

Walrasian Dynamics in Multi-Unit Markets

Simina Brânzei
Purdue University

Aris Filos-Ratsikas
École Polytechnique Fédérale de Lausanne

Abstract

In a multi-unit market, a seller brings multiple units of a good and tries to sell them to a set of buyers that have monetary endowments. While a Walrasian equilibrium does not always exist in this model, natural relaxations of the concept that retain its desirable fairness properties do exist. We study the dynamics of (Walrasian) envy-free pricing mechanisms in this environment, showing that for *any* such pricing mechanism, the best response dynamic starting from truth-telling converges to a pure Nash equilibrium with small loss in revenue and welfare. Moreover, we generalize these bounds to capture all the (reasonable) Nash equilibria for a large class of (monotone) pricing mechanisms. We also identify a natural mechanism, which selects the minimum Walrasian envy-free price, in which for $n=2$ buyers the best response dynamic converges from any starting profile. We conjecture convergence of the mechanism for any number of buyers and provide simulation results to support our conjecture.

Introduction

The question of allocating valuable resources is one of the central problems faced by human society, starting from the division of land to modern variants such as allocating computational resources to the users of an organization or selling goods in online markets. A crucial development in answer to the question of how to exchange goods (optimally) occurred with the invention of money, i.e., of pricing mechanisms that aggregate information about the supply and demand of the goods in order to facilitate trade. The price mechanism was formalized and studied systematically starting with the 19th century, in the works of Fisher (2000), Walras (1874) and Arrow and Debreu (1954), but its complexity was understood only much more recently (Nisan et al. editors 2007). The basic setting in a free market consists of a set of agents that come to the market with their initial endowments (e.g. money or products for sale) and aim to purchase goods¹ in a way that maximizes their utility subject to the initial budget constraints. Walrasian (market) equilibria are outcomes where demand and supply are perfectly balanced, satisfy desirable fairness and efficiency properties, and are guaranteed to exist under certain conditions (Arrow and Debreu 1954).

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Possibly after selling their initial bundles.

While the market equilibrium concept is a powerful abstraction applicable to any economic environment, other driving forces of the modern digital economy are less well understood using traditional economic theory. In particular, online markets have been growing at an extremely high rate in the recent past, but they pose many conceptual challenging questions such as what guiding principles should be used for designing them, how to model and use information about the behavior of the consumers, and even how to understand the product being sold in the first place (as is arguably the case with the products sold in ad auctions).

In this paper we will be concerned with the high level question of deriving a theory for how consumers behave. A scenario motivating our analysis is an online marketplace where a seller posts a number of goods for sale, while interested buyers compete for obtaining them by submitting bids that signal their interest in the goods. The buyers have time to repeatedly update their bids after seeing the bids submitted by others (i.e., there is a clock that counts the time left for changing the bids), in order to get better allocations for themselves or potentially lower the price. This process continues until a stable state is reached (or the time runs out), and then the seller allocates the goods based on the final bids. We are interested in understanding the outcomes at the end states of this dynamic in multi-unit markets, where the units on sale are identical, and in quantifying the social welfare of the participants and revenue extracted by the seller.

Studying the outcomes of economic environments where players repeatedly interact falls under the umbrella of learning how to play games, and in particular, learning how to participate in auctions. There has been an extensive body of work on understanding dynamics in games and auctions under various behavioral models of agents, such as best-response, multiplicative weight updates, no regret learning (e.g., (Freund and Schapire 1999; Roughgarden, Syrgkanis, and Tardos 2017; Daskalakis and Syrgkanis 2016)). Multi-unit auctions have been advocated as a lens to the entire field of (algorithmic) mechanism design (Nisan 2014) and our scenario is in fact “*an environment where repeated best-response provides reasonable approximation guarantees*”, stated explicitly as a research question by Nisan et al. (2011).

We will focus on studying fair pricing mechanisms as captured by (Walrasian) envy-free pricing, where the seller posts a unique price p per unit and each buyer purchases a

bundle of goods that maximizes its utility given this price. This type of pricing refers to the fact that no buyer should “envy” any bundle that it could afford in the market (Guruswami et al. 2005). Fairness has often not been a consideration in auction design, and in fact for the purpose of maximizing revenue of the seller it is useful to impose higher payments to the buyers that are more interested in the goods. However, there are studies showing that customers are unhappy with such discriminatory prices (see, e.g., (Anderson and Simester 2008)), which has led to a body of literature focused on achieving fair pricing (Guruswami et al. 2005; Feldman et al. 2012; Cohen-Addad et al. 2016; Feldman, Gravin, and Lucier 2015; Anshelevich and Sekar 2017). This situation is not unlike that in learning, where fairness has emerged as an important property in classification tasks (Zafar et al. 2017; Kleinberg, Mullainathan, and Raghavan 2017), with multiple notions of fairness emerging in this context.

Our Contributions

We study the (Walrasian) envy-free pricing problem in one of the most basic scenarios possible, namely linear multi-unit markets with budgets. There is one seller who comes equipped with m units of some good (e.g. chairs), while the buyers bring their budgets. The seller has no value for the items, while the buyers value both their money and the goods. The buyers are strategic and a mechanism for envy-free pricing will have to elicit the valuations from the buyers. The budgets are known. In the economics literature, budgets are viewed as hard information (quantitative), as opposed to preferences, which represent soft information and are more difficult to verify (see, e.g., (Petersen 2004)). Bulow, Levin, and Milgrom (2009) provide an example of such an inference in a real-life actions. Known budgets are also studied in the auctions literature (Dobzinski, Lavi, and Nisan 2012; Fiat et al. 2011; Laffont and Robert 1996).

Our goal is to understand the best response dynamics as well as the entire Nash equilibrium set of such a mechanism, together with the social welfare and revenue attained at the stable states. For the best response dynamic, we assume that at each step of the process a buyer observes the bids of others, then updates its own bid optimally given the current state of the market. Our results are for (Walrasian) envy-free pricing mechanisms and can be summarized as follows.

First, we show that for *any* such mechanism, the best-response dynamic starting from the truth-telling profile converges to a pure Nash equilibrium, with the property that all buyers end up with higher utility and all but possibly one buyer end up with more items than what they received at the starting point of the dynamic (Theorem 1). We also show that given some standard market power metrics, the welfare and revenue loss because of the dynamic is small (and actually goes to 0 as the market becomes fully competitive) (Theorems 3 and 4). Furthermore, we show that while convergence can generally be slow (Theorem 2), there are classes of natural mechanism that convergence to an equilibrium in at most n steps (where n is the number of buyers).

Then, we consider any possible profile as the starting point and we show that there are mechanisms for which the

dynamic sometimes fails to converge (Theorem 5). On the positive side, we show that if it does converge to a reasonable equilibrium, then for many mechanisms of interest, the good welfare and revenue guarantees are retained (Theorems 6 and 7). We also prove that a known mechanism for the problem, ALL-OR-NOTHING (Brânzei et al. 2017), actually converges from any starting point, in the case of two buyers (Theorem 8) and provide experimental evidence supporting our conjecture that it converges for any number of buyers. The omitted proofs are in the full version of the paper ².

Related Work

The notion of the Walrasian equilibrium was formulated and studied systematically as early as the beginning of the 19th century, with the foundational works of Fisher (Brainard and Scarf 2000), Walras (1874), and Arrow and Debreu (1954). Its simple pricing scheme and its superior fairness properties, have led to the employment of the Walrasian equilibrium as an auction mechanism for selling goods (Guruswami et al. 2005; Babaioff et al. 2014). The 2000 SMRA auction for 3G licences to UK providers (coined as “The Biggest Auction Ever” in (Binmore and Klemperer 2002)) as well as the Product-Mix auction (Klemperer 2010) run by the Bank of England are prime real-life examples of the employment of Walrasian pricing mechanisms.

The literature on best-response dynamics on games is rich (Roughgarden 2009; Chien and Sinclair 2007; Awerbuch et al. 2008), with questions raised about convergence and convergence time and the properties of the equilibrium which is reached as the results of the dynamic. A line of work has also considered best-response dynamics in auctions (Dütting and Kesselheim 2017; Nisan et al. 2011); this is relevant to this work as pricing mechanisms in our setting can be seen as simple and fair auction mechanisms for selling indivisible goods. Our work is also related to the literature on understanding dynamics of auctions under behavioral assumptions such as no-regret learning and equilibria of auctions (e.g. for Nash, correlated, and coarse-correlated equilibria) with corresponding price of anarchy bounds, e.g. see (Freund and Schapire 1999; Roughgarden, Syrgkanis, and Tardos 2017; Daskalakis and Syrgkanis 2016).

Contrary to some of the related work on dynamics (e.g. see (Kesselheim 2016)), we do not need to make any assumptions about the order in which the buyers best-respond; our results are independent of this order. A summary on dynamics for deciding the allocation of public goods can be found in (Laffont 1987). Best response dynamics starting from the true profile have been studied before in voting (see, e.g. (Brânzei et al. 2013; Meir et al. 2010)). We emphasize here that while a truthful mechanism with good welfare and revenue guarantees exists for our setting (Brânzei et al. 2017), providing guarantees for very general classes of mechanisms as we do here is quite important, as the mechanisms that might actually be employed in reality might not be truthful. A prominent example of this phenomenon is the Generalized Second Price Auction, which is actually used

²The full version is at <https://arxiv.org/abs/1712.08910>.

in practice in favour of its famous truthful counter-part, the VCG mechanism (see (Nisan et al. editors 2007)) and in fact, understanding such non-truthful mechanisms theoretically has been a focus of the related literature (Caragiannis et al. 2015; de Keijzer et al. 2013).

Finally, the question of understanding what can be achieved with item pricing is particularly important given the recent emphasis on the theme of complexity versus simplicity in mechanism design (Hartline and Roughgarden 2009), which is inspired by the computational perspective. Item pricing is a qualitative notion of simple pricing (as opposed to more complex pricing schemes, such as pricing different bundles arbitrarily), and is frequently used for selling goods in supermarkets.

Model and Preliminaries

A linear *multi-unit market* is composed of a set $N = \{1, \dots, n\}$ of buyers and one seller that brings m (indivisible) units of a good. Each buyer i has a budget³ $B_i > 0$, and a valuation v_i per unit, such that the value of buyer i for k units is $v_i \cdot k$. The seller has no value for the units and its goal is to extract money from the buyers, while the buyers value both the money and the good and their goal is to maximize their utility by purchasing units of the good as long as the transaction is profitable.

The seller will set a price p per unit and any buyer can purchase k units at a price of $k \cdot p$. The outcome of the market will be a pair (\mathbf{x}, p) , where p is the unit price and $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{Z}_+^n$ is an *allocation*, with the property that x_i is the number of units received by buyer i and $\sum_{i=1}^n x_i \leq m$.

The *utility* of buyer i for receiving k units at price p is:

$$u_i(p, k) = \begin{cases} v_i \cdot k - p \cdot k, & \text{if } p \cdot k \leq B_i \\ -\infty, & \text{otherwise} \end{cases}$$

The *demand* of a buyer i at price p is defined as the set of optimal bundles for the buyer given the price and its budget. Since the units are indistinguishable, the demand set will simply contain all the bundle sizes that the buyer would maximally prefer:

$$D_i(p) = \begin{cases} \left\{ \min \left\{ \lfloor \frac{B_i}{p} \rfloor, m \right\} \right\}, & \text{if } v_i > p \\ \left\{ 0, 1, \dots, \min \left\{ \lfloor \frac{B_i}{p} \rfloor, m \right\} \right\}, & \text{if } v_i = p \\ \{0\}, & \text{if } v_i < p. \end{cases}$$

Thus, if the valuation is higher than the price per unit, the buyer demands as many units as it can afford (up to exhausting all the units), while if the valuation is lower, the buyer wants zero units.

Terminology: A buyer i is said to be *hungry* at a price p if $v_i > p$, *semi-hungry* if $v_i = p$, *interested* if $v_i \geq p$, and *uninterested* if $v_i < p$.

³Monetary endowment of intrinsic value to all the buyers and the seller.

Walrasian (Envy-Free) Pricing: An allocation and price (\mathbf{x}, p) represent an *envy-free pricing*⁴ if the price per unit is p and $x_i \in D_i(p)$ for all $i \in N$, i.e. each buyer gets a number of units in its demand set at p . A price p is an *envy-free price* if there exists an allocation \mathbf{x} such that (\mathbf{x}, p) is an envy-free pricing. An allocation \mathbf{x} can be *supported* at price p if there is an envy-free pricing (\mathbf{x}, p) .

While an envy-free pricing can always be obtained (e.g. set $p = 1 + B_1 + \dots + B_n$), it is not always possible to sell *all* the units in an envy-free way as can be seen next.

Example 1 (Non-existence of envy-free clearing prices)

Consider a market with $m = 3$ units and two buyers, Alice and Bob, with valuations $v_{\text{Alice}} = v_{\text{Bob}} = 1.1$ and budgets $B_{\text{Alice}} = B_{\text{Bob}} = 1$. At any price $p > 0.5$, no more than 2 units can be sold in total due to budget constraints. At $p \leq 0.5$, both Alice and Bob are interested and want at least 2 units each, but there are only 3 units in total.

Mechanisms and Objectives: We study envy-free pricing mechanisms and will be interested in the social welfare of the buyers and the revenue of the seller. A *mechanism* A for *envy-free pricing* receives as input a market $\mathcal{M} = (\mathbf{v}, \mathbf{B}, m)$ and outputs an envy-free pricing (\mathbf{x}, p) . The valuations are *private* and will be elicited by A from the agents. The budgets and number of units are known.

The *social welfare* at an envy-free pricing (\mathbf{x}, p) is the total value of the buyers for the goods allocated to them, while the *revenue* is the amount of money received by the seller, respectively: $\mathcal{SW}(\mathbf{x}, p) = \sum_{i=1}^n v_i \cdot x_i$, and $\mathcal{REV}(\mathbf{x}, p) = \sum_{i=1}^n x_i \cdot p$.

Incentives: Buyers are rational agents who strategize when reporting their valuations, in order to gain better allocations at a lower price. The truthful valuation profile will be denoted by $\mathbf{v} = (v_1, \dots, v_n)$. However, each buyer i can report any number $v'_i \in \mathbb{R}_+$ as a value, and this will be its *strategy*. Given a mechanism A and market $\mathcal{M} = (\mathbf{v}, \mathbf{B}, m)$, a strategy profile $\mathbf{s} = (s_1, \dots, s_n)$ is a *pure Nash equilibrium* (or simply an equilibrium) if no buyer i can find an alternative strategy s'_i that would improve its utility under mechanism A , given that the strategies of the other buyers are fixed.

Discrete Domain: The valuations and the budgets come from a discrete domain (an infinite grid): $\mathbb{V} = \{k \cdot \epsilon \mid k \in \mathbb{N}\}$ for some $\epsilon > 0$. An mechanism will be allowed to choose the price from an output grid $\mathbb{W} = \{k \cdot \delta \mid k \in \mathbb{N}\}$, for some $\delta > 0$. We note here that if the input and output domains are continuous, the best-response may not be well-defined.

Dynamics Starting from the Truth

In this section, we consider best-response dynamics and will focus on the truth-telling profile as a natural starting point of the dynamic. Given a market $\mathcal{M} = (\mathbf{v}, \mathbf{B}, m)$, we will denote by $\mathcal{H}_p(\mathbf{v})$, $\mathcal{S}_p(\mathbf{v})$, and $\mathcal{I}_p(\mathbf{v})$ the hungry, semi-hungry, and interested buyers at p , respectively. For a strategy profile

⁴We note that more complex forms of envy-free pricing are possible, such as setting a different price per bundle, but our focus will be on the simplest type of unit pricing.

s, s_{-i} will denote the strategy profile of all other buyers except the i 'th one. Our main theorem is that the best response dynamic always converges to a pure Nash equilibrium when starting from the truth telling profile.

Theorem 1 (Convergence) *Let A be any mechanism. Then the best response dynamic starting from the truth-telling profile converges to a pure Nash equilibrium of A . Compared to the truth-telling outcome, the Nash equilibrium reached has the property that*

- the utility of each buyer is (weakly) higher, and
- the number of units received by each buyer is (weakly) larger, possibly with the exception of the last deviator.

Proof: For any step k in the best response process, let $\mathbf{s}^k = (s_1^k, \dots, s_n^k)$ be the vector of valuations reported by all the buyers at k , and p^k the price output by A in this round. For $k = 0$ we have $\mathbf{s}^0 = \mathbf{v}$ as the initial (true) valuations. We show by induction the following properties:

1. The price strictly decreases throughout the best response process, i.e. $p^{k-1} > p^k$ for every step $k \geq 0$.
2. The buyers that appear hungry at p^k , i.e. $\mathcal{H}_{p^k}(\mathbf{s}^k)$, are also hungry at this price with respect to their true valuations, i.e. $v_i > p^k$ for all $i \in \mathcal{H}_{p^k}(\mathbf{s}^k)$. The buyers that appear semi-hungry at p^k , i.e. $\mathcal{S}_{p^k}(\mathbf{s}^k)$, have not changed their inputs, possibly with the exception of the last deviator, whose true valuation is greater or equal to p^k . The buyers that appear uninterested at p^k are honest, i.e. $N \setminus \mathcal{I}_{p^k}(\mathbf{s}^k)$ is such that $s_i^k = v_i$ for all $i \in \mathcal{I}_{p^k}(\mathbf{s}^k)$.

At $k = 0$ property 1) holds trivially since there is no previous price (we can define $p^{-1} = \infty$), while 2) holds for the truth-telling valuations. Suppose properties (1 – 2) hold for all steps up to $k - 1$. We show they also hold at step k . Let i be the buyer that deviates in step k to some valuation s_i^k . Let $v_j^k = v_j^{k-1}$ for all buyers $j \neq i$ and p^k the new price selected by A on input \mathbf{v}^k . We consider a few cases depending on the deviating buyer.

Case 1: Buyer i appears hungry in round $k - 1$: $s_i^{k-1} > p^{k-1}$. By the induction hypothesis, $i \in \mathcal{H}_{p^{k-1}}(\mathbf{s}^{k-1})$, and so $v_i > p^{k-1}$. Then buyer i receives a maximal allocation in round $k - 1$ given the price. If i increases the price in round k , i.e. $p^k > p^{k-1}$, this can only result in buyer i getting (weakly) fewer units at a higher price, which worsens i 's utility compared to round $k - 1$. This cannot be a best response. Similarly, if i 's deviation resulted in the same price for round k , $p^k = p^{k-1}$, buyer i would receive at most as many units at the same price, which is also not an improvement. Thus it must be the case that $p^k < p^{k-1}$.

Moreover, buyer i 's new valuation cannot be smaller than p^k since that would result in i getting no units in round k , which cannot be better than round $k - 1$ where buyer i was hungry with respect to its true value. Since $v_i > p^{k-1} > p^k$, we get $s_i^k \geq p^k$ and $v_i > p^k$. The set $\mathcal{H}_{p^k}(\mathbf{s}^k)$ of buyers that appear hungry in round k contains, in addition to possibly buyer i , also the following buyers:

- the buyers in $\mathcal{H}_{p^{k-1}}(\mathbf{s}^{k-1}) \setminus \{i\}$, which remain hungry with respect to their true valuations in round k since the price decreased compared to round $k - 1$ while their reports have not changed,
- the buyers with valuations $s_i^{k-1} \in (p^k, p^{k-1})$, which are honest in round $k - 1$ and remain honest in round k , with true values strictly above p^k , since none of them deviated in between.

The set $\mathcal{S}_{p^k}(\mathbf{s}^k)$ of buyers that appear semi-hungry in round k contains only buyers that were honest in round $k - 1$, possibly with the exception of i . Thus all the buyers in $\mathcal{S}_{p^k}(\mathbf{s}^k) \setminus \{i\}$ have their true (and reported) valuations equal to p^k . If buyer i deviates to $s_i^k = p^k$, then by previous arguments, $v_i > p^k$. Finally, the buyers $N \setminus \mathcal{I}_{p^k}(\mathbf{s}^k)$ that appear uninterested in round k were honest in round $k - 1$ and remain so in round k , since the deviator i is not one of them. Thus properties (1-2) hold in round k .

Case 2: Buyer i appears semi-hungry in round $k - 1$: $s_i^{k-1} = p^{k-1}$. By the description of $\mathcal{S}_{p^{k-1}}(\mathbf{s}^{k-1})$, if buyer i is not honest in round $k - 1$, it must be the case that i was the deviator in round $k - 1$. If buyer i manages to strictly increase its utility from $k - 1$ to k , then i 's input in round $k - 1$ was not a best response to round $k - 2$, which is a contradiction. By the induction hypothesis, it follows that buyer i is honest in round $k - 1$. Then i cannot wish to keep the price the same or increase it, since that would go above its true valuation, which is $v_i = s_i^{k-1}$. Thus i can only decrease the price, so $p^k < p^{k-1}$ and $v_i > p^k$. Similarly, for s_i^k to be an improvement, buyer i should appear interested on the new instance, so $s_i^k \geq p^k$.

The set of buyers $\mathcal{H}_{p^k}(\mathbf{s}^k)$, which appear hungry in round k , is a superset of $\mathcal{H}_{p^{k-1}}(\mathbf{s}^{k-1})$ and contains all the buyers with valuations s_i^{k-1} in (p^k, p^{k-1}) , as well as i in case $s_i^k > p^k$. The set of semi-hungry buyers in round k , $\mathcal{S}_{p^k}(\mathbf{s}^k)$, contains only buyers that were honest in round $k - 1$, possibly with the exception of buyer i in case $s_i^k = p^k$. Finally, the uninterested buyers in round k , $N \setminus \mathcal{I}_{p^k}(\mathbf{s}^k)$, form a subset of $N \setminus \mathcal{I}_{p^{k-1}}(\mathbf{s}^{k-1})$, were honest in round $k - 1$, and have not changed their inputs in between. Thus the sets in round k satisfy the required properties and the price decreased compared to round $k - 1$.

Case 3: Buyer i appears uninterested in round $k - 1$: $s_i^{k-1} < p^{k-1}$. Then $i \in N \setminus \mathcal{I}_{p^{k-1}}(\mathbf{s}^{k-1})$, which by the induction hypothesis only contains honest buyers. Then buyer i 's utility can only improve by decreasing the price and appearing hungry or semi-hungry in the next round, so $p^k < p^{k-1}$, $s_i^k \geq p^k$, and $v_i > p^k$. Similarly to the previous cases, $\mathcal{H}_{p^k}(\mathbf{s}^k)$ contains all of $\mathcal{H}_{p^{k+1}}(\mathbf{s}^{k+1})$ (which continue to be hungry in round k with respect to both their true and reported valuations), the buyers with valuations $s_i^{k-1} \in (p^k, p^{k-1})$ (which were honest in round $k - 1$ and have not changed in between), and buyer i in case $s_i^k > p^k$ (recall $v_i > p^k$). The set $\mathcal{S}_{p^k}(\mathbf{s}^k)$ contains buyers which were honest in round $k - 1$ and have the same reports in round k , possibly with the exception of buyer i in case $s_i^k = p^k$. The set $N \setminus \mathcal{I}_{p^k}(\mathbf{s}^k)$

is a subset of $N \setminus \mathcal{I}_{p^{k-1}}(\mathbf{s}^{k-1})$, all of which were honest at $k-1$ and kept the same valuations at k .

In all three cases, we obtain that properties (1 – 2) are maintained in round k , which completes the proof by induction. Then the price strictly decreases in every iteration. Since the price output is chosen from a discrete grid of non-negative entries, the best response process either stops or reaches the smallest grid point available above zero. At this point the buyers cannot decrease the price further, which implies there are no more best responses.

The improvement properties follow from the description of the sets $\mathcal{H}_{p^k}(\mathbf{s}^k)$, $\mathcal{S}_{p^k}(\mathbf{s}^k)$, $N \setminus \mathcal{I}_{p^k}(\mathbf{s}^k)$ in the round k where the best response process stops, which completes the argument. \square

In the full version, we also prove that the loss in the allocation of the last deviator is sometimes inevitable.

In terms of the convergence rate, in the worst case, convergence can be slow: the buyers can force a mechanism to traverse the entire input grid by slowly decreasing their reported valuations with each iteration.

Theorem 2 (Lower bound on convergence time) *There is a mechanism such that for arbitrary (fixed) market parameters $\mathcal{M} = (\mathbf{v}, \mathbf{B}, m)$, the best response dynamic takes $\Omega(1/\epsilon)$ steps to converge, where the distance between consecutive entries in both input and output domain is at most $\epsilon > 0$ ⁵.*

However, many natural mechanisms, such as ones that maximize or approximate welfare and revenue (Brânzei et al. 2017; Feldman et al. 2012)) converge faster, which we capture under the umbrella of a *consistency* property, for which we obtain that each buyer best responds at most once, and so the dynamic converges to a pure Nash equilibrium in at most n steps. Informally, consistency means that the price is dictated by the sets of hungry and semi-hungry buyers and not their relative reports.

Social Welfare and Revenue

Next, we will evaluate the effect of the dynamic on the social welfare and revenue obtained by envy-free pricing mechanisms. We will show that the loss in the objectives is small if the market is sufficiently competitive. To measure the level of competition, we define the *budget share* of a market $\mathcal{M} = (\mathbf{v}, \mathbf{B}, m)$ as $\alpha(\mathcal{M}) = \max_{i \in N} \frac{B_i}{\mathcal{REV}(\mathcal{M})}$, where $\mathcal{REV}(\mathcal{M})$ is the maximum possible revenue for valuation profile \mathbf{v} . Intuitively, the budget share measures the fraction of the total revenue that a single buyer can be responsible for; this number is small when the market is competitive. Similar notions of market competitiveness have been studied before (see, e.g., (Dobzinski, Lavi, and Nisan 2012; Cole and Tao 2016)).

Theorem 3 (Welfare of the dynamic from truth-telling)

Let A be any mechanism. Then the best response dynamic starting from the truth-telling profile converges to a pure

⁵E.g., the entries are at most ϵ -apart when the input and output domain are discretized to have the form $\{k \cdot \epsilon \mid k \in \mathbb{N}\}$.

Nash equilibrium of A , whose loss in welfare compared to the truth-telling profile is at most the maximum budget.

We also measure the revenue loss due to the dynamic. A mechanism A is a β approximation for the revenue objective if for every market $\mathcal{M} = (\mathbf{v}, \mathbf{B}, m)$, the revenue achieved by A on \mathcal{M} is at least β times the maximum revenue achieved by any envy-free pricing pair (\mathbf{x}, p) on \mathcal{M} .

Theorem 4 (Revenue of the dynamic from truth-telling)

Let A be a mechanism that approximates the optimal revenue within a factor of $\beta \in [0, 1]$ for every market. Consider any market $\mathcal{M} = (\mathbf{v}, \mathbf{B}, m)$. Then in every Nash equilibrium of A reached through a best-response dynamic from truth-telling, the revenue is a $(\beta - \alpha)/2$ approximation of the optimal revenue for \mathcal{M} , where α is the budget share of the market.

Proof: Given the market $\mathcal{M} = (\mathbf{v}, \mathbf{B}, m)$, for any strategy profile $\tilde{\mathbf{v}}$ we define the following notation. Let $p_A(\tilde{\mathbf{v}})$ denote the price and $x_i^A(\tilde{\mathbf{v}})$ the allocation to buyer i computed by mechanism A on input $\tilde{\mathcal{M}} = (\tilde{\mathbf{v}}, \mathbf{B}, m)$. Additionally, let $\mathcal{REV}_A(\tilde{\mathbf{v}}, \mathbf{B})$ denote the revenue extracted by A on $\tilde{\mathcal{M}}$. For any envy-free price p of $\tilde{\mathcal{M}}$, let $\mathcal{REV}(\tilde{\mathbf{v}}, \mathbf{B}, p)$ denote the maximum possible revenue at the price p . This is achieved by serving fully all the hungry buyers at p and allocating as many units as possible to the semi-hungry buyers. Finally, let $\mathcal{REV}_0(\tilde{\mathbf{v}}, \mathbf{B}, p)$ denote the least possible revenue at price p , which is attained by allocating zero units to each semi-hungry buyer, and $\mathcal{REV}(\tilde{\mathbf{v}}, \mathbf{B})$ denote the maximum possible revenue from the market $\tilde{\mathcal{M}}$. For ease of notation, we will write $\mathcal{H}_{p_A}(\mathbf{v})$ to denote the set $\mathcal{H}_{p_A(\mathbf{v})}(\mathbf{v})$ of hungry buyers at the price $p_A(\mathbf{v})$ computed by A on input \mathbf{v} : (and similarly for the sets $\mathcal{S}_{p_A(\mathbf{v})}(\mathbf{v})$ and $\mathcal{I}_{p_A(\mathbf{v})}(\mathbf{v})$).

Let A be any mechanism that approximates the optimal revenue within a factor of β . We have two cases:

Case 1. $p_A(\mathbf{s}) = p_A(\mathbf{v})$. A deviating buyer can find a strictly improving deviation if and only if it can strictly lower the price. Since the price reached is equal to the price output on the true valuations, it follows that no deviation has taken place. Then \mathbf{v} is the Nash equilibrium to which the dynamic converges and there is no loss in revenue.

Case 2. $p_A(\mathbf{s}) < p_A(\mathbf{v})$. The proof of Theorem 1 implies that in the Nash equilibrium reached from the truth-telling profile, each buyer receives at least as many units as they did on the true input \mathbf{v} , except possibly the last deviator. Let ℓ be the last deviator in the best response path from \mathbf{v} to \mathbf{s} .

Consider the *reduced market* $\mathcal{M}' = (\mathbf{v}_{-\ell}, \mathbf{B}_{-\ell}, m)$, which is obtained from the market \mathcal{M} by removing buyer ℓ . Let $p_{min}^{-\ell}$ be the minimum envy-free price in \mathcal{M}' . Our goal is to lower bound the revenue attained by A in the Nash equilibrium⁶, for which it will be sufficient to provide a lower bound on $\mathcal{REV}_0(\mathbf{s}, \mathbf{B})$ since $\mathcal{REV}_A(\mathbf{s}, \mathbf{B}, p_A(\mathbf{s})) \geq \mathcal{REV}_0(\mathbf{s}, \mathbf{B}, p_A(\mathbf{s}))$.

⁶Note that the revenue objective is not a function of the real values and therefore it can be measured by simply the reports and the corresponding prices.

First, we claim that $p_A(\mathbf{s}) \geq p_{\min}^{-\ell}$, or equivalently, that $p_A(\mathbf{s})$ is an envy-free price for the market \mathcal{M}' . This follows from the following fact established by Theorem 1:

During the best response process, the price always decreases and at any point in time, the only buyer that appears semi-hungry from the set $\mathcal{I}_{p_A}(\mathbf{v})$ (i.e. the set of interested buyers at \mathbf{v}), is the last deviator.

In our case, the last deviator is buyer ℓ . Then for each buyer $i \in \mathcal{I}_{p_A}(\mathbf{v}) \setminus \{\ell\}$, we have that $s_i > p_A(\mathbf{s})$, while for buyer ℓ we have $s_\ell \geq p_A(\mathbf{s})$. This implies that

$$\mathcal{REV}_0(\mathbf{s}, \mathbf{B}, p_A(\mathbf{s})) = \mathcal{REV}_0(\mathbf{v}_{-\ell}, \mathbf{B}_{-\ell}, p_A(\mathbf{s})).$$

Thus, it suffices to bound the minimum revenue attainable at the price $p_A(\mathbf{s})$, that is, $\mathcal{REV}_0(\mathbf{v}_{-\ell}, \mathbf{B}_{-\ell}, p_A(\mathbf{s}))$. Next, note that since $p_A(\mathbf{s}) \in [p_{\min}^{-\ell}, p_A(\mathbf{v})]$, we get $\mathcal{REV}_0(\mathbf{v}_{-\ell}, \mathbf{B}_{-\ell}, p_A(\mathbf{s})) \geq (1/2)\mathcal{REV}_0(\mathbf{v}_{-\ell}, \mathbf{B}_{-\ell}, p_A(\mathbf{v}))$. For ease of exposition, let $\alpha_i = B_i/p_A(\mathbf{s})$ and $\alpha_i^* = B_i/p_A(\mathbf{v})$, $\forall i \in N$. Denote by L the set of buyers with valuations at least $p_A(\mathbf{v})$ in the market $\mathcal{M}' = (\mathbf{v}_{-\ell}, \mathbf{B}_{-\ell}, m)$ (i.e. buyers j with $v_j \geq p_A(\mathbf{v})$) that can afford at least one unit at $p_A(\mathbf{v})$; note that the set of buyers that get allocated any items at $p_A(\mathbf{s})$ is a superset of L . Additionally, since $p_A(\mathbf{v}) > p_A(\mathbf{s})$, the set L does not contain any buyers that are semi-hungry at $p_A(\mathbf{s})$ on \mathcal{M}' . Moreover, the revenue at $p_A(\mathbf{v})$ is bounded by the revenue attained at the (possibly infeasible) allocation where all the buyers in L get the maximum number of units in their demand. These observations give the next inequalities:

$$\mathcal{REV}_0(\mathbf{v}_{-\ell}, \mathbf{B}_{-\ell}, p_A(\mathbf{s})) \geq \sum_{i \in L} \lfloor \alpha_i \rfloor \cdot p_A(\mathbf{s})$$

$$\mathcal{REV}_0(\mathbf{v}_{-\ell}, \mathbf{B}_{-\ell}, p_A(\mathbf{v})) \leq \sum_{i \in L} \lfloor \alpha_i^* \rfloor \cdot p_A(\mathbf{v}).$$

Then the revenue loss, denoted by

$$r = \frac{\mathcal{REV}_0(\mathbf{v}_{-\ell}, \mathbf{B}_{-\ell}, p_A(\mathbf{s}))}{\mathcal{REV}_0(\mathbf{v}_{-\ell}, \mathbf{B}_{-\ell}, p_A(\mathbf{v}))},$$

can be bounded as follows:

$$\begin{aligned} r &\geq \frac{\sum_{i \in L} \lfloor \alpha_i \rfloor \cdot p_A(\mathbf{s})}{\sum_{i \in L} \lfloor \alpha_i^* \rfloor \cdot p_A(\mathbf{v})} \geq \frac{\sum_{i \in L} \lfloor \alpha_i \rfloor \cdot p_A(\mathbf{s})}{\sum_{i \in L} \alpha_i^* \cdot p_A(\mathbf{v})} \\ &= \frac{\sum_{i \in L} \lfloor \alpha_i \rfloor \cdot p_A(\mathbf{s})}{\sum_{i \in L} B_i} = \frac{\sum_{i \in L} \lfloor \alpha_i \rfloor}{\sum_{i \in L} \alpha_i} \geq \frac{\sum_{i \in L} \lfloor \alpha_i \rfloor}{\sum_{i \in L} 2 \lfloor \alpha_i \rfloor}. \end{aligned}$$

Therefore, the revenue loss is at most $1/2$, where we used the fact that the for any buyer $i \in L$, $\lfloor \alpha_i \rfloor \geq 1$, and so $\alpha_i \leq \lfloor \alpha_i \rfloor + 1 \leq 2 \lfloor \alpha_i \rfloor$, since by construction, all buyers in L can afford at least one unit at $p_A(\mathbf{v})$ and therefore at $p_A(\mathbf{s})$. Note that in case the set L is empty, then Mechanism A extracts zero revenue at the truth-telling profile and the theorem follows trivially. Finally, observe that $\mathcal{REV}_0(\mathbf{v}_{-\ell}, \mathbf{B}_{-\ell}, p_A(\mathbf{v})) \geq \mathcal{REV}_0(\mathbf{v}, \mathbf{B}, p_A(\mathbf{v})) - B_\ell$. This inequality holds because $p_A(\mathbf{v})$ is an envy-free price on input $\mathcal{M}' = (\mathbf{v}_{-\ell}, \mathbf{B}_{-\ell}, m)$ and outputting this price results in a loss of revenue of at most the budget of the removed buyer. Since Mechanism A outputs price $p_A(\mathbf{v})$ on (\mathbf{v}, \mathbf{B}) and since it is a β -approximation mechanism for the revenue objective,

it also holds that $\mathcal{REV}_0(\mathbf{v}, \mathbf{B}, p_A(\mathbf{v})) \geq \mathcal{REV}_A(\mathbf{v}, \mathbf{B}) \geq \beta \cdot \mathcal{REV}_0(\mathbf{v}, \mathbf{B})$.

Tying everything together we have:

$$\begin{aligned} \mathcal{REV}_A(\mathbf{s}, \mathbf{B}) &\geq \mathcal{REV}_0(\mathbf{s}, \mathbf{B}, p_A(\mathbf{s})) \\ &= \mathcal{REV}_0(\mathbf{v}_{-\ell}, \mathbf{B}_{-\ell}, p_A(\mathbf{s})) \\ &\geq \frac{1}{2} \cdot \mathcal{REV}_0(\mathbf{v}_{-\ell}, \mathbf{B}_{-\ell}, p_A(\mathbf{v})) \\ &\geq \frac{\beta}{2} \cdot \mathcal{REV}_0(\mathbf{v}, \mathbf{B}) - \frac{B_\ell}{2} \\ &\geq \frac{1}{2} (\beta - \alpha) \mathcal{REV}_0(\mathbf{v}, \mathbf{B}), \end{aligned}$$

where the last inequality follows from the budget share definition. This completes the proof. \square

Note that as β approaches 1 (that is, A approaches a revenue-optimal mechanism) and α approaches 0 (that is the market becomes fully competitive), the revenue attained in the equilibrium is at least half of the optimum. We can also show the budget share is necessary. If there is one buyer with a budget share of 100%, then the revenue obtained as a result of this dynamic can be arbitrarily worse compared to the truth-telling profile (even for a revenue maximizing mechanism).

Dynamics Starting from Arbitrary Profiles

The natural next question is whether we can get similar convergence and welfare and revenue guarantees when the best-response dynamic initiates from arbitrary profiles. If we are interested in the class of all possible mechanisms, then we prove that this is not always the case.

Theorem 5 (Best-response cycles) *There exist mechanisms for which the best response dynamic can cycle when starting from non-truthful profiles.*

Despite the negative nature of the result above, we will provide several theoretical and experimental guarantees for natural mechanisms. Before discussing the convergence properties, we provide the following general result, which states that, for a large class of mechanisms, if the dynamic does converge to a reasonable equilibrium, then the theoretical guarantees of Theorem 3 and Theorem 4 extend to this case as well. We state the two extension theorems below and first provide the following definition of *monotone* mechanisms.

Monotone Mechanism: A mechanism is monotone if the following two properties hold:

- (Price Monotonicity). For any two valuation profiles \mathbf{v} and \mathbf{v}' such that $v'_i \leq v_i$ for all buyers $i \in N$ (with the other market parameters fixed), it holds that the price on \mathbf{v} is not lower than the price on \mathbf{v}' .
- (Supply Monotonicity). Keeping the valuation of a buyer and the price unchanged, while increasing the supply, results in that buyer receiving (weakly) more units.

Theorem 6 (Welfare in non-overbidding equilibria)

Let A be a monotone mechanism. Then in any pure Nash

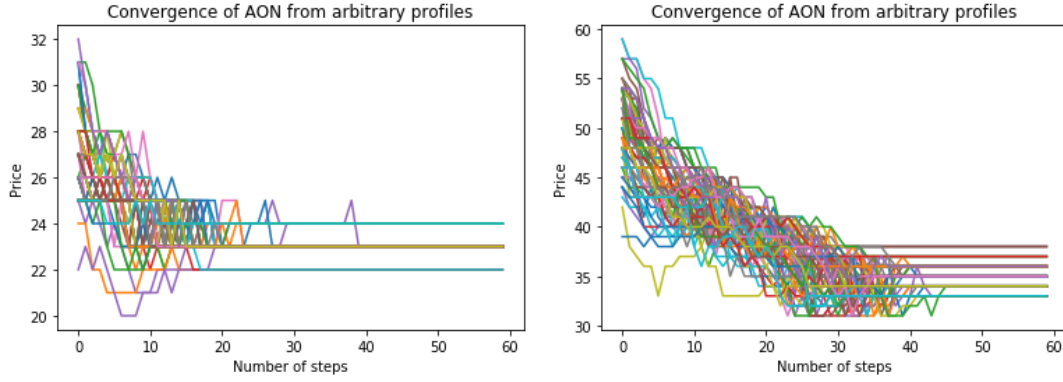


Figure 1: The convergence properties of the ALL-OR-NOTHING mechanism when the dynamic starts from 100 different arbitrary profiles, for $n = 25, m = 20$. The budgets are drawn uniformly at random from $[1, 50]$ (left) and $[1, 125]$ (right). The differently coloured lines indicate different starting profiles $s^{(0)}$.

equilibrium of A where buyers do not overbid, the loss in welfare (compared to the truth-telling outcome of A) is at most $\gamma_A \cdot B^*$, where γ_A is the maximum number of semi-hungry buyers that receive partial allocations by A and B^* the maximum budget.

Theorem 7 (Revenue in non-overbidding equilibria) *Let A be a monotone mechanism that approximates the optimal revenue within a factor of β (with $0 \leq \beta \leq 1$). Then in every non-overbidding Nash equilibrium of A , the revenue is a $(\beta - \gamma_A \cdot \alpha) / 2$ approximation of the optimal revenue for that instance, where α is the budget share of the market and γ_A is the maximum number of buyers that receive partial allocations by A .*

The notion of “partial” allocations refers to buyers that are indifferent between multiple bundles. Here, the buyers that are indifferent among multiple bundles are precisely those whose valuation exactly matches the price, and their demand set contains all the bundle sizes that they can afford at that price (semi-hungry buyers). Note that there are many natural mechanisms with allocation rules that guarantee $\gamma_A = 1$, and for those, the guarantees of Theorem 4 extend verbatim.

“Overbidding” refers to the behavior in which a buyer i , with true value v_i , declares a bid (value) higher than its true value per unit. Any overbidding strategy is weakly dominated by truth-telling and can be ruled out using arguments about uncertainty, risk-aversion or trembling-hand considerations (Lucier and Borodin 2010; Caragiannis et al. 2015). For this reason, overbidding has been ruled out as “unnatural” in the literature (e.g. see (Feldman et al. 2013)) and the study of no-overbidding equilibria is common (see e.g. the incentive properties of the second-price auction) (Christodoulou, Kovács, and Schapira 2008; Caragiannis et al. 2015; Bhawalkar and Roughgarden 2011).⁷

Our final positive theoretical result of the section regards a known-mechanism from the literature, the ALL-OR-NOTHING MECHANISM (Brânzei et al. 2017). Informally,

⁷For a more detailed discussion on the no-overbidding assumption and why it is natural, the reader is referred to (Lucier and Borodin 2010).

the mechanism first computes the minimum envy-free price and then allocates fully to the hungry buyers (as required by envy-freeness) and then considers the semi-hungry buyers in some fixed ordering and it allocates to them either as many units as they can afford or no units at all. We prove that the mechanism converges from any starting profile when we have two buyers.

Theorem 8 *For $n = 2$ buyers, the best response dynamic of the ALL-OR-NOTHING mechanism converges to a Nash equilibrium from any initial strategy profile.*

We remark here that ALL-OR-NOTHING is monotone and actually guarantees $\gamma_A = 1$, therefore by Theorems 6 and 7, it is a mechanism that achieves the guarantees of Theorems 3 and 4 in all reasonable equilibria and which converges from any starting profile, at least when $n = 2$. Whether it always converges for any number of buyers is an open problem.

To this end, in Figure 1, we show experimentally that for randomly generated profiles, the mechanism actually converges for more buyers as well. We present the convergence results for values $n = 25$ and $m = 20$ and for two different choices of budgets, either drawn from $[1, 50]$ or from $[1, 120]$, but the mechanism actually always converges to a pure Nash equilibrium from any initial profile, for any choice of n and m that we have considered. We conclude with the following conjecture.

Conjecture 1 *The best response dynamic of ALL-OR-NOTHING converges to a Nash equilibrium for any initial strategy profile, for any number of buyers.*

Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 740282), from the ISF grant 1435/14 administered by the Israeli Academy of Sciences and Israel-USA Bi-national Science Foundation (BSF) grant 2014389. Aris Filos-Ratsikas was supported by the Swiss National Science Foundation under contract No. 200021_165522 and the ERC Advanced Grant 321171 (ALGAME).

References

- Anderson, E., and Simester, D. 2008. Price stickiness and customer antagonism. *Available at SSRN 1273647*.
- Anshelevich, E., and Sekar, S. 2017. Price doubling and item halving: Robust revenue guarantees for item pricing. In *EC*, 325–342. ACM.
- Arrow, K. J., and Debreu, G. 1954. Existence of an equilibrium for a competitive economy. *Econometrica* 22(3):265–290.
- Awerbuch, B.; Azar, Y.; Epstein, A.; Mirrokni, V. S.; and Skopalik, A. 2008. Fast convergence to nearly optimal solutions in potential games. In *EC*, 264–273. ACM.
- Babaioff, M.; Lucier, B.; Nisan, N.; and Paes Leme, R. 2014. On the efficiency of the walrasian mechanism. In *EC*, 783–800.
- Bhawalkar, K., and Roughgarden, T. 2011. Welfare guarantees for combinatorial auctions with item bidding. In *SODA*.
- Binmore, K., and Klemperer, P. 2002. The biggest auction ever: the sale of the british 3g telecom licences. *Econ. J* 112(478).
- Brainard, W., and Scarf, H. 2000. How to compute equilibrium prices in 1891. *Cowles Foundation Discussion Paper*.
- Brânzei, S.; Caragiannis, I.; Morgenstern, J.; and Procaccia, A. D. 2013. How bad is selfish voting? In *AAAI*, 138–144.
- Brânzei, S.; Filos-Ratsikas, A.; Miltersen, P. B.; and Zeng, Y. 2017. Walrasian pricing in multi-unit auctions. In *MFCS*, 80:1–80:14.
- Bulow, J.; Levin, J.; and Milgrom, P. 2009. Winning play in spectrum auctions. Tech report.
- Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; Kyropoulou, M.; Lucier, B.; Leme, R. P.; and Tardos, E. 2015. Bounding the inefficiency of outcomes in generalized second price auctions. *JET* 156:343–388.
- Chien, S., and Sinclair, A. 2007. Convergence to approximate nash equilibria in congestion games. In *SODA*, 169–178.
- Christodoulou, G.; Kovács, A.; and Schapira, M. 2008. Bayesian combinatorial auctions. In *ICALP*, 820–832.
- Cohen-Addad, V.; Eden, A.; Feldman, M.; and Fiat, A. 2016. The invisible hand of dynamic market pricing. In *EC*, 383–400.
- Cole, R., and Tao, Y. 2016. When does the price of anarchy tend to 1 in large walrasian auctions and fisher markets? In *EC*.
- Daskalakis, C., and Syrgkanis, V. 2016. Learning in auctions: Regret is hard, envy is easy. In *FOCS*, 219–228.
- de Keijzer, B.; Markakis, E.; Schäfer, G.; and Telelis, O. 2013. Inefficiency of standard multi-unit auctions. In *European Symposium on Algorithms*, 385–396. Springer.
- Dobzinski, S.; Lavi, R.; and Nisan, N. 2012. Multi-unit auctions with budget limits. *GEB* 74(2):486–503.
- Dütting, P., and Kesselheim, T. 2017. Best-response dynamics in combinatorial auctions with item bidding. In *SODA*, 521–533.
- Feldman, M.; Fiat, A.; Leonardi, S.; and Sankowski, P. 2012. Revenue maximizing envy-free multi-unit auctions with budgets. In *EC*, 532–549.
- Feldman, M.; Fu, H.; Gravin, N.; and Lucier, B. 2013. Simultaneous auctions are (almost) efficient. In *STOC*, 201–210.
- Feldman, M.; Gravin, N.; and Lucier, B. 2015. Combinatorial auctions via posted prices. In *SODA*, 123–135.
- Fiat, A.; Leonardi, S.; Saia, J.; and Sankowski, P. 2011. Single valued combinatorial auctions with budgets. In *EC*, 223–232.
- Freund, Y., and Schapire, R. E. 1999. Adaptive game playing using multiplicative weights. *GEB* 29(1-2):79–103.
- Guruswami, V.; Hartline, J.; Karlin, A.; Kempe, D.; Kenyon, C.; and McSherry, F. 2005. On profit-maximizing envy-free pricing. In *SODA*, 1164–1173.
- Hartline, J. D., and Roughgarden, T. 2009. Simple versus optimal mechanisms. In *EC*, 225–234.
- Kesselheim, T. 2016. Combinatorial auctions with item bidding: Equilibria and dynamics. In *WINE*.
- Kleinberg, J. M.; Mullainathan, S.; and Raghavan, M. 2017. Inherent trade-offs in the fair determination of risk scores. In *ITCS*, 43:1–43:23.
- Klemperer, P. 2010. The product-mix auction: A new auction design for differentiated goods. *J. Eur. Econ. Assoc* 8.
- Laffont, J., and Robert, J. 1996. Optimal auction with financially constrained buyers. *Econ. Letters* 52(2):181–186.
- Laffont, J. 1987. Incentives and the allocation of public goods. *Handbook of Public Economics* 2:537–569.
- Lucier, B., and Borodin, A. 2010. Price of anarchy for greedy auctions. In *SODA*, 537–553.
- Meir, R.; Polukarov, M.; Rosenschein, J.; and Jennings, N. 2010. Convergence to equilibria in plurality voting. In *AAAI*.
- Nisan, N.; Schapira, M.; Valiant, G.; and Zohar, A. 2011. Best-response auctions. In *EC*, 351–360. ACM.
- Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. (editors) 2007. *Algorithmic Game Theory*. Cambridge Univ. Press.
- Nisan, N. 2014. Algorithmic mechanism design through the lens of multi-unit auctions. *Handbook of Game Theory* 4.
- Petersen, M. 2004. Information: Hard and soft. Working paper.
- Roughgarden, T.; Syrgkanis, V.; and Tardos, E. 2017. The price of anarchy in auctions. *J. Artif. Intell. Res.* 59:59–101.
- Roughgarden, T. 2009. Intrinsic robustness of the price of anarchy. In *STOC*, 513–522. ACM.
- Walras, L. 1874. *Elements d’économie politique pure, ou théorie de la richesse sociale*.
- Zafar, M. B.; Valera, I.; Gomez-Rodriguez, M.; and Gummedi, K. P. 2017. Fairness constraints: Mechanisms for fair classification. In *AISTATS*, 962–970.