split her budget across goods.

Next we derive the result for both linear and Leontief valuation functions separately. Recall that we have assumed *perfect competition* where very good is liked by at least two players, *i.e.*, for each good j, there exist $i \neq k \in N$ such that $v_{i,j} \neq 0, v_{k,j} \neq 0$.

B.1 Perfect Substitutes (Linear valuations)

Function $u_i^{\sigma_{-i}}$ is as follows for the case of linear utilities.

$$u_i^{\sigma_{-i}}(\boldsymbol{s}_i) = \sum_{\boldsymbol{s}_{-i} \in S_{-i}} (\Pi_{k\neq i}\sigma_i(\boldsymbol{s}_k)) \sum_j v_{ij} \frac{s_{ij}}{\sum_{k\neq i} s_{kj} + s_{ij}}, \ \forall \boldsymbol{s}_i \in S_i \tag{15}$$

Next we will show that function $u_i^{\sigma_{-i}}$ is strictly concave on the entire domain of $\mathbb{R}^m$, therefore it is strictly concave on S_i as well. For this, we will show that Hessian of $u_i^{\sigma_{-i}}$ is negative definite.

Lemma 7. *Given a mixed-strategy profile σ_{-i} of all agents $k \neq i$ such that every good is bought by at least one of them with non-zero probability, payoff function of agent i, namely $u_i^{\sigma_{-i}}$, is strictly concave.*

Proof. We first show that $u_i^{\sigma_{-i}}$ is strictly concave on ${\mathbb{R}_+}^m$. Take the derivative of $u_i^{\sigma_{-i}}$ with respect to s_{ij} for each j,

$$\frac{\partial u_i^{\sigma_{-i}}}{\partial s_{ij}} = \sum_{\boldsymbol{s}_{-i} \in S_{-i}} (\Pi_{k\neq i}\sigma_i(\boldsymbol{s}_k)) v_{ij} \frac{\sum_{k\neq i} s_{kj}}{(\sum_{k\neq i} s_{kj} + s_{ij})^2}$$

Differentiating the above w.r.t. s_{ig}s for each good g, we get:

$$\frac{\partial^2 u_i^{\sigma_{-i}}}{\partial s_{ij} \partial s_{ig}} = 0, \ \forall g \neq j; \qquad \frac{\partial^2 u_i^{\sigma_{-i}}}{\partial s_{ij}^2} = -2 \sum_{\boldsymbol{s}_{-i} \in S_{-i}} (\Pi_{k\neq i}\sigma_i(\boldsymbol{s}_k)) v_{ij} \frac{\sum_{k\neq i} s_{kj}}{(s_{ij} + \sum_{k\neq i} s_{kj})^2}$$

Since bids are non-negative and at least one agent other than i is bidding on good j with non-zero probability as per σ_{-i}, the second term in above expression is strictly negative. Therefore, the hessian is diagonal matrix with -ve entries in diagonal, and hence negative definite. Thus, function $u_i^{\sigma_{-i}}$ is strictly concave on $\mathbb{R}^m$. Therefore it remains strictly concave on convex subset $S_i \subset \mathbb{R}^m$ as well. $\square$

Strict concavity of $u_i^{\sigma_{-i}}$ established in Lemma 7 implies that every agent has a unique best response against any strategy of the opponents and it is pure. At Nash equilibrium since only strategies in best response can have non-zero probability, it follows that all Nash equilibria have to be pure. Thus we get the next theorem.

Theorem 12. *In a market with linear utilities, every Nash equilibrium of the corresponding trading post game is pure.*

Proof. At any given strategy profile, if only agent i is biding on a good j with non-zero probability, then it can not be Nash equilibrium since that i can reduce its bid on good j to very small amount while still getting it fully, and use this saved money to buy other goods. $\square$

We observe that the equilibria of Trading Post with linear utilities are not necessarily unique.

Example 3. *Consider a market with four buyers and two goods. Players 1 and 2 only want the first and second good, respectively, while players 3 and 4 like both goods equally. Then the bid profiles* $\boldsymbol{s}_1 = (1, 0), \boldsymbol{s}_2(0, 1), \boldsymbol{s}_3 = (1-\epsilon, \epsilon), \boldsymbol{s}_4 = (\epsilon, 1-\epsilon)$ *are in NE for any* $0 \leq \epsilon \leq 1$.

B.2 Perfect Complements (Leontief utilities)

In case of Leontief utilities, the payoff of agent i is $\min_j \frac{x_{ij}}{v_{ij}}$, where x_{ij} is the amount of good j agent i gets. Like the linear case, for Leontief valuations too we will show uniqueness of best response, however the approach is different because the utility function $u_i^{\sigma_{-i}}$ (defined below) is no more strictly concave.

$$u_i^{\sigma_{-i}}(\boldsymbol{s}_i) = \sum_{\boldsymbol{s}_{-i} \in S_{-i}} (\Pi_{k \neq i} \sigma_i(\boldsymbol{s}_k)) \min_j \frac{1}{v_{ij}} \frac{s_{ij}}{\sum_{k \neq i} s_{kj} + s_{ij}}, \quad \forall \boldsymbol{s}_i \in S_i \tag{16}$$

We can show that $u_i^{\sigma_{-i}}$ is concave using the concavity of $\frac{s_{ij}}{\sum_{k \neq i} s_{kj} + s_{ij}}$ w.r.t. s_{ij}, however it is not strictly concave as we take min of the latter. Instead we will show that function $u_i^{\sigma_{-i}}$ has a unique optimum over all the possible strategies of player i in game $TP(\Delta)$. This will suffice to show that all equilibria are pure.

Since every good is liked by at least two agents, relevant bid profiles $\boldsymbol{s}_{-i}$ are only those where every good is bid on by some agent k other than i since bidding zero fetches zero amount of the good. To capture this we we define *valid* strategy profiles.

Definition 3. *Profile* $\boldsymbol{s}_{-i}$ *is* valid *if* $\forall j, \ \sum_{k \neq i} s_{kj} > 0$. *Similarly, mixed-profiles* σ_{-i} *is said to be* valid *if each pure profile* $\boldsymbol{s}_{-i}$ *with* $P(\boldsymbol{s}_{-i}) = \Pi_{k \neq i} \sigma_k(\boldsymbol{s}_k) > 0$ *is valid.*

Lemma 8. *If* σ *is a Nash equilibrium of game* $TP(\Delta)$ *for any* $\Delta > 0$, *then profile* σ_{-i} *is valid for each* $i \in \mathcal{A}$.

Proof. For some i if σ_{-i} is not valid, then there exists $\boldsymbol{s}_{-i}$ with $P(\boldsymbol{s}_{-i}) > 0$ and good j such that no one is bidding on it at $\boldsymbol{s}_{-i}$. Let $k \neq i$ is an agent with $v_{kj} > 0$. Clearly she gets zero utility whenever she plays $\boldsymbol{s}_k$. Instead if she replaces $\boldsymbol{s}_k$ with $\boldsymbol{t}$ where $t_j = \Delta$, $t_{j'} = s_{kj'} - \Delta$ where $s_{kj'} \geq 2\Delta$ (assuming Δ to be small there is such a good), and $t_d = s_{kd}, \forall d \neq j, j'$, will give her strictly better utility. $\square$

Due to Lemma 8 it suffice to consider only valid strategy profiles, both mixed as well as pure.

Lemma 9. *Given a valid strategy profile* $\boldsymbol{s}_{-i}$ *of* $TP(\Delta)$ *for* $\Delta > 0$ *and good* j, *consider function* $f_j = \frac{1}{v_{ij}} \frac{s_{ij}}{\sum_{k \neq i} s_{kj} + s_{ij}}$. f_j *seen as a function of* s_{ij} *is strictly concave, and as a function of* $(s_{i1}, \ldots, s_{in})$ *it is concave.*

Proof. Taking double derivative of f_j we get,

$$\frac{\partial^2 f_j}{\partial s_{ij}^2} = \frac{-2}{v_{ij}} \frac{\sum_{k \neq i} s_{kj}}{(s_{ij} + \sum_{k \neq i} s_{kj})^2}$$
$$\forall k \neq j, \ \frac{\partial^2 f_j}{\partial s_{ij} \partial s_{ik}} = 0$$

First equality implies strict concavity w.r.t. s_{ij} since $\sum_{k \neq i} s_{kj} \geq \Delta$ given that there is another agent who would want good j. Both equalities together imply that Hessian of f_j is negative semi-definite and there by concavity w.r.t. $\boldsymbol{s}_i$. $\square$

Next we show uniqueness of best response for agent i against any given σ_{-i}.

Lemma 10. *Given a valid mixed-strategy profile* σ_{-i} *of all agents* $k \neq i$ *in trading post game* $TP(\Delta)$ *for any* $\Delta > 0$, *payoff function of agent* i, *namely* $u_i^{\sigma_{-i}}$, *has a unique optimum.*

Proof. First it is easy to see using Lemma 9 that each of the term inside summation of $u_i^{\sigma-i}$ is a concave function w.r.t. $\boldsymbol{s}_i$, since minimum of a set of concave functions is a concave function. And by the same argument $u_i^{\sigma-i}$ itself is concave because it is summation of a set of concave functions.

To the contrary suppose there are two optimum $\boldsymbol{s}_i, \boldsymbol{s}'_i \in S_i$, $\boldsymbol{s}_i \neq \boldsymbol{s}'_i$. Clearly $\sum_j s_{ij} = \sum_j s'_{ij} = m_i$. Therefore, there exists a good where bid at $\boldsymbol{s}$ is more than that at $\boldsymbol{s}'$. Let $j^* \in \{j \mid s_{ij} > s'_{ij}\}$. Clearly, $s_{ij^*} > \Delta$.

Claim 1. *Given $\boldsymbol{s}$, if there is a good j with $s_{ij} > \Delta > 0$ and $u_i(\boldsymbol{s}_i, \boldsymbol{s}_{-i}) < \frac{1}{v_{ij}} \frac{s_{ij}}{\sum_{k\neq i} s_{kj} + s_{ij}}$, then $\boldsymbol{s}_i$ is not a best response to $\boldsymbol{s}_{-i}$.*

Proof. It is easy to see that there exists δ, such that for $t_j = s_{ij} - n\delta \geq \Delta$, $t_k = s_{ik} + \delta$, $\forall g \neq j$, we have $u_i(\boldsymbol{s}) < u_i(\boldsymbol{t}, \boldsymbol{s}_{-i})$. □

$\forall \boldsymbol{s}_{-i}$ with $P(\boldsymbol{s}_{-i}) > 0$ if we have $u_i(\boldsymbol{s}_i, \boldsymbol{s}_{-i}) < \frac{1}{v_{ij^*}} \frac{s_{ij^*}}{\sum_{k\neq i} s_{kj} + s_{ij^*}}$, then from the above claim it follows that $\boldsymbol{s}_i$ is not an optimal solution of $u_i^{\sigma-i}$.

Otherwise for some $\boldsymbol{s}_{-i}$ we have $u_i(\boldsymbol{s}_i, \boldsymbol{s}_{-i}) = \frac{1}{v_{ij^*}} \frac{s_{ij^*}}{\sum_{k\neq i} s_{kj} + s_{ij^*}}$. Due to concavity of $u_i^{\sigma-i}$ we have that entire line-segment from $\boldsymbol{s}_i$ to $\boldsymbol{s}'_i$ is optimal. Call this line segment $\mathcal{L}$. On this line-segment bid on good j is strictly changing (decreasing). Since $u_i(\boldsymbol{s}_i, \boldsymbol{s}_{-i}) = \min_j \frac{1}{v_{ij}} \frac{s_{ij}}{\sum_{k\neq i} s_{kj} + s_{ij}}$ is governed by the utility obtained from good j^*, it is strictly concave on $\mathcal{L}$ at $\boldsymbol{s}_i$.

Furthermore, at a point if a set of functions are concave and one of them is strictly concave then their summation is strictly concave. Therefore, $u_i^{\sigma-i}$ is strictly concave at $\boldsymbol{s}_i$ on $\mathcal{L}$ and hence either $\boldsymbol{s}_i$ is not optimum or other points on $\mathcal{L}$ are not optimum. In either case we get a contradiction. □

Since at Nash equilibrium only optimal strategies can have non-zero probability, the next theorem follows using Lemmas 8 and 10.

Theorem 13. *For market with Leontief utilities, every Nash equilibrium of the corresponding $\Delta > 0$ trading post game $TP(\Delta)$ is pure.*

The Trading Post game with Leontief valuations does not always have unique PNE, as can be seen from the next example.

Example 4. *Let there be an instance of Trading Post with players $N = \{1, 2\}$, items $M = \{1, 2\}$, and Leontief valuations $v_{i,j} = 1$, $\forall i, j \in \{1, 2\}$. Then every strategy profile of the form $v_{1,1} = v_{2,1} = a$, with $a \in (0, 1)$ is a pure Nash equilibrium of Trading Post, since both players get the items in the optimal ratios.*