



Social Welfare and Profit Maximization from Revealed Preferences

Ziwei Ji^(✉), Ruta Mehta, and Matus Telgarsky

University of Illinois at Urbana-Champaign, Urbana, USA
{ziwei2, rutameht, mjt}@illinois.edu

Abstract. Consider the *seller's problem* of finding optimal prices for her n (divisible) goods when faced with a set of m consumers, given that she can only observe their purchased bundles at posted prices, i.e., *revealed preferences*. We study both social welfare and profit maximization with revealed preferences. Although social welfare maximization is a seemingly non-convex optimization problem in prices, we show that (i) it can be reduced to a dual convex optimization problem in prices, and (ii) the revealed preferences can be interpreted as supergradients of the concave conjugate of valuation, with which subgradients of the dual function can be computed. We thereby obtain a simple subgradient-based algorithm for strongly concave valuations and convex cost, with query complexity $O(m^2/\epsilon^2)$, where ϵ is the additive difference between the social welfare induced by our algorithm and the optimum social welfare. We also study social welfare maximization under the online setting, specifically the random permutation model, where consumers arrive one-by-one in a *random order*. For the case where consumer valuations can be arbitrary continuous functions, we propose a price posting mechanism that achieves an expected social welfare up to an additive factor of $O(\sqrt{mn})$ from the maximum social welfare. Finally, for profit maximization (which may be non-convex in simple cases), we give nearly matching upper and lower bounds on the query complexity for separable valuations and cost (i.e., each good can be treated independently).

1 Introduction

In consumer theory, it is standard to assume that the preferences of a consumer are captured by a *valuation* function, which is often assumed to be known to the mechanism designer. However, in a real market, one can only observe what buyers buy at given prices, the *revealed preferences*. Research on revealed preferences within TCS has two primary objectives: (i) learning valuation functions from revealed preferences, with the goal of having predictive properties [6, 7, 10, 36]; (ii) directly learning the prices that maximize social welfare or profit [4, 5, 8, 9, 11, 13, 28, 29, 35].

The latter problem is of importance to sellers in today's online economies, where a large amount of data about consumers' buying patterns is available. For a seller, profit maximization is the primary goal in general, while she may also

want to maximize social welfare in an effort to earn the goodwill of consumers, with increased market share as a byproduct.

In this paper, we consider social welfare and profit maximization using only revealed preferences.

1.1 Our Model, Results, and Techniques

Consider a market with m consumers and a producer (seller) who produces and sells a set of n *divisible* goods. In the most general case, the preferences of consumer i over bundles of goods are defined by a valuation function $v_i : \mathcal{C} \rightarrow \mathbb{R}_+$ ($\mathcal{C} \subset \mathbb{R}^n$ is called the feasible set), which is her private information and *unknown*. At prices $\mathbf{p}$ she demands bundle $\mathbf{x}_i(\mathbf{p})$ that maximizes her value minus payment, i.e., her *quasilinear utility*

$$\mathbf{x}_i(\mathbf{p}) \in \arg \max_{\mathbf{x} \in \mathcal{C}} (v_i(\mathbf{x}) - \langle \mathbf{p}, \mathbf{x} \rangle).$$

Given prices $\mathbf{p}$, the *revealed preference* refers to the purchased bundle $\mathbf{x}_i(\mathbf{p})$ of each consumer in the market (*demand oracle information*), or even only $\sum_{i=1}^m \mathbf{x}_i(\mathbf{p})$ (*aggregate demand oracle information*). No other information of the valuations is revealed.

Producing the demanded goods incurs cost to the producer, which is represented by a convex cost function c . The producer, or the algorithm, posts prices and makes observations repeatedly, trying to maximize the *social welfare* or *profit*, as described below.

Social Welfare Maximization. The social welfare of bundles $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{C}$ is the sum of consumers' valuations minus production cost, i.e.,

$$\text{SW}(\mathbf{x}_1, \dots, \mathbf{x}_m) = \sum_{i=1}^m v_i(\mathbf{x}_i) - c\left(\sum_{i=1}^m \mathbf{x}_i\right). \quad (1)$$

The benchmark used in this paper is the maximum social welfare SW^* and corresponding maximizing bundles $(\mathbf{x}_1^*, \dots, \mathbf{x}_m^*)$, defined as

$$\text{SW}^* = \max_{\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{C}} \text{SW}(\mathbf{x}_1, \dots, \mathbf{x}_m) \text{ and } (\mathbf{x}_1^*, \dots, \mathbf{x}_m^*) \in \arg \max_{\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{C}} \text{SW}(\mathbf{x}_1, \dots, \mathbf{x}_m). \quad (2)$$

In Sect. 3, offline social welfare maximization is considered. The producer tries to find good prices $\mathbf{p}$ such that $\text{SW}(\mathbf{x}_1(\mathbf{p}), \dots, \mathbf{x}_m(\mathbf{p}))$ is maximized. Although there exist many methods to maximize a concave function, the social welfare is usually a non-concave function in $\mathbf{p}$ [29]. Moreover, the producer only has access to the aggregate demand oracle; the true valuations $v_i(\mathbf{x}_i(\mathbf{p}))$ are unknown.

We first show using duality theory that the maximum social welfare SW^* , which is larger than or equal to any social welfare that can be induced by some prices, can in fact be induced by a single price vector $\mathbf{p}^*$, which is the minimizer of a convex dual function $f(\mathbf{p}) = c^*(\mathbf{p}) - \sum_{i=1}^m v_i^*(\mathbf{p})$, where c^* and v_i^* are respectively *convex* and *concave conjugates* [27], as reviewed in Sect. 2. Moreover, the revealed preferences are supergradients of v_i^* , with which subgradients of

f can be computed. Finally, to get a faster algorithm, we apply a smoothing technique to f and then invoke the accelerated gradient descent method. These ideas are formalized in Algorithm 1, whose guarantee is given below.

Theorem 1 (Informal statement of Theorem 4). *The additive error between the social welfare induced by Algorithm 1 and the maximum social welfare Eq. (2) is at most $O(m/\sqrt{T})$, where T is the number of queries to the aggregate demand oracle.*

In other words, to ensure an additive ϵ approximation of the maximum social welfare, Algorithm 1 needs $O(m^2/\epsilon^2)$ queries to the aggregate demand oracle.

[29] and [28] are the most relevant prior work. [29] studies profit maximization instead of social welfare maximization in a market with one consumer. However, it is assumed that the valuation function is homogeneous, under which profit maximization can be reduced to social welfare maximization. Assumptions made in [29] and this paper are basically identical. Key differences are: (i) [29] proposes a two-level algorithm, where there is an outer iterative algorithm maximizing social welfare, and for each outer iterate, the supergradient of the unknown valuation function is computed by solving a dual optimization problem. In this paper, we only need to solve a single (different) dual optimization problem. Therefore, this gives a simpler approach which may be of independent interest. (ii) The subgradient of the dual objective function in this paper can be interpreted as the excess supply (see the discussion around Eq. (8) in Sect. 3.1), which gives our algorithm a natural interpretation as a Tâtonnement process. (iii) The query complexity given in [29] can be as large as $O(1/\epsilon^6)$ to ensure an additive error of ϵ between the induced social welfare and the maximum social welfare; one reason is that they use subgradient descent, which works for non-smooth convex functions but converges slowly. In this paper, by combining a smoothing technique and accelerated gradient descent, Theorem 1 only needs $O(1/\epsilon^2)$ queries to the aggregate demand oracle. [28] assumes that the valuation is stochastic, but only considers a linear cost. It also considers unit demand consumers with indivisible goods, which is out of the scope of this paper.

Next in Sect. 4 we consider *online* social welfare maximization under the *random permutation* model. In this model, m consumers come to make purchases one by one, and correspondingly the producer is allowed to post prices dynamically, i.e., to update prices from $\mathbf{p}_i$ to $\mathbf{p}_{i+1}$ after the purchase of consumer i . Random permutation here means that those m consumers are first chosen potentially by an adversary, and then come and make purchases one by one in a uniformly random order. (The random permutation model has been extensively studied within online optimization [2, 14, 17], and is more general than the i.i.d. model where each valuation is an independent sample from an unknown distribution.) The objective is to maximize the expected online social welfare $\mathbb{E}[\text{SW}(\mathbf{x}_1(\mathbf{p}_1), \dots, \mathbf{x}_m(\mathbf{p}_m))]$, where the expectation is taken over random orders.

The idea to solve the online social welfare maximization problem is to run an online convex optimization algorithm on a dual problem $f_i(\mathbf{p}) = c^*(\mathbf{p})/m - \langle \mathbf{x}_i(\mathbf{p}_i), \mathbf{p} \rangle$. See Algorithm 2 for details; an introduction to online convex optimization is given in Sect. 4.

Theorem 2 (Informal statement of Theorem 5). *The expected additive error between the online social welfare achieved by Algorithm 2 and the maximum offline social welfare Eq. (2) is bounded by $O(\sqrt{mn})$, where the expectation is taken over random orders of valuations.*

For a given producer, the number of goods n can be thought of as fixed. As a result, the loss of social welfare induced by Algorithm 2 is sublinear in the number of consumers m .

The idea of Algorithm 2 comes from [2], where a general online stochastic convex programming problem is considered. It has many other advantages when applied to online social welfare maximization. First, it is enough to assume that the valuations are continuous; the consumer demand oracle may potentially need to solve some non-convex quasi-linear utility maximization problem, but our focus is on the producer side. Since f_i only depends on the revealed preference $\mathbf{x}_i(\mathbf{p}_i)$, not on v_i , it is still convex. Second, Algorithm 2 is robust, in the sense that it is not sensitive to the potential error in quasi-linear utility maximization. For details, see the discussion at the end of Sect. 4.

Profit Maximization. Next we consider profit maximization with access to the aggregate demand oracle. Given prices $\mathbf{p}$, the profit of producer is the revenue minus production cost, i.e.,

$$\text{Profit}(\mathbf{p}) = \langle \mathbf{p}, \sum_{i=1}^m \mathbf{x}_i(\mathbf{p}) \rangle - c\left(\sum_{i=1}^m \mathbf{x}_i(\mathbf{p})\right). \quad (3)$$

Although it is more reasonable for the producer to maximize the profit, this problem is hard due to non-convexity. The social welfare maximization problems are solved by making a reduction to some convex optimization problem on the space of prices. However, for profit maximization, both the set of optimal bundles and the set of optimal prices may be non-convex, as shown by Example 1 in Sect. 5.

We then consider the case where both the valuations and cost are separable. A separable valuation $v_i(\mathbf{x}) = \sum_{j=1}^n v_{ij}(x_j)$, while similarly a separable cost $c(\mathbf{y}) = \sum_{j=1}^n c_j(y_j)$. Under this assumption, in Sect. 5 we give upper and lower bounds on the query complexity for profit maximization and revenue maximization (i.e., the cost is 0). These upper and lower bounds match for revenue maximization.

Theorem 3 (Informal statement of Theorems 6 and 7). *Consider a market with m consumers and n goods. If the valuations are strongly concave, and both the valuations and cost are separable and Lipschitz continuous, then Algorithm 3 maximizes the profit up to an additive ϵ error with $O(mn/\epsilon)$ queries to the aggregate demand oracle. If the cost is zero, then the strongly concave assumption on valuations can be dropped.*

On the other hand, for concave, separable and Lipschitz continuous valuations, any algorithm requires $\Omega(n/\epsilon)$ queries to the aggregate demand oracle in order to maximize the revenue up to an ϵ additive error.

1.2 Related Work

Samuelson started the theory of *revealed preferences* in 1938 [30] to facilitate mapping observed data to valuation functions, which led to extensive work within economics on “rationalization” or “fitting the samples” [1, 15, 20, 22, 23, 26, 33, 34]. In TCS, there have been a lot of work on learning valuations from revealed preferences with which predictions can be made [6, 7, 10, 36].

Another line of research is on learning prices directly that can maximize social welfare or profit, usually known as the dynamic pricing problem [4, 5, 8, 9, 11, 13]. Some prior works assume nice properties of the demand function (oracle) itself, such as linearity in case of large number of goods [12, 21], concavity [5], Lipschitz continuity [8, 9, 35]. However, these properties may not be satisfied by demands that come from typical concave valuation functions. In [4], the valuation function is assumed to be linear, and is first partially inferred and then used in a price optimization step. However, if the valuation is general concave, such a learning phase is not possible [7].

Recently, [16] studies an online linear classification problem under the revealed preference model.

2 Preliminaries

Market Model. Our model consists of one producer (seller) who produces and sells n divisible goods, and m consumers. Consumer i ’s preferences are represented by an *unknown* valuation function $v_i : \mathcal{C} \rightarrow \mathbb{R}_+$. The feasible consumption set $\mathcal{C} \subset \mathbb{R}_+^n$ is typically assumed to be convex and compact with non-empty interior. It is assumed that $\mathcal{C}$ is known to the algorithm, and let $D = \max_{\mathbf{x}, \mathbf{y} \in \mathcal{C}} \|\mathbf{x} - \mathbf{y}\|_2$ denote the ℓ_2 diameter of $\mathcal{C}$. Note that our algorithm can be extended to the case where v_i ’s have different domains with different diameters; a common domain $\mathcal{C}$ is used here only for convenience.

Given prices $\mathbf{p} = (p_1, \dots, p_n) \in \mathbb{R}_+^n$ of goods, the quasi-linear utility of a bundle $\mathbf{x} \in \mathbb{R}_+^n$ is defined as

$$u_i(\mathbf{x}, \mathbf{p}) = v_i(\mathbf{x}) - \langle \mathbf{x}, \mathbf{p} \rangle.$$

Naturally, consumer i demands a bundle from $\mathcal{C}$ that maximizes her quasi-linear utility

$$\mathbf{x}_i(\mathbf{p}) \in \arg \max_{\mathbf{x} \in \mathcal{C}} u_i(\mathbf{x}, \mathbf{p}),$$

which is known as the *revealed preference* of consumer i at prices $\mathbf{p}$. Once the seller sets prices $\mathbf{p}$, we only get to see $\mathbf{x}_i(\mathbf{p})$ for each consumer i , where every consumer can be thought of as a *demand oracle*, or even only $\mathbf{x}(\mathbf{p}) = \sum_{i=1}^m \mathbf{x}_i(\mathbf{p})$, where the market can be seen as an *aggregate demand oracle*. v_i is always assumed to be continuous to ensure that $\mathbf{x}_i(\mathbf{p})$ exists. This is the only assumption needed for the online social welfare maximization part of this paper; the

offline social welfare maximization part and the profit maximization part further assume that the valuations are strongly concave, which will be introduced later.

The production cost is represented by a convex, Lipschitz continuous, non-decreasing cost function $c : m\mathcal{C} \rightarrow \mathbb{R}_+$, where $m\mathcal{C} = \{m\mathbf{x} | \mathbf{x} \in \mathcal{C}\} = \{\sum_{i=1}^m \mathbf{x}_i | \mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{C}\}$ since $\mathcal{C}$ is convex. Note that the domain of c is big enough to allow production of any aggregate demand. Let λ denote the modulus of Lipschitz continuity of c with respect to the ℓ_2 -norm. It is assumed that the cost function is known to the algorithm.

The producer, or the algorithm, can only post prices and observe the purchased bundles repeatedly, trying to maximize the *social welfare* Eq. (1) or *profit* Eq. (3). Note that if valuations are only continuous, $\mathbf{x}_i(\mathbf{p})$ and the induced social welfare and profit may not be uniquely defined. In this paper, the online social welfare result holds for any $\mathbf{x}_i(\mathbf{p})$, while in offline social welfare and profit maximization, $\mathbf{x}_i(\mathbf{p})$ is unique since strong concavity is assumed.

Convex and Concave Conjugates. The notion of convex and concave conjugates are crucial in our algorithms. Given a convex function $f : \mathcal{D} \rightarrow \mathbb{R}$ where $\mathcal{D} \subset \mathbb{R}^n$ is non-empty, its convex conjugate f^* is defined as:

$$f^*(\mathbf{y}) = \sup_{\mathbf{x} \in \mathcal{D}} (\langle \mathbf{y}, \mathbf{x} \rangle - f(\mathbf{x})),$$

where the domain of f^* is given by $\mathbf{dom} f^* = \{\mathbf{y} \in \mathbb{R}^n | f^*(\mathbf{y}) < \infty\}$. Similarly, given a concave function $f : \mathcal{D} \rightarrow \mathbb{R}$, its concave conjugate is defined as

$$f^*(\mathbf{y}) = \inf_{\mathbf{x} \in \mathcal{D}} (\langle \mathbf{y}, \mathbf{x} \rangle - f(\mathbf{x})),$$

where the domain of f^* is given by $\mathbf{dom} f^* = \{\mathbf{y} \in \mathbb{R}^n | f^*(\mathbf{y}) > -\infty\}$. Since we only compute convex conjugates of convex functions and concave conjugates of concave functions, the above notation is fine. ∂f denotes the set of subgradients of convex f or supergradients of concave f .

Note that in our case, since $\mathbf{dom} c = m\mathcal{C}$ is non-empty and compact and c is continuous, $\mathbf{dom} c^* = \mathbb{R}^n$. Similarly, for every i , $\mathbf{dom} v_i^* = \mathbb{R}^n$.

Lemma 1 is crucial in our algorithm: One key observation in this paper is that revealed preferences are actually supergradients of the concave conjugate of valuation, which is given by Lemma 1. Lemma 1 can be derived from [19] Corollary E.1.4.4 immediately. Although it is stated for convex functions and convex conjugates, corresponding properties hold for concave functions and concave conjugates.

Lemma 1. *Suppose f is convex continuous with non-empty domain. For every pair $(\mathbf{x}, \mathbf{y}) \in \mathbf{dom} f \times \mathbf{dom} f^*$,*

$$\begin{aligned}
& \mathbf{y} \in \partial f(\mathbf{x}) \\
& \iff \mathbf{x} \in \partial f^*(\mathbf{y}) \\
& \iff \mathbf{x} \in \arg \max_{\mathbf{x}' \in \text{dom } f} (\langle \mathbf{y}, \mathbf{x}' \rangle - f(\mathbf{x}')) \\
& \iff \mathbf{y} \in \arg \max_{\mathbf{y}' \in \text{dom } f^*} (\langle \mathbf{x}, \mathbf{y}' \rangle - f^*(\mathbf{y}')) \\
& \iff f(\mathbf{x}) + f^*(\mathbf{y}) = \langle \mathbf{x}, \mathbf{y} \rangle.
\end{aligned}$$

3 Offline Social Welfare Maximization

Problem Description. The goal of offline social welfare maximization is to find prices $\mathbf{p} \in \mathbb{R}_+^n$ such that the induced social welfare $\text{SW}(\mathbf{x}_1(\mathbf{p}), \dots, \mathbf{x}_m(\mathbf{p})) = \sum_{i=1}^m v_i(\mathbf{x}_i(\mathbf{p})) - c(\sum_{i=1}^m \mathbf{x}_i(\mathbf{p}))$ is maximized. As introduced below, v_i 's are assumed to be strongly concave, and thus $\mathbf{x}_i(\mathbf{p})$'s are uniquely determined.

Strongly Concave Valuations. In the offline setting, the valuation functions (v_i 's) are further assumed to be α -strongly concave, meaning that $v_i(x) + \alpha\|x\|_2^2/2$ is concave. Concavity is a standard assumption on valuations to capture diminishing marginal returns. Strong concavity, as its name suggests, is a strong assumption; however it is satisfied by many common valuations such as the constant elasticity of substitution functions and Cobb-Douglas functions (c.f. [29]). Furthermore, a common modulus of strong concavity α is only for convenience; the algorithm can be easily adapted to the case where different v_i have different moduli of strong concavity.

The dual notion to strong concavity (convexity) is strong smoothness. $f : \mathcal{D} \rightarrow \mathbb{R}$ is β -strongly smooth if f is differentiable and its gradient is β -Lipschitz continuous, or formally, for any $\mathbf{x}, \mathbf{y} \in \mathcal{D}$, $\|\nabla f(\mathbf{y}) - \nabla f(\mathbf{x})\|_2 \leq \beta\|\mathbf{y} - \mathbf{x}\|_2$.

The following lemma can be immediately derived from [19] Theorem E.4.2.1 and E.4.2.2.

Lemma 2. *Suppose $f : \mathcal{D} \rightarrow \mathbb{R}$ is concave continuous and $\mathcal{D} \subset \mathbb{R}^n$ is non-empty. Then f is α -strongly concave if and only if f^* is $1/\alpha$ -strongly smooth and concave on $\mathbb{R}^n$.*

Accelerated Gradient Descent. The accelerated gradient descent algorithm, which was first introduced in [25], gives the optimal convergence rate for smooth convex optimization problems. There have been many extensions to AGD, including [3, 32]. In this paper, one variant called the AGM algorithm given in [3] will be invoked.

Lemma 3 ([3] Theorem 4.1). *Suppose $f : \mathcal{D} \rightarrow \mathbb{R}$ is β -strongly smooth and convex and $\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathcal{D}} f(\mathbf{x})$. Given $\mathbf{x}_0 \in \mathcal{D}$, for any $t \geq 1$, The AGM algorithm outputs $\mathbf{x}_t \in \mathcal{D}$ such that*

$$f(\mathbf{x}_t) - f(\mathbf{x}^*) \leq \frac{2\beta\|\mathbf{x}_0 - \mathbf{x}^*\|_2^2}{t^2}.$$

3.1 Algorithm and Analysis

We propose Algorithm 1 to solve the offline social welfare maximization problem.

Algorithm 1. Offline SW Maximization

$\mu \leftarrow \frac{2\lambda}{mD\sqrt{T}}$, $c_\mu(\mathbf{y}) = c(\mathbf{y}) + \frac{\mu}{2}\|\mathbf{y}\|_2^2$, $f_\mu = c_\mu^*(\mathbf{p}) - \sum_{i=1}^m v_i^*(\mathbf{p})$.
 $\mathcal{P} \leftarrow \{\mathbf{p} \geq \mathbf{0} \mid \|\mathbf{p}\|_2 \leq \lambda\}$, $\mathbf{p}_\mu^0 = \mathbf{0}$.
 Give T , the total number of rounds, and $\mathbf{p}_\mu^0$ to the AGM algorithm.
 Output $\mathbf{p}_\mu^T$ returned by the AGM algorithm.

Theorem 4. *The social welfare induced by $\mathbf{p}_\mu^T$ given by Algorithm 1 is within $9\lambda mD/\sqrt{T} + 16\lambda^2 m/\alpha T$ from the maximum offline social welfare.*

The proof is based on two observations. The first one, formalized in Lemma 4, says that the optimal solution of a dual optimization problem can induce the maximum social welfare SW^* .

Lemma 4. *Given concave continuous valuations $v_1, \dots, v_m : \mathcal{C} \rightarrow \mathbb{R}$ and a convex continuous cost $c : m\mathcal{C} \rightarrow \mathbb{R}$,*

$$\text{SW}^* = \min_{\mathbf{p} \in \mathbb{R}^n} \left(c^*(\mathbf{p}) - \sum_{i=1}^m v_i^*(\mathbf{p}) \right),$$

and for any dual optimal solution $\mathbf{p}^$, $(\mathbf{x}_1(\mathbf{p}^*), \dots, \mathbf{x}_m(\mathbf{p}^*))$ maximizes social welfare. If furthermore the cost is non-decreasing and λ -Lipschitz continuous, then*

$$\text{SW}^* = \min_{\mathbf{p} \in \mathbb{R}^n} \left(c^*(\mathbf{p}) - \sum_{i=1}^m v_i^*(\mathbf{p}) \right) = \min_{\mathbf{p} \geq \mathbf{0}, \|\mathbf{p}\|_2 \leq \lambda} \left(c^*(\mathbf{p}) - \sum_{i=1}^m v_i^*(\mathbf{p}) \right),$$

and for any optimal solution $\mathbf{p}^$ of the rightmost dual problem, $(\mathbf{x}_1(\mathbf{p}^*), \dots, \mathbf{x}_m(\mathbf{p}^*))$ maximizes social welfare.*

Proof. Maximizing social welfare is equivalent to solving the following problem

$$\begin{aligned} \max_{\substack{\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{C} \\ \mathbf{y} \in m\mathcal{C}}} \quad & \sum_{i=1}^m v_i(\mathbf{x}_i) - c(\mathbf{y}) \\ \text{s.t.} \quad & \sum_{i=1}^m \mathbf{x}_i = \mathbf{y}. \end{aligned} \tag{4}$$

The Lagrangian is $L(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}, \mathbf{p}) = \sum_{i=1}^m v_i(\mathbf{x}_i) - c(\mathbf{y}) + \langle \mathbf{p}, \mathbf{y} - \sum_{i=1}^m \mathbf{x}_i \rangle$. Equation (4) equals

$$\begin{aligned} & \max_{\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{C}, \mathbf{y} \in m\mathcal{C}} \min_{\mathbf{p} \in \mathbb{R}^n} L(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}, \mathbf{p}) \\ &= \min_{\mathbf{p} \in \mathbb{R}^n} \max_{\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{C}, \mathbf{y} \in m\mathcal{C}} L(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}, \mathbf{p}) \\ &= \min_{\mathbf{p} \in \mathbb{R}^n} \left(c^*(\mathbf{p}) - \sum_{i=1}^m v_i^*(\mathbf{p}) \right). \end{aligned} \tag{5}$$

Here the second line is due to Slater's condition, and the third line is due to the definition of convex conjugate and concave conjugate.

Due to the continuity of v_i 's and c and the compactness of $\mathcal{C}$, the optimal primal solution $(\mathbf{x}_1^*, \dots, \mathbf{x}_m^*, \mathbf{y}^*)$ exists. By Slater's condition, the dual optimal solution $\mathbf{p}^*$ also exists. By the minimax property, we know that $\mathbf{y}^* = \sum_{i=1}^m \mathbf{x}_i^*$, and $\mathbf{x}_i^*$ maximizes $v_i(\mathbf{x}) - \langle \mathbf{x}, \mathbf{p}^* \rangle$, $1 \leq i \leq m$.

For the second part, due to the monotonicity of c , maximizing social welfare is equivalent to

$$\begin{aligned} & \max_{\substack{\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{C} \\ \mathbf{y} \in m\mathcal{C}}} \sum_{i=1}^m v_i(\mathbf{x}_i) - c(\mathbf{y}) \\ & \text{s.t. } \sum_{i=1}^m \mathbf{x}_i \leq \mathbf{y}. \end{aligned} \quad (6)$$

The Lagrangian is still $L(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}, \mathbf{p}) = \sum_{i=1}^m v_i(\mathbf{x}_i) - c(\mathbf{y}) + \langle \mathbf{p}, \mathbf{y} - \sum_{i=1}^m \mathbf{x}_i \rangle$. Equation (6) equals

$$\begin{aligned} & \max_{\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{C}, \mathbf{y} \in m\mathcal{C}} \min_{\mathbf{p} \geq \mathbf{0}, \|\mathbf{p}\|_2 \leq \lambda} L(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}, \mathbf{p}) \\ &= \min_{\mathbf{p} \geq \mathbf{0}, \|\mathbf{p}\|_2 \leq \lambda} \max_{\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{C}, \mathbf{y} \in m\mathcal{C}} L(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}, \mathbf{p}) \\ &= \min_{\mathbf{p} \geq \mathbf{0}, \|\mathbf{p}\|_2 \leq \lambda} \left(c^*(\mathbf{p}) - \sum_{i=1}^m v_i^*(\mathbf{p}) \right). \end{aligned} \quad (7)$$

Here the first line is due to the Lipschitz continuity of c , the second line is due to Sion's theorem [31]. $\square$

Lemma 4 tells us that the minimizer $\mathbf{p}^*$ of $f(\mathbf{p}) = c^*(\mathbf{p}) - \sum_{i=1}^m v_i^*(\mathbf{p})$ induces SW^* , and thus it is natural to try to solve this dual optimization problem. However, v_i 's are unknown to us, and so are v_i^* 's and f . The second observation is that the revealed preference, $\mathbf{x}_i(\mathbf{p})$, actually gives a supergradient of v_i^* at $\mathbf{p}$. Formally, given a concave continuous valuation $v : \mathcal{C} \rightarrow \mathbb{R}$, for any $\mathbf{x} \in \mathcal{C}$, by Lemma 1,

$$\mathbf{x} \in \arg \max_{\mathbf{x}' \in \mathcal{C}} (v(\mathbf{x}') - \langle \mathbf{x}', \mathbf{p} \rangle) \iff \mathbf{x} \in \arg \min_{\mathbf{x}' \in \mathcal{C}} (\langle \mathbf{x}', \mathbf{p} \rangle - v(\mathbf{x}')) \iff \mathbf{x} \in \partial v^*(\mathbf{p}). \quad (8)$$

Similarly, given a convex continuous cost $c : m\mathcal{C} \rightarrow \mathbb{R}_+$, for any $\mathbf{y} \in \mathbb{R}_+^n$, $\mathbf{y} \in \partial c^*(\mathbf{p}) \iff \mathbf{y} \in \arg \max_{\mathbf{y}'} (\langle \mathbf{y}', \mathbf{p} \rangle - c(\mathbf{y}'))$. In other words, the subgradient of c^* at $\mathbf{p}$ gives a bundle which maximizes the producer's profit, assuming everything produced can be sold.

As a result, we can run subgradient-based optimization algorithms to minimize f . (c is known to the algorithm, and so is c^* ; the computation of subgradients of c^* is another problem, but does not require access to the consumer demand oracles.) Since v_i is α -strongly concave, by Lemma 2, v_i^* is $1/\alpha$ -strongly smooth and concave. However, there is no guarantee on c^* , and thus in general, f is not strongly smooth. In this case the standard optimization algorithm is subgradient descent. However, for strongly smooth and convex functions, accelerated gradient descent converges much faster than subgradient descent. To invoke

accelerated gradient descent, the *smoothing technique* given in [24] is used. We minimize f_μ as given in Algorithm 1 and tune the parameter μ , which finally gives Theorem 4. The detailed proof is given in the full version of this paper at <https://arxiv.org/abs/1711.02211>.

4 Online Social Welfare Maximization

Problem Description. In online social welfare maximization, m consumers come one by one and the producer/algorithm can post prices dynamically. Specifically, at step i , prices $\mathbf{p}_i$ are posted, and then consumer i comes and makes a purchase $\mathbf{x}_i(\mathbf{p}_i)$. Then the algorithm updates $\mathbf{p}_i$ to $\mathbf{p}_{i+1}$, based on past information. The goal is to maximize the online social welfare $\text{SW}(\mathbf{x}_1(\mathbf{p}_1), \dots, \mathbf{x}_m(\mathbf{p}_m)) = \sum_{i=1}^m v_i(\mathbf{x}_i(\mathbf{p}_i)) - c(\sum_{i=1}^m \mathbf{x}_i(\mathbf{p}_i))$.

To model the randomness in the real world, it is usually assumed that valuations are sampled i.i.d. from some unknown distribution. Here we consider a slightly stronger model, called the random permutation model. In the random permutation model, an adversary chooses m valuations $\tilde{v}_1, \dots, \tilde{v}_m$ in advance, which then come in a uniformly random order. Formally, let $\gamma = (\gamma_1, \dots, \gamma_m)$ be a random permutation of $(1, \dots, m)$, then at step i the consumer with valuation $v_i = \tilde{v}_{\gamma_i}$ comes and makes a purchase, after $\mathbf{p}_i$ is posted. Note that in the random permutation model, the corresponding offline problem is fixed (with valuations $\tilde{v}_1, \dots, \tilde{v}_m$). We still let $\text{SW}^* = \max_{\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{C}} (\sum_{i=1}^m \tilde{v}_i(\mathbf{x}_i) - c(\sum_{i=1}^m \mathbf{x}_i))$, and let

$$(\tilde{\mathbf{x}}_1^*, \dots, \tilde{\mathbf{x}}_m^*) = \arg \max_{\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{C}} \left(\sum_{i=1}^m \tilde{v}_i(\mathbf{x}_i) - c \left(\sum_{i=1}^m \mathbf{x}_i \right) \right).$$

Our goal is to show that the expected online social welfare $\mathbb{E}_\gamma[\text{SW}(\mathbf{x}_1(\mathbf{p}_1), \dots, \mathbf{x}_m(\mathbf{p}_m))]$ is close to SW^* , where the expectation is taken over the random permutation γ .

Note that no more assumption is made; valuations are only required to be continuous.

Online Convex Optimization. The algorithm for online social welfare maximization invokes an online convex optimization (OCO) algorithm as a subroutine. In an OCO problem, there is a feasible domain $\mathcal{D}$ and T steps. At step t , the OCO algorithm determines $\mathbf{x}_t \in \mathcal{D}$, and then a convex function $f_t : \mathcal{D} \rightarrow \mathbb{R}$ is chosen (potentially by an adversary) and a loss of $f_t(\mathbf{x}_t)$ is induced. Based on the past information (formally, $\mathbf{x}_1, \dots, \mathbf{x}_t$ and $f_1, \dots, f_t$), the algorithm updates $\mathbf{x}_t$ to $\mathbf{x}_{t+1}$, and tries to minimize the *regret*

$$R(T) = \sum_{t=1}^T f_t(\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{D}} \sum_{t=1}^T f_t(\mathbf{x}).$$

The regret of an OCO algorithm $\mathcal{A}$ is denoted by $R_{\mathcal{A}}(T)$.

The *online (sub)gradient descent* algorithm performs the following update at step t :

$$\mathbf{x}_{t+1} = \Pi_{\mathcal{D}}[\mathbf{x}_t - \eta_t g_t(\mathbf{x}_t)],$$

where η_t is the step size, $g_t(\mathbf{x}_t) \in \partial f_t(\mathbf{x}_t)$, and $\Pi_{\mathcal{D}}$ is the ℓ_2 projection onto $\mathcal{D}$.

Lemma 5 ([18] Theorem 3.1). *Let $D = \max\{\|\mathbf{x}_1 - \mathbf{x}_2\|_2 \mid \mathbf{x}_1, \mathbf{x}_2 \in \mathcal{D}\}$, $G = \max\{\|\partial f_t(\mathbf{x})\|_2 \mid 1 \leq t \leq T, \mathbf{x} \in \mathcal{D}\}$, and $\eta_t = D/G\sqrt{T}$. Then*

$$R_{\text{OGD}}(T) = \sum_{t=1}^T f_t(\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{D}} \sum_{t=1}^T f_t(\mathbf{x}) \leq DG\sqrt{T}.$$

4.1 Algorithm and Analysis

Algorithm 2 is proposed to solve the online social welfare maximization problem. The idea of Algorithm 2 comes from [2].

Algorithm 2. Online SW Maximization

$\mathcal{P} \leftarrow \{\mathbf{p} \geq \mathbf{0} \mid \|\mathbf{p}\|_2 \leq \lambda\}$.

Give $\mathcal{P}$ to an OCO algorithm $\mathcal{A}$, and let $\mathbf{p}_1 \in \mathcal{P}$ be the initial prices chosen by $\mathcal{A}$.

for $i = 1$ **to** m **do**

 Post prices $\mathbf{p}_i$.

 Observe $\mathbf{x}_i(\mathbf{p}_i)$, the choice of the buyer who shows up in the i -th step.

 Give $f_i(\mathbf{p}) = \frac{1}{m}c^*(\mathbf{p}) - \langle \mathbf{x}_i(\mathbf{p}_i), \mathbf{p} \rangle$ with domain $\mathcal{P}$ to $\mathcal{A}$, and observe an updated $\mathbf{p}_{i+1}$ from $\mathcal{A}$.

end for

Theorem 5. *The expected social welfare of Algorithm 2 with respect to a uniformly random permutation of continuous valuations, is within $R_{\mathcal{A}}(m) + 2\lambda D_{\infty}\sqrt{nm}$ from the offline optimum social welfare, where $D_{\infty} = \max_{\mathbf{x}, \mathbf{y} \in \mathcal{C}} \|\mathbf{x} - \mathbf{y}\|_{\infty}$ is the ℓ_{∞} diameter of $\mathcal{C}$. Specifically, for online gradient descent, the difference is bounded by $4\lambda D_{\infty}\sqrt{nm}$.*

Proof. For convenience, let $\mathbf{x}_i$ denote $\mathbf{x}_i(\mathbf{p}_i)$. By the regret bound of the OCO algorithm,

$$\begin{aligned} & \sum_{i=1}^m f_i(\mathbf{p}_i) - \min_{\mathbf{p} \geq \mathbf{0}, \|\mathbf{p}\|_2 \leq \lambda} \sum_{i=1}^m f_i(\mathbf{p}) \\ &= \sum_{i=1}^m (f_i(\mathbf{p}_i) + v_i(\mathbf{x}_i)) - \min_{\mathbf{p} \geq \mathbf{0}, \|\mathbf{p}\|_2 \leq \lambda} \sum_{i=1}^m (f_i(\mathbf{p}) + v_i(\mathbf{x}_i)) \\ &\leq R_{\mathcal{A}}(m). \end{aligned} \tag{9}$$

First, we examine the second term of Eq. (9):

$$\begin{aligned} \min_{\mathbf{p} \geq \mathbf{0}, \|\mathbf{p}\|_2 \leq \lambda} \sum_{i=1}^m (f_i(\mathbf{p}) + v_i(\mathbf{x}_i)) &= \min_{\mathbf{p} \geq \mathbf{0}, \|\mathbf{p}\|_2 \leq \lambda} \left(c^*(\mathbf{p}) - \left\langle \sum_{i=1}^m \mathbf{x}_i, \mathbf{p} \right\rangle \right) + \sum_{i=1}^m v_i(\mathbf{x}_i) \\ &= \sum_{i=1}^m v_i(\mathbf{x}_i) - c \left(\sum_{i=1}^m \mathbf{x}_i \right). \end{aligned}$$

The first equality comes from the definition of f , while the second inequality is due to the definition of c^* and the monotonicity and Lipschitz continuity of c . Thus the second term of Eq. (9) always equals the social welfare achieved by Algorithm 2. In the following we show that the first term of Eq. (9) is within $O(\sqrt{mn})$ from the offline maximum social welfare SW^* .

For a permutation $(\gamma_1, \dots, \gamma_m)$ of $1, \dots, m$, let Γ_i denote $(\gamma_1, \dots, \gamma_i)$. Note that $\mathbf{p}_i$ is determined by Γ_{i-1} ($\Gamma_0 = \emptyset$), v_i is determined by γ_i , and $\mathbf{x}_i$ depends on $\mathbf{p}_i$ and γ_i . Fix $1 \leq i \leq m$ and Γ_{i-1} , note that the revealed preference $\mathbf{x}_i$ maximizes the quasi-linear utility given $\mathbf{p}_i$, we have

$$\begin{aligned} \mathbb{E}_{\gamma_i} [f_i(\mathbf{p}_i) + v_i(\mathbf{x}_i) | \Gamma_{i-1}] &= \mathbb{E}_{\gamma_i} \left[\frac{1}{m} c^*(\mathbf{p}_i) - \langle \mathbf{x}_i, \mathbf{p}_i \rangle + v_i(\mathbf{x}_i) \middle| \Gamma_{i-1} \right] \\ &= \frac{1}{m} c^*(\mathbf{p}_i) + \mathbb{E}_{\gamma_i} [-\langle \mathbf{x}_i, \mathbf{p}_i \rangle + \tilde{v}_{\gamma_i}(\mathbf{x}_i) | \Gamma_{i-1}] \\ &\geq \frac{1}{m} c^*(\mathbf{p}_i) + \mathbb{E}_{\gamma_i} [-\langle \tilde{\mathbf{x}}_{\gamma_i}^*, \mathbf{p}_i \rangle + \tilde{v}_{\gamma_i}(\tilde{\mathbf{x}}_{\gamma_i}^*) | \Gamma_{i-1}] \\ &= \frac{1}{m} c^*(\mathbf{p}_i) - \langle \mathbb{E}_{\gamma_i} [\tilde{\mathbf{x}}_{\gamma_i}^* | \Gamma_{i-1}], \mathbf{p}_i \rangle + \mathbb{E}_{\gamma_i} [\tilde{v}_{\gamma_i}(\tilde{\mathbf{x}}_{\gamma_i}^*) | \Gamma_{i-1}]. \end{aligned} \quad (10)$$

Consider the last term of Eq. (10) and take expectation with respect to Γ_{i-1} :

$$\mathbb{E}_{\Gamma_{i-1}} \left[\mathbb{E}_{\gamma_i} [\tilde{v}_{\gamma_i}(\tilde{\mathbf{x}}_{\gamma_i}^*) | \Gamma_{i-1}] \right] = \frac{1}{m} \sum_{i=1}^m \tilde{v}_i(\tilde{\mathbf{x}}_i^*). \quad (11)$$

Then consider the first two terms of Eq. (10):

$$\begin{aligned} &\frac{1}{m} c^*(\mathbf{p}_i) - \langle \mathbb{E}_{\gamma_i} [\tilde{\mathbf{x}}_{\gamma_i}^* | \Gamma_{i-1}], \mathbf{p}_i \rangle \\ &= \frac{1}{m} c^*(\mathbf{p}_i) - \langle \mathbb{E}_{\gamma_i} [\tilde{\mathbf{x}}_{\gamma_i}^*], \mathbf{p}_i \rangle + \langle \mathbb{E}_{\gamma_i} [\tilde{\mathbf{x}}_{\gamma_i}^*], \mathbf{p}_i \rangle - \langle \mathbb{E}_{\gamma_i} [\tilde{\mathbf{x}}_{\gamma_i}^* | \Gamma_{i-1}], \mathbf{p}_i \rangle \\ &\geq \frac{1}{m} c^*(\mathbf{p}_i) - \langle \mathbb{E}_{\gamma_i} [\tilde{\mathbf{x}}_{\gamma_i}^*], \mathbf{p}_i \rangle - \lambda \left\| \mathbb{E}_{\gamma_i} [\tilde{\mathbf{x}}_{\gamma_i}^*] - \mathbb{E}_{\gamma_i} [\tilde{\mathbf{x}}_{\gamma_i}^* | \Gamma_{i-1}] \right\|_2 \\ &= \frac{1}{m} c^*(\mathbf{p}_i) - \frac{1}{m} \left\langle \sum_{i=1}^m \tilde{\mathbf{x}}_i^*, \mathbf{p}_i \right\rangle - \lambda \left\| \mathbb{E}_{\gamma_i} [\tilde{\mathbf{x}}_{\gamma_i}^*] - \mathbb{E}_{\gamma_i} [\tilde{\mathbf{x}}_{\gamma_i}^* | \Gamma_{i-1}] \right\|_2 \\ &\geq -\frac{1}{m} c \left(\sum_{i=1}^m \tilde{\mathbf{x}}_i^* \right) - \lambda \left\| \mathbb{E}_{\gamma_i} [\tilde{\mathbf{x}}_{\gamma_i}^*] - \mathbb{E}_{\gamma_i} [\tilde{\mathbf{x}}_{\gamma_i}^* | \Gamma_{i-1}] \right\|_2. \end{aligned} \quad (12)$$

Here the first inequality is due to Cauchy-Schwarz inequality, while the second inequality is given by Fenchel-Young inequality.

Equation (10), (11) and (12) give us

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^m f_i(\mathbf{p}_i) + v_i(\mathbf{x}_i) \right] &\geq \sum_{i=1}^m \tilde{v}_i(\tilde{\mathbf{x}}_i^*) - c \left(\sum_{i=1}^m \tilde{\mathbf{x}}_i^* \right) \\ &\quad - \lambda \sum_{i=1}^m \mathbb{E}_{\Gamma_{i-1}} \left[\left\| \mathbb{E}_{\gamma_i}[\tilde{\mathbf{x}}_{\gamma_i}^*] - \mathbb{E}_{\gamma_i}[\tilde{\mathbf{x}}_{\gamma_i}^* | \Gamma_{i-1}] \right\|_2 \right]. \end{aligned} \quad (13)$$

Furthermore, Lemma 6 shows that the last sum in Eq. (13) is bounded by $2D_\infty \sqrt{nm}$, and thus Theorem 5 is proved for general OCO algorithms. Finally, to prove the bound for online gradient descent, it is enough to use step size $\eta_i = 2\lambda/D\sqrt{m}$ (recall that D is the ℓ_2 diameter of $\mathcal{C}$) and invoke Lemma 5. $\square$

As we can see from the proof of Theorem 5, it is enough to have continuous valuations. Furthermore, Algorithm 2 still works if consumers only maximize their quasi-linear utilities approximately. Formally, if consumer i finds a bundle $\mathbf{x}_i$ such that $v_i(\mathbf{x}_i) - \langle \mathbf{p}_i, \mathbf{x}_i \rangle \geq \max_{\mathbf{x} \in \mathcal{C}} (v_i(\mathbf{x}) - \langle \mathbf{p}_i, \mathbf{x} \rangle) - \epsilon_i$, then an additive error of ϵ_i will be introduced in Eq. (10). However, as long as the total error $\sum_{i=1}^m \epsilon_i$ is not large, the expected online social welfare of Algorithm 2 will still be close to the offline optimum.

Finally, we state and prove Lemma 6.

Lemma 6. *For any $1 \leq i \leq m$, we have*

$$\mathbb{E}_{\Gamma_{i-1}} \left[\left\| \mathbb{E}_{\gamma_i}[\tilde{\mathbf{x}}_{\gamma_i}^*] - \mathbb{E}_{\gamma_i}[\tilde{\mathbf{x}}_{\gamma_i}^* | \Gamma_{i-1}] \right\|_2 \right] \leq D_\infty \sqrt{\frac{n}{m-i+1}}.$$

Proof. Let $\mathcal{S} = \{s_1, \dots, s_N\}$ denote a finite population of real numbers, and $X_1, \dots, X_n$ ($1 \leq n \leq N$) denote n samples from $\mathcal{S}$ without replacement. Furthermore, let $\mu = \frac{1}{N} \sum_{i=1}^N s_i$ and $\sigma^2 = \frac{1}{N} \sum_{i=1}^N (s_i - \mu)^2$. Then $\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i$ has mean μ and variance $\frac{N-n}{N-1} \frac{\sigma^2}{n}$.

Now come back the the proof of Lemma 6. By Jensen's inequality, we have

$$\mathbb{E}_{\Gamma_{i-1}} \left[\left\| \mathbb{E}_{\gamma_i}[\tilde{\mathbf{x}}_{\gamma_i}^*] - \mathbb{E}_{\gamma_i}[\tilde{\mathbf{x}}_{\gamma_i}^* | \Gamma_{i-1}] \right\|_2 \right] \leq \sqrt{\mathbb{E}_{\Gamma_{i-1}} \left[\left\| \mathbb{E}_{\gamma_i}[\tilde{\mathbf{x}}_{\gamma_i}^*] - \mathbb{E}_{\gamma_i}[\tilde{\mathbf{x}}_{\gamma_i}^* | \Gamma_{i-1}] \right\|_2^2 \right]}.$$

Note that each coordinate of $\mathbb{E}_{\Gamma_{i-1}} \left[\left\| \mathbb{E}_{\gamma_i}[\tilde{\mathbf{x}}_{\gamma_i}^*] - \mathbb{E}_{\gamma_i}[\tilde{\mathbf{x}}_{\gamma_i}^* | \Gamma_{i-1}] \right\|_2^2 \right]$ equals the variance of the average of $m-i+1$ without-replacement samples. Thus we further have

$$\sqrt{\mathbb{E}_{\Gamma_{i-1}} \left[\left\| \mathbb{E}_{\gamma_i}[\tilde{\mathbf{x}}_{\gamma_i}^*] - \mathbb{E}_{\gamma_i}[\tilde{\mathbf{x}}_{\gamma_i}^* | \Gamma_{i-1}] \right\|_2^2 \right]} \leq \sqrt{n \frac{i-1}{m-1} \frac{D_\infty^2}{m-i+1}} \leq D_\infty \sqrt{\frac{n}{m-i+1}}.$$

$\square$

5 Profit Maximization for Separable Valuations and Cost

Previously, social welfare maximization is solved by reducing to a convex optimization problem on the price space. However, profit maximization may be non-convex on both the bundle space and the price space.

Example 1. Consider a market where there is only one consumer, one good, and zero cost. Suppose $v' : [0, 2] \rightarrow \mathbb{R}_+$ is continuous and strictly decreasing, with $v'(1) = 2$, $v'(2) = 1$, and $v'(x) < 2$ for any $x \neq 1, 2$. The integral of v' gives a non-decreasing concave valuation v . It can be shown that the maximum profit is 2, which is attained by price 2 at quantity 1 or price 1 at quantity 2. Thus the set of optimum prices and optimum bundles are both non-convex.

Here we present an algorithm of profit maximization and a nearly matching lower bound when all v_i 's and c are separable. Formally, for every $\mathbf{x} \in \mathcal{C}$ and every $1 \leq i \leq m$, $v_i(\mathbf{x}) = \sum_{j=1}^n v_{ij}(x_j)$, and for every $\mathbf{y} \in m\mathcal{C}$, $c(\mathbf{y}) = \sum_{j=1}^n c_j(y_j)$. Due to the separability assumption, we restate the assumptions on the feasible set, valuation functions and cost function:

- $\mathcal{C} = [0, 1]^n$.
- For every $1 \leq i \leq m$, $1 \leq j \leq n$, v_{ij} is α -strongly concave and λ -Lipschitz continuous.
- For every $1 \leq j \leq n$, c_j is λ -Lipschitz continuous.

The i -th consumer's consumption of good j is completely determined by p_j and is denoted by $x_{ij}(p_j)$. Furthermore, $x_j(p_j) = \sum_{i=1}^m x_{ij}(p_j)$. Our goal is thus to maximize $\text{Profit}_j(p_j) = \sum_{i=1}^n x_j(p_j)p_j - c_j(x_j(p_j))$, for each $1 \leq j \leq n$. Although we can set prices for different goods independently now, to keep consistency, we still consider posting new prices $\mathbf{p} \in \mathbb{R}_+^n$ as one query.

Algorithm 3. Profit Maximization Algorithm for Separable Functions

```

 $r \leftarrow \left\lceil \frac{mn\lambda(\lambda+\alpha)}{\alpha\epsilon} \right\rceil.$ 
 $\tilde{\mathbf{p}} = (0, 0, \dots, 0).$ 
for  $t = 1$  to  $r$  do
  Post prices  $\mathbf{p}_t = \left( \frac{t\alpha\epsilon}{mn(\lambda+\alpha)}, \dots, \frac{t\alpha\epsilon}{mn(\lambda+\alpha)} \right).$ 
  for  $j = 1$  to  $n$  do
    if  $\text{Profit}_j(p_{t,j}) > \text{Profit}_j(\tilde{p}_j)$  then
       $\tilde{p}_j = p_{t,j}$ 
    end if
  end for
end for
Output  $\tilde{\mathbf{p}}.$ 

```

Theorem 6. *The profit given by Algorithm 3 is no less than the optimum profit minus ϵ . The number of queries is $\lceil mn\lambda(\lambda + \alpha)/\alpha\epsilon \rceil$.*

Proof. Fix $1 \leq j \leq n$. Let p_j^* denote the profit-maximizing price for good j . Suppose $\hat{p}_j = \frac{z\alpha\epsilon}{mn(\lambda+\alpha)} \leq p_j^* \leq \frac{(z+1)\alpha\epsilon}{mn(\lambda+\alpha)}$. By the definition of strong smoothness and Lemma 2, we have

$$\begin{aligned}
& x_j(p_j^*)p_j^* - c_j(x_j(p_j^*)) \\
& \leq x_j(\hat{p}_j) \left(\hat{p}_j + \frac{\alpha\epsilon}{mn(\lambda+\alpha)} \right) - c_j(x_j(\hat{p}_j)) + \lambda |x_j(p_j^*) - x_j(\hat{p}_j)| \\
& = x_j(\hat{p}_j) \left(\hat{p}_j + \frac{\alpha\epsilon}{mn(\lambda+\alpha)} \right) - c_j(x_j(\hat{p}_j)) + \lambda \left| \sum_{i=1}^m \left((v_{ij}^*)'(p_j^*) - (v_{ij}^*)'(\hat{p}_j) \right) \right| \\
& \leq x_j(\hat{p}_j) \left(\hat{p}_j + \frac{\alpha\epsilon}{mn(\lambda+\alpha)} \right) - c_j(x_j(\hat{p}_j)) + \lambda \frac{m}{\alpha} \frac{\alpha\epsilon}{mn(\lambda+\alpha)} \\
& = x_j(\hat{p}_j)\hat{p}_j - c_j(x_j(\hat{p}_j)) + x_j(\hat{p}_j) \frac{\alpha\epsilon}{mn(\lambda+\alpha)} + \frac{\lambda\epsilon}{n(\lambda+\alpha)} \\
& \leq x_j(\hat{p}_j)\hat{p}_j - c_j(x_j(\hat{p}_j)) + \frac{\alpha\epsilon}{n(\lambda+\alpha)} + \frac{\lambda\epsilon}{n(\lambda+\alpha)} \\
& = x_j(\hat{p}_j)\hat{p}_j - c_j(x_j(\hat{p}_j)) + \frac{\epsilon}{n}.
\end{aligned}$$

□

Remark 1. Note that if the cost is 0 and thus revenue maximization is considered, then we can set $r = \lceil mn\lambda/\epsilon \rceil$ in Algorithm 3, and it is enough to assume concave valuations.

Theorem 7 shows that the dependency on n and ϵ cannot be improved, even for revenue maximization.

Theorem 7. *The revenue maximization problem needs $\Omega(n/\epsilon)$ queries to get an additive error ϵ , even if the valuations are separable, concave, non-decreasing, and Lipschitz continuous.*

Proof. Let us first consider the case where there is only one consumer and one good. Given $\lambda > 1$, consider ϵ such that there exists an integer q satisfying $(1+\epsilon)^q = \lambda$, and thus $q = \frac{\ln \lambda}{\ln(1+\epsilon)} \geq \frac{\ln \lambda}{\epsilon}$.

Now we are going to define concave functions of the amount of good on $[0, 1]$. It is enough to give a non-decreasing and integrable derivative. Let $v'(x) = \lambda$ on $[0, \frac{1}{\lambda}]$, and $\frac{1}{x}$ on $[\frac{1}{\lambda}, 1]$. One can verify that v' is non-decreasing and integrable, and thus we can integrate it into a concave function v (by shifting, we can ensure $v(0) = 0$). The maximum of $v'(x)x$ on $[0, 1]$ is 1.

We claim that the algorithm has to make at least q queries to ensure an ϵ additive or multiplicative error. If it is not true, there must exist some integer $z \in [0, q-1]$ such that no $x \in \mathcal{I} = (\frac{1}{(1+\epsilon)^{z+1}}, \frac{1}{(1+\epsilon)^z})$ is considered. We can then set $\tilde{v}'(x) = v'(x)$ outside $\mathcal{I}$, $\tilde{v}'(x) = (1+\epsilon)^{z+1}$ on $(\frac{1}{(1+\epsilon)^{z+1}}, \frac{1}{(1+\epsilon)^z})$, and $\tilde{v}'(\frac{1}{(1+\epsilon)^z})$ does not exist. $\tilde{v}'$ is still non-decreasing and integrable, but the optimum revenue is $1 + \epsilon$ now, which is not detected by the algorithm.

In other words, for λ -Lipschitz valuations, at least $\frac{\ln \lambda}{\epsilon}$ queries is required. W.l.o.g., suppose $\lambda = e$ and thus $\frac{1}{\epsilon}$ queries is needed. Now suppose there are n goods, T_j different prices are tested for the j -th good, and the profit we get from good j is within ϵ_j from the maximum profit. Then $T_j \geq \frac{1}{\epsilon_j}$. We have

$$\frac{\sum_{j=1}^n T_j}{n} \geq \frac{n}{\sum_{j=1}^n 1/T_j} \geq \frac{n}{\sum_{j=1}^n \epsilon_j} = \frac{n}{\epsilon}.$$

In other words, $\sum_{j=1}^n T_j \geq \frac{n^2}{\epsilon}$. Since each query can set new prices for n goods, we need $\Omega(\frac{n}{\epsilon})$ queries. $\square$

6 Conclusion and Open Problems

In this paper, we study social welfare and profit maximization with only revealed preferences. The social welfare maximization problem can be solved by reducing to a convex dual optimization problem in both the offline and online case, while profit maximization is essentially non-convex, for which we give nearly matching upper and lower bounds on the query complexity when valuations and cost are separable.

While social welfare maximization is interesting and important, it is still more reasonable for a producer to maximize profit. However, as shown by Example 1, this problem is in general non-convex. While we give an algorithm for the separable case, it is a very interesting open problem to design algorithms for profit maximization in a more general setting or show some hardness result.

References

1. Afriat, S.N.: The construction of utility functions from expenditure data. *Int. Econ. Rev.* **8**(1), 67–77 (1967)
2. Agrawal, S., Devanur, N.R.: Fast algorithms for online stochastic convex programming. In: *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1405–1424. SIAM (2015)
3. Allen-Zhu, Z., Orecchia, L.: Linear coupling: an ultimate unification of gradient and mirror descent. *arXiv preprint arXiv:1407.1537* (2014)
4. Amin, K., Cummings, R., Dworkin, L., Kearns, M., Roth, A.: Online learning and profit maximization from revealed preferences. In: *AAAI*, pp. 770–776 (2015)
5. Babaioff, M., Dughmi, S., Kleinberg, R., Slivkins, A.: Dynamic pricing with limited supply. *ACM Trans. Econ. Comput.* **3**(1), 4 (2015)
6. Balcan, M.-F., Daniely, A., Mehta, R., Urner, R., Vazirani, V.V.: Learning economic parameters from revealed preferences. In: Liu, T.-Y., Qi, Q., Ye, Y. (eds.) *WINE 2014. LNCS*, vol. 8877, pp. 338–353. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-13129-0_28
7. Beigman, E., Vohra, R.: Learning from revealed preference. In: *Proceedings of the 7th ACM Conference on Electronic Commerce*, pp. 36–42. ACM (2006)
8. Besbes, O., Zeevi, A.: Dynamic pricing without knowing the demand function: risk bounds and near-optimal algorithms. *Oper. Res.* **57**(6), 1407–1420 (2009)

9. Besbes, O., Zeevi, A.: Blind network revenue management. *Oper. Res.* **60**(6), 1537–1550 (2012)
10. Blum, A., Mansour, Y., Morgenstern, J.: Learning what's going on: reconstructing preferences and priorities from opaque transactions. In: *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pp. 601–618. ACM (2015)
11. den Boer, A.V.: Dynamic pricing and learning: historical origins, current research, and new directions. *Surv. Oper. Res. Manage. Sci.* **20**(1), 1–18 (2015)
12. den Boer, A.V., Zwart, B.: Simultaneously learning and optimizing using controlled variance pricing. *Manage. Sci.* **60**(3), 770–783 (2013)
13. Broder, J., Rusmevichientong, P.: Dynamic pricing under a general parametric choice model. *Oper. Res.* **60**(4), 965–980 (2012)
14. Devanur, N.R., Hayes, T.P.: The adwords problem: online keyword matching with budgeted bidders under random permutations. In: *Proceedings of the 10th ACM Conference on Electronic commerce*, pp. 71–78. ACM (2009)
15. Diewert, E.: Afriat and revealed preference theory. *Rev. Econ. Stud.* **40**, 419–426 (1973)
16. Dong, J., Roth, A., Schutzman, Z., Waggoner, B., Wu, Z.S.: Strategic classification from revealed preferences. In: *Proceedings of the 2018 ACM Conference on Economics and Computation*, pp. 55–70. ACM (2018)
17. Goel, G., Mehta, A.: Online budgeted matching in random input models with applications to adwords. In: *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 982–991. Society for Industrial and Applied Mathematics (2008)
18. Hazan, E., et al.: Introduction to online convex optimization. *Found. Trends® Optim.* **2**(3–4), 157–325 (2016)
19. Hiriart-Urruty, J.B., Lemaréchal, C.: *Fundamentals of Convex Analysis*. Springer Science and Business Media, Heidelberg (2012). <https://doi.org/10.1007/978-3-642-56468-0>
20. Houthakker, H.S.: Revealed preference and the utility function. *Economica* **17**, 159–174 (1950)
21. Keskin, N.B., Zeevi, A.: Dynamic pricing with an unknown demand model: asymptotically optimal semi-myopic policies. *Oper. Res.* **62**(5), 1142–1167 (2014)
22. Mas-Colell, A.: The recoverability of consumers' preferences from market demand. *Econometrica* **45**(6), 1409–1430 (1977)
23. Mas-Colell, A.: On revealed preference analysis. *Rev. Econ. Stud.* **45**(1), 121–131 (1978)
24. Nesterov, Y.: Smooth minimization of non-smooth functions. *Math. Program.* **103**(1), 127–152 (2005)
25. Nesterov, Y.: A method of solving a convex programming problem with convergence rate $o(1/k^2)$. *Soviet Math. Dokl.* **27**(2), 372–376 (1983)
26. Richter, M.: Revealed preference theory. *Econometrica* **34**(3), 635–645 (1966)
27. Rockafellar, R.T.: *Convex Analysis*. Princeton University Press, Jersey (1970)
28. Roth, A., Slivkins, A., Ullman, J., Wu, Z.S.: Multidimensional dynamic pricing for welfare maximization. In: *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 519–536. ACM (2017)
29. Roth, A., Ullman, J., Wu, Z.S.: Watch and learn: optimizing from revealed preferences feedback. In: *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, pp. 949–962. ACM (2016)
30. Samuelson, P.A.: A note on the pure theory of consumer's behaviour. *Economica* **5**(17), 61–71 (1938)

31. Sion, M., et al.: On general minimax theorems. *Pacific J. Math.* **8**(1), 171–176 (1958)
32. Tseng, P.: On accelerated proximal gradient methods for convex-concave optimization (2008). <http://www.mit.edu/~dimitrib/PTseng/papers/apgm.pdf>
33. Uzawa, H.: Preference and rational choice in the theory of consumption. In: Arrow, K.J., Karlin, S., Suppes, P. (eds.) *Mathematical Models in Social Science* (1960)
34. Varian, H.R.: Revealed preference. In: Szenberg, M., Ramrattand, L., Gottesman, A.A. (eds.) *Samuelsonian Economics and the 21st Century*, pp. 99–115 (2005)
35. Wang, Z., Deng, S., Ye, Y.: Close the gaps: a learning-while-doing algorithm for single-product revenue management problems. *Oper. Res.* **62**(2), 318–331 (2014)
36. Zadimoghaddam, M., Roth, A.: Efficiently learning from revealed preference. In: Goldberg, P.W. (ed.) *WINE 2012. LNCS*, vol. 7695, pp. 114–127. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-35311-6_9