



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Nash Social Welfare Approximation for Strategic Agents

Simina Brânzei, Vasilis Gkatzelis, Ruta Mehta

To cite this article:

Simina Brânzei, Vasilis Gkatzelis, Ruta Mehta (2022) Nash Social Welfare Approximation for Strategic Agents. Operations Research 70(1):402-415. <https://doi.org/10.1287/opre.2020.2056>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2021, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Methods

Nash Social Welfare Approximation for Strategic Agents

Simina Brânzei,^a Vasilis Gkatzelis,^b Ruta Mehta^c

^aComputer Science, Purdue University, West Lafayette, Indiana 47407; ^bComputer Science, Drexel University, Philadelphia, Pennsylvania 19104; ^cComputer Science, University of Illinois at Urbana–Champaign, Champaign, Illinois 61801

Contact: simina@purdue.edu (SB); gkatz@drexel.edu,  <https://orcid.org/0000-0001-7203-438X> (VG); rutamehta@cs.illinois.edu (RM)

Received: September 20, 2018

Accepted: June 11, 2020

Published Online in Articles in Advance:
February 17, 2021

Area of Review: Decision Analysis, Games/
group decision, Noncooperative

<https://doi.org/10.1287/opre.2020.2056>

Copyright: © 2021 INFORMS

Abstract. A central goal in the long literature on fair division is the design of mechanisms that implement fair outcomes, despite the participants' strategic behavior. We study this question by measuring the fairness of an allocation using the geometric mean of the agents' values, known as the *Nash social welfare* (NSW). This objective is maximized by widely known concepts such as the Nash bargaining solution, proportional fairness, and the competitive equilibrium with equal incomes; we focus on (approximately) implementing this objective and analyze the Trading Post mechanism. We consider allocating goods that are substitutes or complements and show that this mechanism achieves an approximation of two for concave utility functions and becomes essentially optimal for complements, where it can reach $(1 + \epsilon)$ for any $(\epsilon > 0)$. Moreover, we show that the Nash equilibria of this mechanism are pure and provide individual fairness in the sense of proportionality.

Funding: This project has received funding from the European Research Council under the European Union's H2020 European Research Council Research and Innovation Programme [Grant 740282]. S. Brânzei was supported in part by the Israel Science Foundation [Grant 1435/14] administered by the Israeli Academy of Sciences and by the United States-Israel Binational Science Foundation [Grant 2014389]. V. Gkatzelis was supported by the Division of Computing and Communication Foundations [Grants 1408635, 1216073, and 1161813] and by the National Science Foundation [Grant CCF-175595]. R. Mehta was supported by the National Science Foundation [CAREER Grant CCF-1750436].

Supplemental Material: The e-companion is available at <https://doi.org/10.1287/opre.2020.2056>.

Keywords: Nash social welfare • Trading Post • price of anarchy

1. Introduction

We study a setting where m divisible goods need to be distributed among n strategic agents with concave valuations, and we analyze the classic Trading Post mechanism, originally introduced by Shapley and Shubik (Shapley and Shubik 1977). Our analysis shows that every Nash equilibrium of the game induced by this mechanism is “near optimal,” in the sense that the Nash social welfare (NSW; the geometric mean of the agents' utilities) is always within a small factor of the best possible, indicating that the equilibria strike a desired balance between fairness and efficiency.

Our work builds on a literature that studies fair division from a mathematical standpoint, which began with the work of Steinhaus during the Second World War. This led to an extensive and growing body of work on fair division protocols within economics and political science (e.g., Young 1995, Brams and Taylor 1996, Robertson and Webb 1998, Moulin 2003, Barbanel 2005). Recent years have seen an increased amount of work on fair division coming from operations research (e.g., Bertsimas et al. 2011, 2012) and computer science (e.g., for a recent survey, see Brandt et al. 2016, part II), which has focused on the

computational and incentive issues that arise when designing resource allocation mechanisms that are fair.

Although identifying an appropriate notion of fairness for each setting can be a nontrivial task, a large body of work in economics, particularly social choice theory, is concerned with this very question, with numerous solution concepts proposed in response. Our fairness measure of choice herein is the *Nash social welfare*, which dates back to the fifties (Nash 1950, Kaneko and Nakamura 1979) and was proposed by Nash as a solution for bargaining problems, using an axiomatic approach. In a setting where n agents are “bargaining” over a set of possible outcomes and each outcome x yields utility $u_i(x)$ for each agent i , the NSW objective points to the outcome with maximum geometric mean of these utilities as a fair compromise.

Using this objective as a measure of fairness, we seek to maximize it, and the main obstacle that we focus on is informational: to optimize this objective, we need to know the preferences of the participants. When these preferences are private information of the agents, we need to use some type of mechanism that interacts with the agents in order to reach a desired outcome. The obvious candidate would be the direct

revelation mechanism known as the *Fisher market mechanism*: ask the agents to report their preferences and then compute the NSW-maximizing allocation based on their reports. Unfortunately, it is well known that agents can feign different interests and eventually get better allocations (Adsul et al. 2010; Chen et al. 2011, 2012; Babaioff et al. 2014; Brânzei et al. 2014). This strategic behavior can cause the mechanism to compute NSW-maximizing outcomes based on misreported preferences that have little to do with reality, leading to unfair allocations. In this work, we address the following basic question.

To what extent can the Nash social welfare be maximized when the participants are strategic?

We analyze a natural mechanism, known as the *Trading Post* mechanism, and show that even for small markets it achieves outcomes that closely approximate the optimal NSW on every instance while simultaneously satisfying additional individual fairness guarantees. To formally evaluate the quality of this mechanism, we use the standard *Price of Anarchy* (PoA) measure (Roughgarden and Tardos 2002, Koutsoupias and Papadimitriou 2009), which corresponds to the ratio between the optimal NSW and the NSW of the worst Nash equilibrium outcome obtained by the mechanism.

Along with other standard welfare objectives, the NSW objective is captured by a 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 $M_0(\mathbf{x})$, the limit of $M_p(\mathbf{x})$ as p goes to zero (i.e., $(\prod_i u_i(\mathbf{x}))^{1/n}$). Note that optimizing the NSW objective is equivalent to maximizing the sum of the logarithms of the agent utilities (i.e., $\sum_i \log(u_i(\mathbf{x}))$; the latter can be derived via a logarithmic transformation of the former), which is a very commonly used objective for achieving a notion called proportional fairness in networking (Kelly 1997, Kelly et al. 1998) and beyond (Bertsimas et al. 2011, 2012).

Although an extended treatment of the NSW objective can be found, for example, in Moulin (2003), we highlight the fact that it achieves a *natural compromise between individual fairness and efficiency*. This is in contrast to two other well-studied functions captured by $M_p(\mathbf{x})$: (i) the *egalitarian* (max-min) objective attained as $p \rightarrow -\infty$ and (ii) the *utilitarian* (average) objective attained at $p = 1$, which correspond to extreme fairness and extreme efficiency, respectively. The former may cause vast inefficiencies, whereas the latter can completely neglect how unhappy some agents might be. The NSW objective lies between these two extremes and strikes a balance between

them because maximizing the geometric mean leads to more balanced valuations, but without neglecting efficiency. For example, given two outcomes $\mathbf{x}$ and $\mathbf{x}'$ where the former has higher NSW than the latter, if some agent i prefers $\mathbf{x}'$ because its utility there is $q > 1$ times higher, then it must be that a transition from $\mathbf{x}$ to $\mathbf{x}'$ must be causing the geometric mean of the other agents' utility to drop by a factor at least q . This suggests that agent i may be harder to satisfy, and it provides a formal justification for choosing $\mathbf{x}$ over $\mathbf{x}'$. Also, note that this rationale naturally prioritizes low utility agents as it is commonly easier to increase their utility by a larger factor.

1.1. Our Results

The Shapley and Shubik mechanism (Shapley and Shubik 1977) (see Section 7 for a discussion of related work) gives each participant a budget of some artificial currency and asks them to distribute this budget across the available goods. After the agents choose how to distribute their budget, they receive a fraction from each good that is proportional to the amount they spent on it. Because the agents are strategic, we assume that they distribute their budget aiming to get an allocation that maximizes their utility, anticipating the strategic behavior of the other agents. Therefore, using a game-theoretic analysis, we evaluate the performance of this mechanism over the Nash equilibria of the induced game. To ensure the existence of pure Nash equilibria, we enhance the mechanism with an arbitrarily small parameter $\epsilon > 0$, such that each agent needs to spend at least ϵ amount of budget on an item in order to be allocated some portion of it.

To model different types of goods and agent preferences over the outcomes, we consider large classes of extensively studied agent valuation functions. Common examples include linear (or additive) valuations and Leontief valuations (Gale 1960, Codenotti and Varadarajan 2004, Nisan et al. 2007). Additive valuations capture the agent preferences in settings where the goods are *perfect substitutes* (i.e., that can replace each other in consumption, such as Pepsi and Coca-Cola). On the other hand, Leontief valuations capture *perfect complements* (i.e., goods that have no value without each other, such as a left shoe and a right shoe). Both of these valuation functions are concave over the set of possible (fractional) bundles allocated to the bidder. We measure the performance of the Trading Post mechanism with respect to the NSW objective, and our main result is:

Theorem (Informal). *Any Nash equilibrium of the Trading Post mechanism approximates the optimal Nash social welfare within a factor of two for all concave valuations. For Leontief valuations, the mechanism achieves an approximation of $1 + \epsilon$ for an arbitrarily small constant $\epsilon > 0$.*

Note that for Leontief valuations, the Trading Post can get arbitrarily close to the optimal solution! As we discuss in Section 7, this result can have a significant impact in computing environments where Leontief valuations are commonly used. Overall, our results testify to the usefulness and robustness of the Trading Post mechanism: a good NSW approximation implies that the geometric mean of agent utilities is high, and so, happiness is well distributed across the participants.

In Section 5, we complement our aforementioned results regarding the NSW (an *aggregate* measure of fairness) by also analyzing the Trading Post mechanism with respect to *proportionality*, one of the fundamental notions of *individual* fairness. Proportionality (not to be confused with proportional fairness) requires that every agent i who has a value of V_i for the bundle containing all of the goods receives utility at least $\frac{1}{n} V_i$. We show that the Trading Post mechanism guarantees proportionality in any equilibrium, even for general concave valuations.

Theorem (Informal). *Every Nash equilibrium of the Trading Post mechanism satisfies proportionality for all concave valuations.*

In conjunction with the bounds regarding the NSW approximation, the proportionality results imply that the Trading Post mechanism guarantees a surprising combination of both individual and aggregate fairness for very general utilities. Furthermore, all of the results regarding the NSW and proportionality also hold for the weighted version of this problem where each agent i may be endowed with a different budget B_i that signifies the agent's importance or priority.

Finally, we prove that under some mild conditions, the set of mixed Nash equilibria of the Trading Post mechanism coincides with the set of pure Nash equilibria even for general concave valuations, which extends our (pure) PoA bounds for this mechanism to mixed PoA. Also, in the process of obtaining a near-optimal bound for Leontief in the Trading Post mechanism, we show that ϵ -approximate market equilibria for Leontief utilities approximate its Nash social welfare by a factor of $\frac{1}{(1+\epsilon)}$. We believe that this result may be of independent interest.

2. Preliminaries

Let $N = \{1, \dots, n\}$ be a set of agents and $M = \{1, \dots, m\}$ a set of divisible goods. An allocation $\mathbf{x}$ is a distribution of the goods to the agents such that $x_{i,j}$ represents the amount of good j received by agent 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 $\mathcal{F} = \{\mathbf{x} \in [0, 1]^{n \times m} \mid \sum_{i=1}^n x_{i,j} = 1 \ \forall j \in M\}$. Agent i 's value (or

utility) for a bundle of goods is represented by a nondecreasing, nonnegative, and concave valuation function $u_i : [0, 1]^m \rightarrow \mathbb{R}_{\geq 0}$.

2.1. Valuation Functions

Two very common and extensively studied valuation functions are *linear valuations* (where the goods are *perfect substitutes*) and *Leontief valuations* (where the goods are *perfect complements*). For both, let $\mathbf{v}_i = (v_{i,1}, \dots, v_{i,m}) \in \mathbb{R}_{\geq 0}^m$ be a vector of valuations for agent i , where $v_{i,j}$ captures how much agent i likes good j .

Perfect substitutes, defined through linear (additive) valuations, represent goods that can replace each other in consumption. The utility of agent i for bundle $\mathbf{x}_i$ is $u_i(\mathbf{x}_i) = \sum_{j=1}^m v_{i,j} \cdot x_{i,j}$.

Perfect complements, represented by Leontief utilities, capture scenarios where one good may have no value without the other, such as a left shoe and a right shoe, or the central processing unit time and computer memory required for the completion of a computing task. Given a matrix of $v_{i,j} \geq 0$ values, for each agent i let $D_i = \{j \in M \mid v_{i,j} > 0\}$ be the set of items that the agent is interested in. In the Leontief model, the utility of agent i for a bundle $\mathbf{x}_i$ is $u_i(\mathbf{x}_i) = \min_{j \in D_i} \{x_{i,j}/v_{i,j}\}$: that is, agent i desires the items in the proportions $v_{i,1} : v_{i,2} : \dots : v_{i,m}$ and gains no additional value from an allocation increase in one good unless the other goods are also increased according to these proportions.

A more general class of valuations, which captures perfect substitutes and complements, is constant elasticity of substitution (CES). CES valuations take the form $u_i(\mathbf{x}_i) = (\sum_{j=1}^m v_{i,j} \cdot x_{i,j}^\rho)^{\frac{1}{\rho}}$ and interpolate between additive and Leontief through the parameter ρ such that $-\infty < \rho \leq 1$, $\rho \neq 0$. Leontief and additive utilities are obtained when ρ approaches $-\infty$ and equals one, respectively.

All CES valuations are, in turn, part of a more general class of *concave* valuation functions. In this paper, we focus on nonnegative, nondecreasing, and concave valuation functions but also consider special classes, like the ones mentioned or ones that arise if we impose additional conditions, such as strict monotonicity (any additional conditions are mentioned in our statements).

Given an allocation $\mathbf{x}$, we may use $u_i(\mathbf{x})$ to denote $u_i(\mathbf{x}_i)$. Also, if some agent's utility $u_i(\mathbf{x})$ is equal to the utility of that agent for being allocated all of the m items fully, then we refer to this agent as *satiated at allocation $\mathbf{x}$* .

2.2. Fisher Market Equilibrium

Our setting is closely related to the Fisher market model, and throughout the paper, we make use of the market equilibrium notion from this model. Given a set of agents with valuation functions u_i and budgets

$B_i \geq 1$, a Fisher market equilibrium is an allocation $\mathbf{x}$ along with prices $(p_j)_{j \in M}$, such that

1. For each agent $i \in N$, $\mathbf{x}_i \in \arg\max_{\mathbf{x} \geq 0, \mathbf{x}^T \mathbf{p} \leq B_i} u_i(\mathbf{x})$ (agents buy the best bundle they can afford).
2. For each good $j \in M$, $\sum_{i \in N} x_{i,j} = 1$ if $p_j > 0$; otherwise, $\sum_{i \in N} x_{i,j} \leq 1$ (demand meets supply).

2.3. The Trading Post Mechanism

The Trading Post mechanism asks the agents to choose how to distribute their budgets. After 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 agent i is $\mathcal{S}_i = \{\mathbf{b}_i \in [0, B_i]^m \mid \sum_{j=1}^m b_{i,j} = B_i\}$. Given a bid profile $\mathbf{b} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$, the induced prices $\mathbf{p}$ are $p_j = \sum_{i \in N} b_{i,j}$ for all $j \in M$, and the allocation $\mathbf{x}$ is such that

$$x_{i,j} = \begin{cases} \frac{b_{i,j}}{\sum_{k=1}^n b_{k,j}} & \text{if } b_{i,j} > 0 \\ 0 & \text{otherwise.} \end{cases}$$

2.4. The Trading Post Mechanism with Entrance Fee Δ

The Trading Post mechanism with entrance fee $\Delta \in (0, 1/m]$, denoted as $\mathcal{TP}(\Delta)$, asks the agents to choose how to distribute their budgets, and the strategy space is the same as that of the standard Trading Post mechanism. Given a bid profile $\mathbf{b} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ by the agents, let $\bar{\mathbf{b}} = (\bar{\mathbf{b}}_1, \dots, \bar{\mathbf{b}}_n)$ be the “effective” bid profile, which for every $i \in N$ and $j \in M$, satisfies $\bar{b}_{i,j} = b_{i,j}$ if $b_{i,j} \geq \Delta$ and $\bar{b}_{i,j} = 0$ otherwise. Then, the bid profile $\mathbf{b}$ yields prices $\mathbf{p}$ such that $p_j = \max\{\Delta, \sum_{i \in N} \bar{b}_{i,j}\}$ for all $j \in M$, and the allocation $\mathbf{x}$ is such that

$$x_{i,j} = \begin{cases} \frac{\bar{b}_{i,j}}{\sum_{k=1}^n \bar{b}_{k,j}} & \text{if } \bar{b}_{i,j} > 0 \\ 0 & \text{otherwise.} \end{cases}$$

2.5. Strategic Behavior and the Nash Equilibrium

In analyzing the performance of the Trading Post mechanism, throughout the paper we consider agents that are strategic and will therefore choose how to distribute their budget aiming to maximize their utility (in response to the corresponding strategic choices of the other agents). A *pure Nash equilibrium* of the Trading Post mechanism is a profile of bids $\mathbf{b} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ such that every agent’s utility in the induced allocation is greater than or equal to the highest utility that they could achieve through a unilateral deviation. A pure Nash equilibrium is not guaranteed to exist for all instances in our setting, but we prove its existence for large families of instances in Section 6.

2.6. Nash Social Welfare and the Price of Anarchy

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:

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

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 that this is the objective maximized by the Fisher market equilibrium solution when the agents have homogeneous valuation functions, which include all of CES and Leontief functions. We slightly override notation and refer to the weighted objective as the NSW as well. If $B_i \geq 1$ is the weight (budget) of agent i and $\mathcal{B} = \sum_{i=1}^n B_i$ is the total weight, the induced market equilibrium in Fisher’s model maximizes the objective:

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

Note that we get back the original definition when all agents have the same weight/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 Trading Post mechanism using the *price of anarchy* with respect to the NSW objective. Given a problem instance $\mathcal{I}$ and a set of pure Nash equilibria E , the PoA of the Trading Post mechanism for $\mathcal{I}$ is the maximum ratio between the optimal NSW—obtained at some allocation $\mathbf{x}^*$ —and the NSW at an allocation $\mathbf{x} \in E$ (i.e., the ratio of the optimal Nash social welfare and the Nash social welfare in the *worst* Nash equilibrium: $\text{PoA}(\mathcal{I}) = \max_{\mathbf{x} \in E} \left\{ \frac{\text{NSW}(\mathbf{x}^*)}{\text{NSW}(\mathbf{x})} \right\}$). The price of anarchy of the mechanism for a class of instances is the maximum of this value over all possible instances in this class: $\max_{\mathcal{I}} \{\text{PoA}(\mathcal{I})\}$.

3. Price of Anarchy of Trading Post for Concave Valuations

In this section, we prove an upper bound of two on the PoA for Trading Post mechanism with concave valuation functions. We first prove a lemma regarding the structure of the Nash equilibria in the Trading Post mechanism.

Lemma 1. *In any Nash equilibrium allocation $\mathbf{x}$ of the Trading Post mechanism, every agent who is not satiated at $\mathbf{x}$ and has a valuation function that is nonnegative, nondecreasing, and concave is spending only on items that at least one other agent is also spending on.*

Proof. Assume that there exists a Nash equilibrium allocation $\mathbf{x}$ where some agent i that is not satiated at $\mathbf{x}$ is spending $b_{i,j} > 0$ on an item j that no other agent is spending on (and is thus receiving all of that item). Given that agent i is not satiated at $\mathbf{x}$, there exists some other bundle $\mathbf{x}'_i$ such that $u_i(\mathbf{x}'_i) > u_i(\mathbf{x}_i)$. Now, consider the segment connecting $\mathbf{x}_i$ to $\mathbf{x}'_i$ in $\mathbb{R}^m$: because the utility function of agent i is concave, the utility for any allocation $\mathbf{y}_i$ on this segment, excluding $\mathbf{x}_i$ itself, satisfies $u_i(\mathbf{y}_i) > u_i(\mathbf{x}_i)$. Also, because agent i is the only one spending on item j , a unilateral deviation where that agent reduces the spending on j to, say, $b_{i,j}/2$ would still allow that agent to keep all of item j . However, then the agent could appropriately use this spending on other items in order to change its allocation to some $\mathbf{y}_i$ such as the one described. Specifically, if $M' = \{j \in M \mid x'_{i,j} > x_{i,j}\}$ is the subset of items that agent i gets more of in $\mathbf{x}'_i$ than in $\mathbf{x}_i$ (this set includes no items on which only agent i is spending because i gets those fully in $\mathbf{x}$) and $p_j = \sum_{i \in N} b_{i,j}$ is the total spending for each item in the Nash equilibrium that induced $\mathbf{x}$, then agent i can compute the amount $\tilde{b}_j$ to spend on each item $j \in M'$ by solving the following program:

$$\max \alpha_i \quad \text{such that} \quad \alpha_i (x'_{i,j} - x_{i,j}) = \frac{\tilde{b}_{i,j}}{p_j + \tilde{b}_{i,j}} \quad \text{and} \\ \tilde{b}_{i,j} \geq 0, \quad \forall j \in M'; \quad \sum_{j \in M'} \tilde{b}_{i,j} = b_{i,j}/2.$$

The solution to this program will return some $\alpha_i > 0$, and the resulting allocation after this deviation will be $y_{ij} = x_{i,j} + \alpha_i(x'_{i,j} - x_{i,j}) = (1 - \alpha_i)x_{i,j} + \alpha x'_{i,j}$ for all $j \in M'$ and $y_{ij} = x_{i,j}$ for all $j \in M \setminus M'$. Because the utility of the agent will be strictly higher after this unilateral deviation, this contradicts the initial assumption that $\mathbf{x}$ is a Nash equilibrium, thus proving the lemma. $\square$

In order to prove the PoA upper bound, we will also be using the folklore weighted arithmetic and geometric mean inequality stated in the following lemma.

Lemma 2. For any nonnegative numbers $\rho_1, \rho_2, \dots, \rho_n$ and nonnegative weights $w_1, w_2, \dots, w_n$ with $W = \sum_{i=1}^n w_i$, we have $(\prod_{i=1}^n \rho_i^{w_i})^{1/W} \leq \frac{\sum_{i=1}^n \rho_i w_i}{W}$ (i.e., the weighted geometric mean of these numbers is always less than or equal to their weighted arithmetic mean).

We are now ready to prove the main result of this section.

Theorem 1. The Trading Post mechanism with concave valuations has price of anarchy at most two.

Proof. Given a problem instance with concave valuations, let $\mathbf{x}^*$ be the allocation that maximizes the Nash social welfare subject to supply constraints. And let $\tilde{\mathbf{x}}$ be the allocation obtained under a Nash equilibrium where $\tilde{\mathbf{b}}$, where the price of each item j is $\tilde{p}_j = \sum_i \tilde{b}_{i,j}$.

Let i be any one of the agents that is not satiated in $\tilde{\mathbf{x}}$, and for each item j , let $p'_j = \tilde{p}_j - \tilde{b}_{i,j}$ be the total bid on item j from agents other than i . Note that, by Lemma 1, $p'_j > 0$ for every j ; otherwise, $\tilde{\mathbf{b}}$ would not be an equilibrium because agent i could bid less on item j and still receive all of it. Now, consider the allocation $\mathbf{x}'$ that arises if every agent $k \neq i$ bids $\tilde{\mathbf{b}}_k$ (i.e., according to the equilibrium), whereas agent i unilaterally deviates to a bid of $b'_{i,j}$ for each item j such that for some $\alpha_i > 0$ and every item j :

$$\frac{b'_{i,j}}{p'_j + b'_{i,j}} = \alpha_i \cdot x^*_{i,j}. \quad (1)$$

This bid, $\mathbf{b}'_i$, corresponds to a solution of the following program:

$$\max \alpha_i \quad \text{s.t.} \quad \alpha_i \cdot x^*_{i,j} = \frac{b'_{i,j}}{p'_j + b'_{i,j}} \quad \text{and} \\ b'_{i,j} \geq 0, \quad \forall j \in M; \quad \sum_{j \in M} b'_{i,j} = B_i.$$

The allocation induced by this unilateral deviation of i is

$$x'_{i,j} = \frac{b'_{i,j}}{p'_j + b'_{i,j}} = \alpha_i \cdot x^*_{i,j}. \quad (2)$$

If $\alpha_i \leq 1$, the utility of agent i after this deviation is $u_i(\alpha_i \cdot \mathbf{x}^*_i) \geq \alpha_i u_i(\mathbf{x}^*)$ because of the concavity of u_i . If, on the other hand, $\alpha_i > 1$, then $u_i(\alpha_i \cdot \mathbf{x}^*_i) \geq u_i(\mathbf{x}^*)$ because the utility is nondecreasing. Therefore, $u_i(\alpha_i \cdot \mathbf{x}^*_i) \geq \min\{\alpha_i, 1\} u_i(\mathbf{x}^*)$. However, $\tilde{\mathbf{x}}$ is a Nash equilibrium, so this deviation of agent i cannot yield a higher utility for i , which implies that

$$u_i(\tilde{\mathbf{x}}) \geq u_i(\alpha_i \cdot \mathbf{x}^*_i) \Rightarrow u_i(\tilde{\mathbf{x}}) \geq \min\{\alpha_i, 1\} u_i(\mathbf{x}^*) \\ \Rightarrow \frac{u_i(\mathbf{x}^*)}{u_i(\tilde{\mathbf{x}})} \leq \frac{1}{\min\{\alpha_i, 1\}}. \quad (3)$$

By Equation (2), we get $\sum_j b'_{i,j} = \sum_j x'_{i,j}(p'_j + b'_{i,j})$. Because $\sum_j b'_{i,j} = B_i$, this implies that $B_i = \sum_{j=1}^m x'_{i,j}(p'_j + b'_{i,j})$. Substituting for $x'_{i,j} = \alpha_i \cdot x^*_{i,j}$, we get

$$\frac{B_i}{\alpha_i} = \sum_{j=1}^m x^*_{i,j}(p'_j + b'_{i,j}). \quad (4)$$

Also, we have $0 \leq x^*_{i,j} \leq 1$, and $p'_j \leq \tilde{p}_j$ for every item j . As a result, for each agent i we have

$$\sum_{j=1}^m x^*_{i,j}(p'_j + b'_{i,j}) \leq \sum_{j=1}^m (x^*_{i,j} \tilde{p}_j) + \sum_{j=1}^m (x^*_{i,j} b'_{i,j}) \\ \leq \sum_{j=1}^m x^*_{i,j} \tilde{p}_j + \sum_{j=1}^m b'_{i,j} \leq \sum_{j=1}^m (x^*_{i,j} \tilde{p}_j) + B_i. \quad (5)$$

To conclude the proof, we will now use the fact that inequalities (3)–(5) hold for every agent i that is not satiated in $\tilde{\mathbf{x}}$. Note that for every agent i that is satiated in $\tilde{\mathbf{x}}$, we have $u_i(\tilde{\mathbf{x}}) \geq u_i(\mathbf{x}^*)$. Let $N_g \subseteq N$ be the set of agents for which $\alpha_i \leq 1$ and N_ℓ be the set of agents with $\alpha_i > 1$ as well as the agents that are satiated in $\tilde{\mathbf{x}}$. Using this partition of the agents, we get

$$\begin{aligned} \sum_{i=1}^n B_i \frac{u_i(\mathbf{x}_i^*)}{u_i(\tilde{\mathbf{x}}_i)} &\leq \sum_{i=1}^n \frac{B_i}{\min\{\alpha_i, 1\}} \\ &\leq \sum_{i \in N_\ell} B_i + \sum_{i \in N_g} \left(\sum_{j=1}^m (x_{ij}^* \tilde{p}_j) + B_i \right) \leq 2\mathcal{B}. \end{aligned} \quad (6)$$

Using Equation (6) and the inequality of weighted arithmetic and geometric means provided by Lemma 2, we conclude that

$$\left(\prod_{i \in N} \left(\frac{u_i(\mathbf{x}_i^*)}{u_i(\tilde{\mathbf{x}}_i)} \right)^{B_i} \right)^{1/\mathcal{B}} \leq 2. \quad \square$$

The next theorem complements the upper bound on the price of anarchy by providing a lower bound of approximately 1.445, even for the special case of linear valuations.

Theorem 2. *The Trading Post mechanism has a price of anarchy no better than $e^{1/e} \approx 1.445$, even for the special case of linear valuations.*

Proof. Using two positive integers q and $k < q$ as parameters, whose values are set later, we construct a market with $q+1$ goods and $q+2$ agents with equal budget $B_i = 1$. Each agent $i \leq k$ has value $v_{i,i} = 1$ for good i and value $v_{i,j} = 0$ for all other goods $j \neq i$. On the other hand, every agent $i \in [k+1, q]$, apart from having $v_{i,i} = 1$, also has some small but positive, value $v_{i,j} = a$ for all items $j \leq k$. The rest of that agent's $v_{i,j}$ values are zero. Agent $(q+1)$ has a small but positive value a' for goods $j \in [k+1, q]$ and value 2 for good $(q+1)$. Finally, agent $(q+2)$ values only good $(q+1)$ at value 2. We will set the values of a and a' later. Consider the following strategy $\mathbf{b}$

$$\forall 1 \leq i \leq (q+1), b_{i,i} = 1, \quad b_{(q+2),(q+1)} = 1,$$

and the rest are zero.

In $\mathbf{b}$, every agent $i \in [1, q]$ gets all of good i , whereas agents $(q+1)$ and $(q+2)$ share good $(q+1)$ equally. Thus, the utility of every agent $i \in [1, (q+2)]$ at $\mathbf{b}$ is $u_i(\mathbf{b}) = 1$, giving NSW of one.

Next, we construct an alternative strategy profile $\mathbf{s}$, which for carefully chosen values of a and a' , is a Nash equilibrium and has an NSW that approaches $(1/e)^{1/e}$ as $q \rightarrow \infty$. In this strategy profile $\mathbf{s}$, each agent $i \leq k$

spends all of its budget on good i , and agent $(q+2)$ spends all of its budget on good $(q+1)$. On the other hand, every agent $i \in [k+1, q]$ spends some small amount ϵ on item i and divides the rest of its budget, namely $(1-\epsilon)$, equally among the items in $[1, k]$. Agent $(q+1)$ spends ϵ' on every good $i \in [k+1, q]$ and spends the remaining $1 - (q-k)\epsilon'$ on good $(q+1)$.

This leads to a price of $1 + \frac{(q-k)(1-\epsilon)}{k} = \frac{q(1-\epsilon)+k\epsilon}{k}$ for the first k goods, a price of $(\epsilon + \epsilon')$ for goods $j \in [k+1, q]$, and a price of $(2 - (q-k)\epsilon')$ for good $(q+1)$. It is easy to verify that the first k agents and agent $(q+2)$ have no incentive to deviate from this strategy profile because they value only the single good on which they are spending their entire budget. Next, we show that none of the other agents want to deviate from $\mathbf{s}$ either when the values of ϵ, ϵ', a , and a' are set as follows:

$$\begin{aligned} \epsilon &= \frac{1}{(q-k)}, \quad \epsilon' = \epsilon^2, \\ a &= \frac{\epsilon'}{k(\epsilon + \epsilon')^2} \frac{q(1-\epsilon) + k\epsilon}{(q-1)(1-\epsilon) + k\epsilon}, \quad a' = \frac{2(\epsilon + \epsilon')^2}{\epsilon(2 - (q-k)\epsilon')^2}. \end{aligned}$$

Note that all the agents in our market have additive utilities. Let us denote the utility of agent i for a unit amount of good j by $v_{i,j}$. We will first show the following property of Nash equilibrium in such markets that will be useful to prove the theorem.

Claim 1. *Let $\mathbf{p}$ and $\mathbf{x}$ be the prices and allocation corresponding to the strategy profile $\mathbf{s}$; then,*

$$\forall (i, j), \quad x_{i,j} > 0 \Rightarrow \frac{v_{i,j}}{p_j} (1 - x_{i,j}) = \arg \max_k \frac{v_{i,k}}{p_k} (1 - x_{i,k}).$$

Proof. To the contrary, suppose $x_{i,j} > 0$ but for some good k , $\frac{v_{i,j}}{p_j} (1 - x_{i,j}) < \frac{v_{i,k}}{p_k} (1 - x_{i,k})$. We will show this implies agent i can deviate and gain. For a small-enough $\delta > 0$, suppose agent i moves δ money from good j to good k . The difference in agent i 's utility in that case is

$$\begin{aligned} &v_{i,j} \frac{(s_{i,j} - \delta)}{p_j - \delta} + v_{i,k} \frac{s_{i,k} + \delta}{p_k + \delta} - v_{i,j} \frac{s_{i,j}}{p_j} - v_{i,k} \frac{s_{i,k}}{p_k} \\ &= \frac{v_{i,k}}{p_k + \delta} \delta \frac{p_k - s_{i,k}}{p_k} - \frac{v_{i,j}}{p_j - \delta} \delta \frac{p_j - s_{i,j}}{p_j} \\ &= \delta \left(\frac{v_{i,k}}{p_k + \delta} (1 - x_{i,k}) - \frac{v_{i,j}}{p_j - \delta} (1 - x_{i,j}) \right). \end{aligned}$$

Clearly, the expression is strictly greater than zero for any δ such that

$$\delta < \frac{v_{i,k} p_j (1 - x_{i,k}) - v_{i,j} p_k (1 - x_{i,j})}{v_{i,k} (1 - x_{i,k}) + v_{i,j} (1 - x_{i,j})}.$$

Note that the upper bound on δ is strictly positive given our initial assumption, a contradiction. $\square$

At strategy $\mathbf{s}$, each agent $i \in [k+1, q+1]$ is allocating positive amount on every good for which she has some positive utility (i.e., $v_{i,j} > 0$); then, $s_{i,j} > 0$ for every good $j \in [1, (q+1)]$. Therefore, using the claim, it suffices to show that the value of $\frac{v_{i,j}}{p_j}(1 - x_{i,j})$ is the same for all the goods on which she is spending money. Let us denote this quantity by $t_{i,j}$ for ease of notation:

$$\forall(i, j), \quad \text{let } t_{i,j} = \frac{v_{i,j}}{p_j}(1 - x_{i,j}).$$

For agent $(q+1)$, $t_{(q+1),j}$ is the same for all $j \in [k+1, q]$. So, we only need to compare this value with $t_{(q+1),(q+1)}$:

$$\begin{aligned} t_{(q+1),(q+1)} &= \frac{2}{2 - (q-k)\epsilon'} \left(1 - \frac{1 - (q-k)\epsilon'}{2 - (q-k)\epsilon'} \right) \\ &= \frac{2}{(2 - (q-k)\epsilon')^2}. \end{aligned}$$

On the other hand, by construction, for any $j \in [q+1, k]$, we have

$$\begin{aligned} t_{(q+1),j} &= \frac{a'}{(\epsilon + \epsilon')} \left(1 - \frac{\epsilon'}{\epsilon + \epsilon'} \right) = \frac{a'\epsilon}{(\epsilon + \epsilon')^2} = \frac{2}{(2 - (q-k)\epsilon')^2} \\ &= t_{(q+1),(q+1)} \quad \left(\because a' = \frac{2(\epsilon + \epsilon')^2}{\epsilon(2 - (q-k)\epsilon')^2} \right). \end{aligned}$$

For any agent $i \in [k+1, q]$ and good i , $t_{i,i} = \frac{1}{(\epsilon + \epsilon')} \times (1 - \frac{\epsilon}{\epsilon + \epsilon'}) = \frac{\epsilon'}{(\epsilon + \epsilon')^2}$. For good $j \in [1, k]$, we have

$$\begin{aligned} t_{i,j} &= \frac{a*k}{q(1-\epsilon) + k\epsilon} \left(1 - \frac{(1-\epsilon)}{q(1-\epsilon) + k\epsilon} \right) \\ &= \frac{a*k*((q-1)(1-\epsilon) + k\epsilon)}{(q(1-\epsilon) + k\epsilon)^2} = \frac{\epsilon'}{(\epsilon + \epsilon')^2} \\ &= t_{i,i} \quad \left(\because a = \frac{\epsilon'}{k(\epsilon + \epsilon')^2} \frac{(q(1-\epsilon) + k\epsilon)^2}{(q-1)(1-\epsilon) + k\epsilon} \right). \end{aligned}$$

Thus, using Claim 1, it follows that $\mathbf{s}$ is a Nash equilibrium. Next, it remains to show that the ratio of total payoff at $\mathbf{b}$ and $\mathbf{s}$ goes to $e^{1/e}$ as $q \rightarrow \infty$. Let

$$\begin{aligned} k &= q/e \Rightarrow (q-k) = q(e-1)/e, \quad \text{and} \\ \epsilon &= 1/(q-k) = (1/q)(e/(e-1)). \end{aligned}$$

At $\mathbf{s}$, the utility of every agent $i \in [k+1, q]$ is

$$\begin{aligned} u_i &= \frac{\epsilon}{\epsilon + \epsilon'} + ak \frac{(1-\epsilon)}{q(1-\epsilon) + k\epsilon} \\ &= \frac{1}{1+\epsilon} + \frac{\epsilon}{1+\epsilon} \frac{(q-1)(1-\epsilon) + k\epsilon}{q(1-\epsilon) + k\epsilon} \\ &= \frac{1}{1+\epsilon} + \frac{\epsilon}{(1+\epsilon)} \frac{(1-1/q)(1-\epsilon) + e\epsilon}{(1-\epsilon) + e\epsilon}. \end{aligned}$$

Replacing $\epsilon = 1/(q-k)$ in the $\lim_{q \rightarrow \infty} u_i \rightarrow 1$, $\forall i \in [k+1, q]$. The utility of agent $(q+1)$ is

$$\begin{aligned} u_{(q+1)} &= \frac{(q-k)a'\epsilon'}{\epsilon + \epsilon'} + 2 \frac{1 - (q-k)\epsilon'}{2 - (q-k)\epsilon'} \\ &= (q-k) \frac{2(\epsilon + \epsilon')^2}{\epsilon(2 - (q-k)\epsilon')^2} \frac{\epsilon'}{\epsilon + \epsilon'} + 2 \frac{1 - (q-k)\epsilon'}{2 - (q-k)\epsilon'} \\ &= \frac{2(1+\epsilon)}{(q-k)(2-\epsilon)^2} + 2 \frac{1-\epsilon}{2-\epsilon} (\epsilon' = \epsilon^2, \epsilon = 1/(q-k)). \end{aligned}$$

Again, replacing $\epsilon = 1/(q-k)$ here and letting $\lim_{q \rightarrow \infty} \times u_{(q+1)} \rightarrow (0 + (2/2)) = 1$. The utility of agent $(q+2)$ is $\frac{2}{2-(q-k)\epsilon'} = \frac{2}{2-(1/(q-k))}$, and therefore, $\lim_{q \rightarrow \infty} u_{(q+2)} \rightarrow 1$. Finally,

$$\forall i \in [1, k], u_i = \frac{k}{q(1-\epsilon) + k\epsilon} = \frac{k}{q - \epsilon(q-k)} = \frac{k}{q-1}.$$

Using the utility values of all the agents, we get that the Nash social welfare at $\mathbf{s}$ as $q \rightarrow \infty$ tends to $(\frac{k}{q})^{\frac{k}{q}}$, and therefore, PoA is at least $(\frac{q}{k})^{\frac{k}{q}}$. Because $k = q/e$, it is at least $e^{\frac{1}{e}}$, which concludes the proof of Theorem 2. $\square$

4. Price of Anarchy of Trading Post for Leontief Valuations

In this section, we focus on the class of Leontief valuations, which is a well-motivated subclass of the concave valuations studied in the previous section, and we show that we can achieve essentially optimal price of anarchy bounds for these valuations.

As we show in Section 6.1, the games induced by the Trading Post mechanism when the agents have Leontief valuations may not always possess a pure Nash equilibrium. We therefore also study the Trading Post mechanism with entrance fee $\Delta > 0$, which alleviates many of the equilibrium existence issues (see Theorem 10). Also, because the value of Δ can be arbitrarily small, the impact of this entrance fee on the outcome of the game is insignificant. The main result of this section is that, for an arbitrarily small constant $\epsilon \in (0, 1/m)$, the Trading Post mechanism $\mathcal{TP}(\Delta)$ with $\Delta \leq \epsilon/m^2$ has a price of anarchy of at most $1 + \epsilon$.

In order to analyze the performance of the Trading Post mechanism for Leontief valuations, we start by defining a notion of approximate market equilibrium that will be useful. Recall the definition of Fisher market equilibrium from Section 2.

Definition 1 (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 (i) all the goods with a positive price are completely sold, (ii) all the agents exhaust their budget, (iii) each agent gets an ϵ -optimal bundle at prices $\mathbf{p}$; that is, for every bundle

$y \in [0, 1]^m$ that i could afford at these prices ($\mathbf{p} \cdot \mathbf{y} \leq B_i$), we have $u_i(\mathbf{y}) \leq u_i(\mathbf{x}_i)(1 + \epsilon)$. Note that when $\epsilon = 0$, it is a Fisher market equilibrium.

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

Theorem 3. *Let $\epsilon > 0$. Then, for every $0 < \Delta < \min\{\frac{\epsilon}{m^2}, \frac{1}{m}\}$, every pure Nash equilibrium of the mechanism $\mathcal{TP}(\Delta)$ with Leontief valuations corresponds to an ϵ -market equilibrium.*

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

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

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

For each agent 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, $u_i(\tilde{\mathbf{b}}) = \min_{j \in D_i} \phi_{i,j}$. Sort the items in D_i increasingly by their contribution to i 's utility: $\phi_{i,i_1} \leq \phi_{i,i_2} \leq \dots \leq \phi_{i,i_{m_i}}$; it follows that $u_i(\tilde{\mathbf{b}}) = \phi_{i,i_1}$. 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, agent 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_i . Because $\tilde{\mathbf{b}}$ 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 $\mathbf{p} = (p_1, \dots, p_m)$, 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 $\tilde{\mathbf{b}}$ 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 agent exhausts their budget. Moreover, observe that all the prices are strictly positive. We must additionally show that each agent gets an ϵ -optimal bundle at $(\mathbf{p}, \mathbf{x})$.

Fix an arbitrary agent 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 agent i shifts all the money spent on purchasing items outside S_i to purchase instead higher fractions from the items in S_i . Because the strategy profile $\mathbf{b}$ is an exact equilibrium in the game $\mathcal{TP}(\Delta)$, the amount of money spent by agent 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: that is,

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

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 because 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 agent gets an ϵ -optimal bundle, and so, $(\mathbf{p}, \mathbf{x})$ is an ϵ -market equilibrium. $\square$

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 4. *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 $(\mathbf{p}', \mathbf{x}')$ be an ϵ -market equilibrium, and let $(\mathbf{p}^*, \mathbf{x}^*)$ be exact market equilibrium prices and allocation. By abuse of notation,

let $u_i(\mathbf{p}')$ denote the optimal utility agent i can obtain at prices $\mathbf{p}'$ (i.e., $u_i(\mathbf{p}') = \max\{u_i(\mathbf{y}) \mid \mathbf{y} \geq 0; \mathbf{p}' \cdot \mathbf{y} \leq B_i\}$).

For Leontief utility functions, the following Eisenberg and Gale (1959) convex formulation captures the market equilibrium allocation:

$$\begin{aligned} \max \quad & \sum_{i=1}^n B_i \cdot \log u_i \\ \text{such that} \quad & u_i \leq \frac{x_{ij}}{v_{ij}}, \forall i \in N, j \in M \\ & \sum_{i=1}^n x_{ij} \leq 1, \forall j \in M \\ & x_{ij} \geq 0, \forall i \in N, j \in M. \end{aligned} \quad (7)$$

Note that, if $(\mathbf{u}^*, \mathbf{x}^*)$ is the optimal of the formulation, then $u_i^* = u_i(\mathbf{x}_i^*)$ for every agent i . In order to get a utility of one at prices $\mathbf{p}$, agent i would need to buy at least v_{ij} units of good j and thereby, spend a total amount of money equal to $\phi_i(\mathbf{p}) = \sum_j v_{ij} p_j$. Using this, Devanur (2010) derived the dual of this convex program:

$$\begin{aligned} \min \quad & \sum_j p_j - \sum_i B_i \log(\phi(\mathbf{p})) + \sum_i B_i \log(B_i) - \sum_i B_i \\ \text{such that} \quad & \forall j: p_j \geq 0. \end{aligned}$$

Note that the term $(\sum_i B_i \log(B_i) - \sum_i B_i)$ is a constant for a given market because B_i 's are constants and hence, is omitted in Devanur (2010). Because $(\mathbf{p}^*, \mathbf{x}^*)$ is a market equilibrium, using strong duality and the fact that agents spend all their money at equilibrium (i.e., $\sum_j p_j^* = \sum_i B_i$):

$$\begin{aligned} \sum_i B_i \log(u_i(\mathbf{x}_i^*)) &= - \sum_i B_i \log(\phi(\mathbf{p}^*)) \\ &+ \sum_i B_i \log(B_i). \end{aligned} \quad (8)$$

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

$$\begin{aligned} & - \sum_i B_i \log(\phi(\mathbf{p}')) + \sum_i B_i \log(B_i) \\ & \leq - \sum_i B_i \log(\phi(\mathbf{p}')) + \sum_i B_i \log(B_i). \end{aligned}$$

Substituting the left-hand side using Equation (8) and taking an antilogarithm on both sides yield

$$\prod_i u_i(\mathbf{x}_i^*)^{B_i} \leq \prod_i \left(\frac{B_i}{\phi(\mathbf{p}')} \right)^{B_i}. \quad (9)$$

Because agent i has total of B_i budget out of which she needs $\phi(\mathbf{p}')$ to derive one unit of utility, the maximum utility she can derive from the optimal bundle at prices $\mathbf{p}'$ is

$$\forall i: u_i(\mathbf{p}') = \frac{B_i}{\phi(\mathbf{p}')} \quad (10)$$

Because $(\mathbf{x}', \mathbf{p}')$ is an ϵ -market equilibrium, each agent gets an ϵ -optimal bundle, so $u_i(\mathbf{p}') \leq u_i(\mathbf{x}_i')(1 + \epsilon)$. According to (10), this implies $\frac{B_i}{\phi(\mathbf{p}')} \leq u_i(\mathbf{x}_i')(1 + \epsilon)$, which combined with (9) gives

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

Because the Nash social welfare at $\mathbf{x}$ is $(\prod_i u_i(\mathbf{x}_i)^{B_i})^{\frac{1}{B}}$, the result follows. $\square$

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

Theorem 5. For every $\epsilon > 0$, the Trading Post game $\mathcal{TP}(\Delta)$ with entrance fee $0 < \Delta < \min\{\frac{\epsilon^2}{m}, \frac{1}{m}\}$ has a price of anarchy of $1 + \epsilon$, even for arbitrary budgets.

Proof. By Theorem 3, every pure Nash equilibrium of $\mathcal{TP}(\Delta)$ corresponds to an ϵ -market equilibrium. By Theorem 4, 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. $\square$

5. Individual Fairness Guarantees

The price of anarchy bounds of the previous sections provided us with an aggregate measure of (approximate) fairness for the outcomes of the Trading Post mechanism. We now show that the Nash equilibria of Trading Post also satisfy an important notion of individual fairness: proportionality. That is, each agent i gets a fraction of at least B_i/B of its maximum utility (i.e., the utility of the agent for receiving all items in full); for the case of equal budgets, the guarantee is $1/n$. Recall that B_i is the agent's budget, B is the sum of budgets, and m is the number of items.

Theorem 6. For each $\Delta \geq 0$, in each Nash equilibrium of $\mathcal{TP}(\Delta)$ with concave utilities, each agent i gets a value of $\frac{B_i}{B}(1 - \epsilon_i)$ of its maximum utility (taken over all bundles), where $\epsilon_i = \frac{\Delta(m-1)}{B_i}$.

Proof Sketch. The proof can be found in the e-companion and is based on showing that at any Nash equilibrium $\tilde{\mathbf{b}}$, every agent i has a “safe” strategy profile $\mathbf{y}_i$ that guarantees the agent a fraction of its maximum possible utility that is (approximately) proportional to its budget. Because $\tilde{\mathbf{b}}$ is a Nash equilibrium, strategy $\mathbf{y}_i$ must to be dominated by $\tilde{\mathbf{b}}_i$, so the agent is getting its proportionality value in the equilibrium. A key tool for establishing proportionality is a combination of Lemma 1 and a lemma stating that for all concave utilities, and every agent i , the price of every item j that agent i values at $\mathbf{x}$ is strictly positive (see Lemma EC.1 in the e-companion). $\square$

Note that for the standard Trading Post mechanism (i.e., for $\Delta = 0$), this result yields exact proportionality, but it also provides approximate proportionality for $\mathcal{TP}(\Delta)$ with $\Delta > 0$.

6. The Trading Post Mechanism and Equilibrium Existence

The previous sections proved aggregate and individual fairness guarantees satisfied by any pure Nash equilibrium of the games induced by the Trading Post mechanism. In this section, we discuss the conditions under which such equilibria are guaranteed to exist. The missing proofs of this section can be found in the e-companion.

Even for the case of additive valuations, if there do not exist at least two agents with positive value for each good, then a pure Nash equilibrium may not exist. For instance, consider a game with two agents, 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 agents will compete for item 1, whereas agent 2 is the only one that wants item 2. As a result, agent 2 will successively reduce its bid for item 2 in order to raise its bid for item 1 and get a higher fraction of it. However, in the limit of its bid for the second item going to zero, agent 2 loses the item.

To address this issue, Feldman et al. (2009) introduced a mild technical condition known as *perfect competition* and proved the existence of pure Nash equilibria in the Trading Post mechanism for additive valuations under perfect competition. This condition states that for each good j , there exist at least two bidders that demand a strictly positive amount of j when the good is priced at zero. For concave valuation functions that are differentiable, the condition can be generalized as $\frac{\partial u_i}{\partial x_{i,j}}(\mathbf{x}), \frac{\partial u_k}{\partial x_{k,j}}(\mathbf{x}) > 0$ for any allocation $\mathbf{x} = (x_g)_{g \in M}$ where $x_j = 0$. Using this generalization, we extend the result of Feldman et al. (2009) and prove the existence of pure Nash equilibria for the Trading Post mechanism and CES valuations with $\rho \in (-\infty, 1]$.

Theorem 7. *The Trading Post game with no minimum bid has exact pure Nash equilibria for all CES utilities with perfect competition and $\rho \in (-\infty, 1]$.*

Using the same sequence of arguments, we also show equilibrium existence for all concave, continuous, and strictly increasing valuations.

Theorem 8. *The Trading Post game with no minimum bid has exact pure Nash equilibria for all concave, continuous, and strictly increasing utilities.*

Proof. The proof can be found in the e-companion and uses Reny's theorem for discontinuous games. In particular, we show that even though a best response may not exist, the game satisfies a property of “better-reply

security”—at all points in the closure of the graph of the payoff function, there is an agent that can change its bidding strategy to strictly improve compared not only against the current strategy profile but also against any small deviations of the other agents around the current profile. $\square$

6.1. Leontief Valuations

Although the existence proofs capture several important classes of valuations, they do not capture Leontief valuations. Leontief valuations are concave and continuous, but they are not strictly increasing. Also, the adaptation of the perfect competition condition would require that for every item j , there exist at least two distinct agents $i, k \in N$ such that $v_{i,j} > 0$, $v_{k,j} > 0$. However, even after imposing this condition, the existence of pure Nash equilibria in Trading Post is not guaranteed for Leontief valuations as demonstrated by the following example.

Example 1. Consider a game with two agents and two items, where agent 1 has values $v_{1,1} = v_{1,2} = 0.5$ and agent 2 has $v_{2,1} = 0.9$, $v_{2,2} = 0.1$. Assume there is a pure Nash equilibrium profile $\mathbf{b}$. Because both agents require a nonzero 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 agent must receive the two items in the same ratio relative to its valuation; that is,

$$u_i(\mathbf{b}) = \left(\frac{b_i}{b_1 + b_2} \right) \frac{1}{v_{i,1}} = \left(\frac{1 - b_i}{b_1 + b_2} \right) \frac{1}{v_{i,2}}. \quad (11)$$

Otherwise, if the two ratios were not equal, then an agent could transfer weight among the items to improve the smaller fraction. Then, the requirement in (11) is equivalent to the following equations:

$$\left(\frac{b_1}{b_1 + b_2} \right) \frac{1}{0.5} = \left(\frac{1 - b_1}{b_1 + b_2} \right) \frac{1}{0.5} \iff b_1 = b_2 \quad (12)$$

and

$$\left(\frac{b_2}{b_1 + b_2} \right) \frac{1}{0.9} = \left(\frac{1 - b_2}{b_1 + b_2} \right) \frac{1}{0.1} \iff 8b_2^2 + 8b_1b_2 = 9b_1 + 7b_2. \quad (13)$$

Combining Equations (12) and (13), we get that $b_1 = 1$ and $b_2 = 1$, which contradict 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. Our next theorem generalizes this observation by identifying the precise conditions under which the Trading Post mechanism has pure Nash equilibria for Leontief valuations,

depending on whether the corresponding market equilibrium has prices that are strictly positive or not.

Theorem 9. *The Trading Post mechanism with Leontief valuations and perfect competition 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 one.*

This theorem shows a correspondence between the pure Nash equilibria of Trading Post and the corresponding market equilibria with respect to the agents' valuations. We now show that the Trading Post mechanism with entrance fee always has a pure Nash equilibrium for Leontief valuations for any $\Delta > 0$. The proof uses an application of Glicksberg's theorem for continuous games.

Theorem 10. *For instances with Leontief valuations, the Trading Post mechanism with entrance fee, $\mathcal{TP}(\Delta)$, is guaranteed to have a pure Nash equilibrium for every $\Delta > 0$.*

Proof. Let $\mathcal{TP}(\Delta)$ be the Trading Post game with entrance fee Δ . We first show that a variant of the game, $\mathcal{TP}'(\Delta)$, where the strategy space of each agent 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 agent i is a nonempty compact convex subset of a Euclidean space. Moreover, the utility function of each agent is continuous in $\mathbf{x}$ and quasiconcave in the agent's own strategy (e.g., Cambini and Martein 2009, chapter 2). We therefore use the following theorem from Debreu (1952), Fan (1952), and Glicksberg (1952).

Lemma 3 (Debreu 1952, Fan 1952, Glicksberg 1952). *Consider an n -player 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 $\mathbf{s} = (s_1, \dots, s_n)$ and quasiconcave in s_i , then there exists a pure strategy Nash equilibrium.*

The conditions of Lemma 3 apply, and so, $\mathcal{TP}'(\Delta)$ has a pure Nash equilibrium $\tilde{\mathbf{b}}$. Consider now the Trading Post game with minimum bid Δ , $\mathcal{TP}(\Delta)$. Note that the strategy profile $\tilde{\mathbf{b}}$ dominates every other strategy in $\mathcal{TP}(\Delta)$, including those that allow the agents to bid zero on items of interest to them because such strategies can only decrease utility. Thus, $\tilde{\mathbf{b}}$ is also a pure Nash equilibrium in $\mathcal{TP}(\Delta)$, which completes the proof. $\square$

6.2. Beyond Pure Nash Equilibria

Finally, we show an interesting fact regarding the *mixed* Nash equilibria of the Trading Post mechanism.

The set of mixed Nash equilibria of a game is a generalization of pure Nash equilibria where agents can use randomization in choosing their strategy, and unlike pure Nash equilibria, mixed Nash equilibria are guaranteed to exist in every game induced by the Trading Post. Although in most games, the set of mixed Nash equilibria is a strict superset of pure Nash equilibria, we show that for the games induced by the Trading Post mechanism, these sets coincide when no agent is satiated. Details of this result may be found in the e-companion, where we show the following theorem.

Theorem 11. *For instances with additive valuations, every Nash equilibrium of the corresponding Trading Post game where no agent is satiated is pure. Further, the result extends to concave valuations if a mild condition of enough competition is satisfied. For markets with Leontief valuations, the result holds for the Trading Post mechanism with entrance fee, $\mathcal{TP}(\Delta)$, for any $\Delta > 0$.*

One implication of this result is that all the results of the previous sections, which focus on the quality of pure Nash equilibria, directly extend to mixed Nash equilibria where no agent is satiated. Given the low price of anarchy bounds achieved by the Trading Post mechanism, one may wonder if this mechanism in fact has a unique equilibrium. However, we rule this out through constructions for both additive and Leontief valuations provided in the e-companion.

7. Related Work

The paper most closely related to our work is that of Cole et al. (2013), which proposes *truthful* mechanisms for 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 Nash social welfare-maximizing allocation. However, in order to ensure truthfulness, this mechanism is forced to keep some of the goods unallocated, which makes it impractical for many settings. Very recent work by Abebe et al. (2020) showed that this mechanism can be modified to achieve truthfulness without keeping any goods unallocated, but this comes at a cost of a superconstant approximation and a more complicated mechanism. Complementing this work, our paper analyzes a simple and well-studied mechanism that allocates everything.

The Trading Post mechanism has been studied in an extensive body of literature over the years, sometimes under very different names, such as the Shapley–Shubik game (Shapley and Shubik 1977), the Chinese auction (Matros 2007), the proportional

sharing mechanism (see, e.g., Brams and Taylor 1996 and Feldman et al. 2009), and the Tullock contest in rent seeking (Buchanan et al. 1980, Moldovanu and Sela 2001, Fang 2002), the latter being a variant of the game with a different success probability for items that nobody bid on. The Trading Post can also be interpreted as a congestion game (e.g., Georgiou et al. 2006) or an all-pay auction when the budgets are intrinsically valuable to the agents. The dynamics of Trading Post in exchange market settings were studied in Brânzei et al. (2018, 2019).

The Trading Post mechanism is also closely related to the proportional mechanism that received a lot of attention in the networking literature, leading to one of the classic papers on the price of anarchy by Johari and Tsitsiklis (2004). Just like in our setting, this work studies a mechanism that asks the agents to bid on a set of divisible goods and then allocates these goods in proportion to the bids. However, unlike in our setting, where the budget that the agents are using for the bids is artificial, the proportional mechanism of the networking literature uses real monetary bids. This difference has a crucial impact on the equilibrium structure because the agents in the games that we analyze wish to spend all of their budget (as it has no value beyond this setting), whereas the agents facing the proportional mechanism may prefer to keep their money if the prices are too high.

The question of equilibrium existence also arises in the proportional mechanism literature, as the games that it induces may not have pure Nash equilibria because of discontinuities in the payoff functions of the agents. To address this issue, Johari and Tsitsiklis (2004) augment the strategy set of each agent by allowing them to request a nonzero amount of a good provided that no other agent wants that good. Feldman et al. (2009) study the same setting as ours, for additive valuations, and resolve the issue of equilibrium nonexistence by considering only settings where each item is valued by at least two agents. Very recently, Plaut (2019) considers an alternative augmentation of the game to ensure pure Nash equilibrium existence by allowing nonlinear constraints on the bids as an alternative to the entrance fee that we consider in this paper.

Part of the motivation for studying the class of Leontief valuations is the fact that they are considered to be natural valuation abstractions for computing settings where jobs need resources in fixed ratios, exhibiting extreme complementarity. In a very impactful paper, Ghodsi et al. (2011) 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 Leontief valuations. Parkes et al. (2015) assessed DRF in terms of the resulting efficiency, showing that it performs poorly. Dolev et al. (2012) proposed an alternate fairness criterion called Bottleneck-Based Fairness, which was subsequently showed by Gutman and Nisan (2012) to be satisfied by the proportionally fair allocation. Gutman and Nisan (2012) also posed the study of incentives related to this latter notion as an interesting open problem. It is worth noting that Ghodsi et al. (2011) acknowledge that the Fisher market equilibrium (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 this outcome arbitrarily well, thus shedding new light on this setting.

Finally, the NSW objective has played a central role in the literature on fair division of indivisible goods. This work has used a list of novel algorithmic techniques to show that a constant approximation of this objective can be computed in polynomial time (Anari et al. 2017, 2018; Cole et al. 2017; Barman et al. 2018; Cole and Gkatzelis 2018; Garg et al. 2018, 2020) and that its optimal allocation is approximately envy free (Caragiannis et al. 2019b). Also, more recent work has studied the extent to which approximately maximizing the NSW can be combined with approximate envy-freeness properties when some of the items could be donated (Caragiannis et al. 2019a, Chaudhury et al. 2020).

8. Discussion and Future Directions

A very natural alternative to the Trading Post mechanism, studied in this paper, is the Fisher market mechanism, which is a well-studied direct revelation mechanism. In the Fisher market mechanism, the agents are asked to directly report their valuations to the mechanism, and then, the mechanism computes the Fisher market equilibrium based on these valuations. Prior work on this mechanism has revealed that the agents may gain by strategically misreporting their valuations. Adsul et al. (2010) studied the agents' incentives and proved the existence and structural properties of Nash equilibria for this mechanism. Extending this work, Chen et al. (2011, 2012) proved bounds on the extent to which an agent can gain by misreporting for various classes of valuation functions, including additive and Leontief. Brânzei et al. (2014) showed bounds for the price of anarchy of this mechanism with respect to the social welfare objective, and Cole and Tao (2016) studied large markets

under mild randomness and showed that this price of anarchy converges to one. The price of anarchy of the Fisher market mechanism with respect to the Nash social welfare was studied in Brânzei et al. (2017), where it was shown to be at most two for the case of additive valuations but equal to n for Leontief valuations.

The most compelling difference between the Fisher market mechanism and the Trading Post mechanism is for Leontief valuations, where Trading Post approximates the NSW arbitrarily close while still guaranteeing fair outcomes to each individual. A possible justification for the significantly stronger Leontief bounds 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 budget. On the other hand, an agent's unilateral deviation in the Fisher market mechanism can lead to a market equilibrium where the other agents' spending and allocation have 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."

In terms of problems that our paper leaves open, the most obvious one would be to close the gap between the price of anarchy upper and lower bounds for concave valuations. Also, as we show in Section 6.2, our price of anarchy upper bounds also applies to mixed Nash equilibria for instances where no agent is satiated, but it would be interesting to understand the extent to which lower-quality mixed Nash equilibria can arise in the presence of satiated agents. Another open problem would be to provide bounds for the price of stability (PoS) of the Trading Post mechanism (i.e., the ratio of the optimal NSW over the NSW in the best (instead of worst) equilibrium). For Leontief valuations, our PoA upper bound can be arbitrarily close to one, thus resolving this question, but for general concave valuations, we are not aware of any known lower bounds for the PoS.

Acknowledgments

This work was done in part while the authors were research fellows at the Simons Institute for the Theory of Computing. An earlier version of some of the results of this paper appeared at the Association for Computing Machinery Conference in Economics and Computation (Brânzei et al. 2017).

References

- Abebe R, Cole R, Gkatzelis V, Hartline JD (2020) A truthful cardinal mechanism for one-sided matching. Chawla S, ed. *Proc. 2020 ACM-SIAM Sympos. Discrete Algorithms, SODA 2020* (SIAM, Philadelphia), 2096–2113.
- Adsul B, Babu CS, Garg J, Mehta R, Sohoni MA (2010) Nash equilibria in Fisher market. *Algorithmic Game Theory—Proc. Third Internat. Sympos. SAGT 2010*, vol. 6386, Lecture Notes in Computer Science (Springer, Berlin), 30–41.
- Anari N, Gharan SO, Saberi A, Singh M (2017) Nash social welfare, matrix permanent, and stable polynomials. Papadimitriou, CH, ed. *8th Innovations Theoret. Comput. Sci. Conf., ITCS 2017*, vol. 67, LIPIcs (Schloss Dagstuhl-Leibniz-Zentrum für Informatik, Wadern, Germany), 36:1–36:12.
- Anari N, Mai T, Gharan SO, Vazirani VV (2018) Nash social welfare for indivisible items under separable, piecewise-linear concave utilities. Czumaj A, ed. *Proc. Twenty-Ninth Annual ACM-SIAM Sympos. Discrete Algorithms, SODA 2018* (SIAM, Philadelphia), 2274–2290.
- Babaioff M, Lucier B, Nisan N, Leme RP (2014) On the efficiency of the Walrasian mechanism. *ACM Conf. Econom. Comput., EC '14* (ACM, New York), 783–800.
- Barbanel JB (2005) *The Geometry of Efficient Fair Division* (Cambridge University Press, Cambridge, United Kingdom).
- Barman S, Krishnamurthy SK, Vaish R (2018) Finding fair and efficient allocations. Tardos E, Elkind E, Vohra R, eds. *Proc. 2018 ACM Conf. Econom. Comput.* (ACM, New York), 557–574.
- Bertsimas D, Farias VF, Trichakis N (2011) The price of fairness. *Oper. Res.* 59(1):17–31.
- Bertsimas D, Farias VF, Trichakis N (2012) On the efficiency-fairness trade-off. *Management Sci.* 58(12):2234–2250.
- Brams S, Taylor A (1996) *Fair Division: From Cake Cutting to Dispute Resolution* (Cambridge University Press, Cambridge, United Kingdom).
- Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (2016) *Handbook of Computational Social Choice* (Cambridge University Press, Cambridge, United Kingdom).
- Brânzei S, Devanur NR, Rabani Y (2019) Proportional dynamics in exchange economies. Preprint, submitted July 11, <https://arxiv.org/abs/1907.05037>.
- Brânzei S, Gkatzelis V, Mehta R (2017) Nash social welfare approximation for strategic agents. Daskalakis C, Babaioff M, Moulin H, eds. *Proc. 2017 ACM Conf. Economics Comput., EC '17* (ACM, New York), 611–628.
- Brânzei S, Mehta R, Nisan N (2018) Universal growth in production economies. Bengio S, Wallach HM, Larochelle H, Grauman K, Cesa-Bianchi N, Garnett R, eds. *Adv. Neural Inform. Processing Systems 31: Annual Conf. Neural Inform. Processing Systems 2018, NeurIPS 2018*.
- Brânzei S, Chen Y, Deng X, Filos-Ratsikas A, Frederiksen SKS, Zhang J (2014) The Fisher market game: Equilibrium and welfare. *Proc. Twenty-Eighth AAAI Conf. Artificial Intelligence (AAAI Press, Palo Alto, CA)*, 587–593.
- Buchanan J, Tullock G, Tollison R (1980) *Toward a Theory of the Rent-Seeking Society*, Texas A & M University Economics Series (Texas A&M University Press, College Station, TX).
- Cambini A, Martein L (2009) *Generalized Convexity and Optimization: Theory and Applications* (Springer, Berlin).
- Caragiannis I, Gravin N, Huang X (2019a) Envy-freeness up to any item with high Nash welfare: The virtue of donating items. Karlin A, Immorlica N, Johari R, eds. *Proc. 2019 ACM Conf. Econom. Comput., EC 2019* (ACM, New York), 527–545.
- Caragiannis I, Kurokawa D, Moulin H, Procaccia AD, Shah N, Wang J (2019b) The unreasonable fairness of maximum Nash welfare. *ACM Trans. Econom. Comput.* 7(3):12:1–12:32.

- Chaudhury BR, Kavitha T, Mehlhorn K, Sgouritsa A (2020) A little charity guarantees almost envy-freeness. *Proc. 2020 ACM-SIAM Sympos. Discrete Algorithms, SODA 2020* (SIAM, Philadelphia), 2658–2672.
- Chen N, Deng X, Zhang J (2011) How profitable are strategic behaviors in a market? *Proc. 19th Annual Eur. Sympos. Algorithms—ESA 2011*, vol. 6942, Lecture Notes in Computer Science (Springer, Berlin), 106–118.
- Chen N, Deng X, Zhang H, Zhang J (2012) Incentive ratios of Fisher markets. *Proc. 39th Internat. Colloquium Automata, Languages Programming, ICALP 2004: Part II*, vol. 7392:464–475, Lecture Notes in Computer Science (Springer, Berlin).
- Codenotti B, Varadarajan KR (2004) Efficient computation of equilibrium prices for markets with Leontief utilities. *Proc. 31st Internat. Colloquium Automata, Languages Programming, ICALP 2004*, vol. 3142, Lecture Notes in Computer Science (Springer, Berlin), 371–382.
- Cole R, Gkatzelis V (2018) Approximating the Nash social welfare with indivisible items. *SIAM J. Comput.* 47(3):1211–1236.
- Cole R, Tao Y (2016) Large market games with near optimal efficiency. Conitzer V, Bergemann D, Chen Y, eds. *Proc. 2016 ACM Conf. Econom. Comput., EC '16* (ACM, New York), 791–808.
- Cole R, Gkatzelis V, Goel G (2013) Mechanism design for fair division: Allocating divisible items without payments. *Proc. Fourteenth ACM Conf. Electronic Commerce, EC 2013* (ACM, New York), 251–268.
- Cole R, Devanur NR, Gkatzelis V, Jain K, Mai T, Vazirani VV, Yazdanbod S (2017) Convex program duality, Fisher markets, and Nash social welfare. Daskalakis C, Babaioff M, Moulin H, eds. *Proc. 2017 ACM Conf. Econom. Comput., EC '17* (ACM, New York), 459–460.
- Debreu G (1952) A social equilibrium existence theorem. *Proc. Natl. Acad. Sci. USA* 38(10):886–893.
- Devanur NR (2010) Fisher markets and convex programs. Preprint, submitted December, https://www.researchgate.net/publication/265153236_Fisher_Markets_and_Convex_Programs.
- Dolev D, Feitelson DG, Halpern JY, Kupferman R, Linial N (2012) No justified complaints: On fair sharing of multiple resources. *Innovations Theoret. Comput. Sci.* 2012 (ACM, New York), 68–75.
- Eisenberg E, Gale D (1959) Consensus of subjective probabilities: The Pari-Mutuel method. *Ann. Math. Statist.* 30:165–168.
- Fan K (1952) Fixed-point and minimax theorems in locally convex topological linear spaces. *Proc. Natl. Acad. Sci. USA* 38(2): 121–126.
- Fang H (2002) Lottery vs. all-pay auction models of lobbying. *Public Choice* 112(3–4):351–371.
- Feldman M, Lai K, Zhang L (2009) The proportional-share allocation market for computational resources. *IEEE Trans. Parallel Distributed Systems* 20(8):1075–1088.
- Gale D (1960) *Theory of Linear Economic Models* (McGraw Hill, New York).
- Garg J, Hoefer M, Mehlhorn K (2018) Approximating the Nash social welfare with budget-additive valuations. Czumaj A, ed. *Proc. Twenty-Ninth Annual ACM-SIAM Sympos. Discrete Algorithms, SODA 2018* (SIAM, Philadelphia), 2326–2340.
- Garg J, Kulkarni P, Kulkarni R (2020) Approximating Nash social welfare under submodular valuations through (un)matchings. Chawla S, ed. *Proc. 2020 ACM-SIAM Sympos. Discrete Algorithms, SODA 2020* (SIAM, Philadelphia), 2673–2687.
- Georgiou C, Pavlides T, Philippou A (2006) Network uncertainty in selfish routing. *Proc. 20th Internat. Parallel Distributed Processing Sympos. (IPDPS 2006)* (IEEE, New York).
- Ghodsi A, Zaharia M, Hindman B, Konwinski A, Shenker S, Stoica I (2011) Dominant resource fairness: Fair allocation of multiple resource types. *Proc. 8th USENIX Sympos. Networked Systems Design Implementation, NSDI 2011* (USENIX Association, Berkeley, CA), 323–336.
- Glicksberg IL (1952) A further generalization of the Kakutani fixed-point theorem. *Proc. Amer. Math. Soc.* 3:170–174.
- Gutman A, Nisan N (2012) Fair allocation without trade. *Internat. Conf. Autonomous Agents Multiagent Systems, AAMAS 2012* (IFAAMAS, Liverpool, United Kingdom), 719–728.
- Johari R, Tsitsiklis JN (2004) Efficiency loss in a network resource allocation game. *Math. Oper. Res.* 29(3):407–435.
- Kaneko M, Nakamura K (1979) The Nash social welfare function. *Econometrica* 47(2):423–435.
- Kelly FP (1997) Charging and rate control for elastic traffic. *Eur. Trans. Telecommunication* 8:33–37.
- Kelly FP, Maulloo AK, Tan DKH (1998) Rate control in communication networks. *J. Oper. Res. Soc.* 49:237–252.
- Koutsoupias E, Papadimitriou C (2009) Worst-case equilibria. *Comput. Sci. Rev.* 3(2):65–69.
- Matros A (2007) Chines auctions. Unpublished manuscript, University of Pittsburgh, Pittsburgh, Pennsylvania.
- Moldovanu B, Sela A (2001) The optimal allocation of prizes in contests. *Amer. Econom. Rev.* 91(3):542–558.
- Moulin H (2003) *Fair Division and Collective Welfare* (The MIT Press, Cambridge, MA).
- Nash J (1950) The bargaining problem. *Econometrica* 18(2):155–162.
- Nisan N, Roughgarden T, Tardos E, Vazirani V (2007) *Algorithmic Game Theory* (Cambridge University Press, Cambridge, United Kingdom).
- Parkes DC, Procaccia AD, Shah N (2015) Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. *ACM Trans. Econom. Comput.* 3(1):3.1–3.22.
- Plaut B (2019) Optimal Nash equilibria for bandwidth allocation. Preprint, submitted April 6, <https://arxiv.org/abs/1904.03322>.
- Robertson J, Webb W (1998) *Cake-Cutting Algorithms—Be Fair If You Can* (CRC Press, Boca Raton, FL).
- Roughgarden T, Tardos É (2002) How bad is selfish routing? *J. ACM* 49(2):236–259.
- Shapley L, Shubik M (1977) Trade using one commodity as a means of payment. *J. Political Econom.* 85(5):937–968.
- Young H (1995) *Equity* (Princeton University Press, Princeton, NJ).

Simina Brânzei is an assistant professor in computer science at Purdue University. Her research focuses on problems in algorithmic game theory, in particular algorithms, complexity, and dynamics of strategic interactions, as well as theoretical aspects of learning.

Vasilis Gkatzelis is an assistant professor in computer science at Drexel University. His research focuses on problems in algorithmic game theory and approximation algorithms.

Ruta Mehta is an assistant professor in the Department of Computer Science at the University of Illinois at Urbana-Champaign. She works on algorithmic, complexity, strategic, and learning aspects of various game-theoretic and economic problems. Her PhD thesis titled “Nash Equilibrium Computation in Various Games” won the Association for Computing Machinery India Doctoral Dissertation Award, and she is a recipient of the National Science Foundation CAREER Award.