Balancing Social Utility with Aggegator Profit in Electric Vehicle Charging

Elliot Anshelevich, Koushik Kar
Rensselaer Polytechnic Institute,
Troy, NY - 12180, USA.
eanshel@cs.rpi.edu,koushik@ecse.rpi.edu

Shreyas Sekar
University of Washington
Seattle, WA - 98105, USA.
sekarshr@uw.edu

Zaid Tariq
Rensselaer Polytechnic Institute,
Troy, NY - 12180, USA.
tariqz@rpi.edu

***Abstract*—In this paper, we study the efficiency of equilibrium charging schedules of self-interested plug-in electric vehicles (PEVs) in an environment where a monopolistic aggregator sets time-dependent prices to maximize its own profit. Computing profit-maximizing prices is not only NP-hard, they can also lead to bad equilibrium charging schedules. In this work, we present a polynomially computable pricing policy that is efficient both in terms of social welfare (economic efficiency of the charging schedules) and seller profit (profit made the aggregator). We show that the economic efficiency and aggregator profit under the proposed pricing scheme (with respect to their maximum possible values) can be expressed in terms of a regularity (convexity) measure of the utility functions associated with the PEVs. Our numerical studies show that in practice these performance measures are within a small constant factor of the corresponding optimum values. Our pricing policy, which is time-dependent but non-discriminatory across users for any given time-slot, can thus be used to provide a good balance between system and aggregator objectives in PEV charging.**

I. Introduction

In recent years, we are witnessing increasing deployment of plug-in (hybrid) vehicles (PEVs), which can reduce the cost of fuel consumption, greenhouse gas emissions, as well as improve vehicle engine performance efficiency. However, large scale penetration of PEVs would imply a considerable increase in the overall load on the electric grid, particularly during peak demand hours. This can overload the distribution network transformers, increase system losses, and reduce operational grid efficiency [1], [2], [3], [4], [5], [6]. Therefore, effective management of the electricity demand from PEVs will be crucial for maintaining the stability and operational efficiency of the power grid in the near future [7]. Fortunately, PEVs provide significant flexibility in terms of their energy consumption rates and schedules, which can be utilized towards reducing the variability of the aggregate demand over time. In recent literature, the question of PEV charging coordination question has been addressed utilizing a variety of techniques, particularly game theoretic models (e.g. [8], [9], [10], [11]) and optimization methods (e.g. [12], [13], [14]).

Coordination of PEV charging under varying levels of PEV penetration through controlling the price dynamics has also recently been studied in [15]. A congestion pricing based distributed framework for controlling PEV charging has also been studied in [16]. While [17] proposes a dual pricing policy for maximizing system utility, [18] explores a non-linear pricing scheme for the same purpose. Auction mechanisms that attains social optimality while ensuring truthful bidding by PEV agents have been analyzed in [19]. While these pricing strategies may result in socially optimal charging schedules (thereby attaining Walrasian equilibria in the system), the question of profit-maximizing pricing for PEV charging has received very little attention in the literature.

Our work is motivated by the consideration that in deregulated markets, a PEV aggregator may be more interested in maximizing its individual profit rather than optimizing system performance (social utility). Furthermore, Walrasian prices that result in socially optimal charging schedules (such as those proposed in prior work [17], [11], [20]) may not lead to maximum profit for the aggregator (seller); see Sections II and V for examples. Conversely, prices that maximize aggregator profit are not only NP-hard to compute, they can also result in poor social utility (economic efficiency) of the charging schedules (see Appendix). Given this inherent dichotomy between maximizing system utility and maximizing aggregator profit, there is a need for a pricing policy that attains a good balance between these two conflicting objectives, in the context of PEV charging.

Towards this end, we make the following contributions in this work. We consider the problem faced by an aggregator that has to choose time-of-the-day prices (in the granularity of 15 mins to an hour, say) for PEV charging in a distribution grid. In response to these per-unit time-of-the-day charging prices set by the aggregator, the PEV agents in the network determine their charging schedules to maximize their individual utilities subject to their charging constraints (i.e., constraints on the times they can be plugged in). We impose the realistic requirement that the price for charging in a time slot is non-discriminative, i.e., PEVs charging in any given time slot pay the same per-unit price. Such realistic constraints crucially distinguish our model from conceptually similar works [21], [22]: for example, in [21], the authors also look at optimizing both profit and social welfare in charging markets albeit by means of user-specific (discriminative) pricing menus in an online setting where users are available at all times.

While the profit-maximizing price setting problem in our model is NP-hard, we present a polynomial-time algorithm for computing prices that guarantee an approximation factor of $O(\frac{1}{1-\alpha})$ with respect to the optimum profit that the aggregator can derive, where α is some measure of the regularity (convexity) of the utility functions associated with the PEV. We further show that the social welfare (economic surplus) attained by our solution is within a factor $O(\frac{1}{1-\alpha})$ of the the optimal social welfare in the system. For a large class of "reasonable" utility functions, this factor is a small constant. Furthermore, our simulation results (in distribution networks with up to several hundred PEVs) show that in practice, the performance attained by our algorithm is within a small constant (about 2 or 3) of the optimum aggregator profit as well as optimum social welfare. Therefore, the proposed pricing policy can provide a way of balancing the two conflicting goals of optimizing system efficiency and maximizing aggregator profit.

The paper is structured as follows. In Sections II and III we describe our system model and outline some analytical properties of the model that will later be used to analyze the performance of our algorithm. Section IV describes the proposed algorithm and provides bounds on its performance, with respect to optimum aggregator revenue as well as optimum social welfare. In Section V we study the proposed algorithm through simulations, with respect to these two performance criteria. We conclude in Section VI.

II. System Model and Preliminary Analysis

We model the market as a bipartite graph whose vertices are $\mathcal{B}$, $\mathcal{T}$, the set of buyers and time slots respectively. An edge (i, t) in this graph models the constraint that buyer $i \in \mathcal{B}$ is available at time slot $t \in \mathcal{T}$. We make no assumptions on the structure or density of the graph. As each time slot represents a fixed, contiguous time window, we assume that each PEV i can only purchase at most a fixed upper bound ℓ_{it} from each slot $t \in B_i$; moreover, we assume that $\ell_{it} > 0$ if $t \in B_i$. We refer to these as the capacity constraints.

Each time slot t is associated with a *doubly convex* production or generation cost function $C_t(x)$ (with derivative $c_t(x)$), which denotes the cost that the seller incurs while procuring x units of energy. That is, both the cost function and its derivative are convex and non-decreasing. Each buyer i is represented by means of a concave valuation or utility function $u_i(x)$, which denotes this buyer's valuation for purchasing x units of energy from any combination of slots. Important to our analysis is the derivative of this function, which is generally referred to as the *inverse demand function*, i.e., $u_i'(x) = \lambda_i(x)$. We assume that λ_i is continuous for all $i \in \mathcal{B}$, and that $C_t(x)$ is continuously differentiable for all $t \in \mathcal{T}$.

One of our objectives is to understand how the nature of the PEV utility curve affects profit and social utility. To obtain a nuanced understanding along these lines, we assume that the derivative of the utility function, i.e., the inverse demand is parameterized by a single quantity α.

Definition ([23]) A buyer i is said to have an α-strongly regular demand function (α-SR) for $\alpha \in [0, 1]$ if for any $x_1 < x_2$, we have $\frac{\lambda_i(x_2)}{|\lambda_i'(x_2)|} - \frac{\lambda_i(x_1)}{|\lambda_i'(x_1)|} \leq \alpha(x_2 - x_1)$.

The notion of α-strong regularity was first introduced in [23] in the context of characterizing buyer demand functions and has since gained popularity as it smoothly interpolates between two well-studied classes of functions: log-concave functions ($\alpha = 0$) and regular functions ($\alpha = 1$). In some sense, α captures the convexity or volatility of $u'(x)$ as α-strong regularity implies that $\frac{d}{dx}\frac{\lambda(x)}{|\lambda'(x)|} \leq \alpha$. A larger value of α (and hence $\lambda(x)$) implies that the utility function $u_i(x)$ is more concave, whereas a small value of α gives us approximate-linearity. Finally, we remark that even when α is small, our framework encapsulates a large number of interesting utility functions; this includes the well-studied class of log-concave or monotone hazard rate demand functions (e.g, $\lambda(x) = a - x$ or $\lambda(x) = e^{-x}$, as well as any concave function).

A. Pricing Strategies and Buyer Response

The fundamental problem studied in this work is that of profit maximization. Towards this end, we assume that the seller posts one price on each time slot based on his estimation of buyer demand and the generation cost. It is crucial that the price on each time slot applies to any and all buyers who consume energy at that time slot. Once the prices are fixed, each buyer's strategy involves purchasing energy from some combination of time slots to maximize her utility. Specifically let $\vec{p}$ denote the vector of fixed prices decided by the seller. Then, the buyer's best-response strategy is determined by the optimal solution to the following convex program.

$$\begin{aligned} \max \quad & u_i(x_i) - \sum_{t \in \mathcal{T}} p_t x_{it} \\ \text{s.t.} \; & \sum_{t: t \in B_i} x_{it} = x_i \\ & 0 \leq x_{it} \leq \ell_{it} \quad \forall t \in B_i \end{aligned} \tag{1}$$

We refer to the triplet $(\vec{p}, \vec{x}, \vec{y})$ as a valid pricing solution if for some price vector $\vec{p}$, $\vec{x}_i$ denotes buyer i's demand or consumption profile as determined by Convex Program (1) and the scalar x_i denotes this buyer's total consumption. The (allocation) vector $\vec{y}$ is used to define the total amount of energy sold at each time slot: $y_t = \sum_{i: t \in B_i} x_{it}$.

The total social welfare (or social utility) of a pricing solution $(\vec{p}, \vec{x}, \vec{y})$ is defined as $SW(\vec{p}, \vec{x}, \vec{y}) = \sum_{i \in \mathcal{B}} u_i(x_i) - C(\vec{y})$. Similarly, the profit that the seller makes at this solution is given by $\pi(\vec{p}, \vec{x}, \vec{y}) = \sum_{t \in \mathcal{T}} p_t y_t - C(\vec{y})$. Note that we use $C(\vec{y})$ as short-hand for $\sum_{t \in \mathcal{T}} C_t(y_t)$.

Walrsian Equilibrium and Social Utility Maximization: An important benchmark in most of this work is the solution that maximizes the social welfare. Since social welfare does not depend on the prices, we can simply use $(\vec{x^*}, \vec{y^*})$ to denote the demand-allocation pair that maximizes the system welfare. This can be efficiently computed using the following convex program.

$$
\begin{aligned}
\max \quad & \sum_{i \in \mathcal{B}} u_i(x_i) - \sum_{t \in \mathcal{T}} C_t(y_t) \\
\text{s.t.} \; & \sum_{t: t \in B_i} x_{it} = x_i \quad \forall i \in \mathcal{B} \\
& \sum_{i: t \in B_i} x_{it} = y_t \quad \forall t \in \mathcal{T} \\
& 0 \leq x_{it} \leq \ell_{it} \quad \forall i \in \mathcal{B}, t \in B_i
\end{aligned} \tag{2}
$$

We will later show that there exists a slot-price vector $\vec{p^*}$ that *implements* the social welfare maximizing allocation, i.e., $(\vec{p^*}, \vec{x^*}, \vec{y^*})$ is a valid pricing solution that maximizes social utility. Therefore, the central seller can always employ welfare-maximizing prices if that is his end-goal. Unfortunately, as discussed earlier, the seller may want to maximize profit and not the overall welfare. To illustrate this point, we now present a simple example that highlights how buyers maximize utility in the face of prices, and show that this behavior can result in severely sub-optimal profit when the seller employs the Walrasian prices. We later reinforce this point in our experiments in Section V based on realistic parameters.

Example Consider a simplified instance with two buyers (i_1, i_2) and two time slots (t_a, t_b) such that buyer i_1 has access to both time slots whereas buyer i_2 only has access to time slot t_a. The utility functions corresponding to the buyers are $u_{i_1} = x - x^2/2$ whereas $u_{i_2} = 2x - x^2$; the corresponding cost functions are given by $C_{t_a}(y) = 0.125y^2$ and $C_{t_b}(y) = 0.15y^2$. The charging constraints are such that $\ell_{i_1 t_a} = \ell_{i_1 t_b} = \ell_{i_2 t_a} = 0.5$. The socially optimal solution is obtained by using the dual prices and involves setting $p_{t_a} = 0.2$ and $p_{t_b} = 0.15$. At this price, buyer i_1 will purchase 0.5 units from slot t_b and 0.3 units from t_a, whereas buyer i_2 would purchase 0.5 units from slot t_a. On the other hand, if the seller increases the prices to $p'_{t_a} = 1$ and $p'_{t_b} = 0.5$, this would improve his profit by a factor of six (from 0.1175 to 0.6825). Therefore, we infer that, even though the Walrasian prices maximize social welfare, they are not desirable from the seller's perspective owing to diminished profits.

Before concluding this section, we wish to highlight the nature of PEV behavior under a fixed price vector. In the example above, at the welfare-maximizing prices, even though t_b is the minimum priced slot available to PEV i_1, this PEV cannot fully satisfy its demand using only slot t_b and must purchase almost one-third of its total battery capacity from a higher-priced time interval. Such non-trivial and 'sharply discrete' behavior crucially distinguishes our setting from other papers with bipartite network markets where each buyer can meet her entire demand from a single slot [22]. This sharp difference in behavior necessitates new tools for analysis, which we provide in the following two sections via simple characterizations of pricing solutions and seller profit.

III. Characterizing Pricing Solutions

In this section, our objective is to provide a partial characterization of how users react to prices by selecting time slots and charging quantities. For that purpose, let $\vec{p}, \vec{x}, \vec{y}$ denote an arbitrary feasible solution for an instance $\mathcal{I}$ of our problem.

Proposition III.1. *A pricing solution $\vec{p}, \vec{x}, \vec{y}$ for a given instance is said to be valid if and only if for every buyer i and slots $t, t' \in B_i$, all of the following conditions are true:*

1) *If $x_{it} > 0$, then $\lambda_i(x_i) \geq p_t$.*
2) *If $x_{it} > 0$ and $x_{i,t'} = 0$, then $p_t \leq p_{t'}$.*
3) *If $x_{it}, x_{it'} > 0$ and $p_t < p_{t'}$, then, $x_{it} = \ell_{it}$.*
4) *If $x_{it'} < \ell_{it'}$, then $\lambda_i(x_i) \leq p_{t'}$.*

Proof. All of these conditions follow directly from the KKT conditions on the convex program corresponding to buyer utility maximization. However, we sketch their proofs for convenience.

Let us begin with the forward direction, i.e., if the given solution is a valid pricing solution, then the above conditions must be true. All of the individual proofs below proceed by contradiction.

1) Suppose that $x_{it} > 0$ and $\lambda_i(x_i) < p_t$, then there exists a sufficiently small $\epsilon > 0$ such that the PEV can strictly increase her utility by dropping an ϵ amount of demand from slot t. Note that upon dropping a small amount of demand from slot t, the buyer's change in utility is given by $p_t\epsilon - \int_{x=x_i-\epsilon}^{x_i} \lambda_i(x)dx$.
2) Suppose that $x_{it} > 0$, $x_{i,t'} = 0$ but $p_t > p_{t'}$. It is not hard to deduce that the buyer's utility strictly increases by transferring all of her flow from slot t to slot t'.
3) If this condition were not true, then the buyer could increase her utility by reducing her consumption of slot t' by some $\epsilon > 0$, and increasing consumption on slot t by the same ϵ ensuring that $x_{it} + \epsilon \leq \ell_{it}$.
4) Finally, if $x_{it'} < \ell_{it'}$ but $\lambda_i(x_i)$ turns out to be larger than $p_{t'}$, then the buyer can strictly increase her utility by consuming (an ϵ amount) more of slot t'.

The reverse direction follows from the same kind of intuition but is more technically involved and we defer it to the Appendix. Suppose that all of the conditions presented above are true, but there exists another consumption vector $\vec{x'}_i$ that yields maximum utility for the given price vector, and that this maximum utility is strictly larger than the one guaranteed by $\vec{x}_i$.

According to the conditions specified, if $x_{it} > 0$, $\lambda_i(x_i) \geq p_t$, so the user cannot increase her utility by (only) strictly reducing her consumption on any slot t. Similarly, since $x_{it} < \ell_{it}$ implies that $\lambda_i(x_i) \leq p_t$, so the user cannot increase her utility by (only) strictly increasing her consumption on some slots. So, we can rule out the fact that one of $\vec{x}_i$ or $\vec{x'}_i$ dominates the other. It follows that there exists a pair of slots t, t' satisfying $x_{it} > x'_{it}$ and $x'_{it'} < x'_{it'}$.

Now, we are ready to complete the proof. By definition $\vec{x}_i$ satisfies the conditions specified in the proposition. Since $\vec{x'}_i$ is a utility-maximizing demand vector, it also satisfies the same set of conditions. Therefore, applying conditions (1) and (4) with respect to both t and t' and the two demand vectors, we get that $\lambda_i(x_i) \geq p_t \geq \lambda_i(x'_i)$ and $\lambda_i(x_i) \leq p_{t'} \leq \lambda_i(x'_i)$. We conclude that $\lambda_i(x_i) = \lambda_i(x_i')$. Moreover, every pair of slots t, t' where the consumption in $\vec{x}_i$, $\vec{x'}_i$ differs must have the exact same price, which equals $p = \lambda_i(x_i) = \lambda_i(x_i')$.

Suppose that T' denotes the set of slots where the consumption in the two vectors $\vec{x}_i$, $\vec{x'}_i$ do not coincide. Subtracting the buyer utility in the two cases, we get that

$$\begin{aligned} & u_i(x_i) - u_i(x_i') - \sum_{t \in \mathcal{T}} p_t(x_{it} - x_{it}') \\ = \;& \lambda_i(x_i')(x_i - x_i') - \sum_{t \in T'} p(x_{it} - x_{it}') \\ = \;& p(x_i - x_i') - p \sum_{t \in T'} (x_{it} - x_{it}') = 0. \end{aligned}$$

This contradicts the fact that $\vec{x}_i$ is not a utility maximizing consumption vector for buyer i. □

Recall that $(\vec{x^*}, \vec{y^*})$ is the social welfare maximizing allocation pair. We now identify a price vector $\vec{p^*}$ to confirm the presence of valid pricing solutions that maximize social welfare.

Proposition III.2. *Given the social welfare maximizing allocation, define the price vector $\vec{p^*}$ such that for every time slot t, $p_t^* = c_t(y_t^*)$. Then, $(\vec{p^*}, \vec{x^*}, \vec{y^*})$ denotes a valid pricing solution of our problem.*

Proof. To prove the claim, it is sufficient to prove that $(\vec{p^*}, \vec{x^*}, \vec{y^*})$ satisfies the four conditions listed in Proposition III.1. Once again, consider some arbitrary buyer i and slots t, t'.

1) Suppose that $x_{it}^* > 0$. We know that $p_t^* = c_t(y_t^*)$. Assume by contradiction that $p_t^* > \lambda_i(x_i^*)$. Let $0 \le \epsilon \le x_{it}^*$ denote some sufficiently small constant such that
$$\lambda_i(x_i^*) \le \lambda_i(x_i^* - \epsilon) < c_t(y_t^* - \epsilon) \le c_t(y_t^*) = p_t.$$
Then, the change in social welfare if buyer i decreases her consumption on slot t by ϵ is at least $c_t(y_t^* - \epsilon)\epsilon - \lambda_i(x_i^* - \epsilon)\epsilon > 0$. Clearly, this violates the fact that $(\vec{x^*}, \vec{y^*})$ is a social welfare maximizing allocation.
2) Suppose that $x_{it}^* > 0$ and $x_{it'}^* = 0$. Then, according to the properties of a social welfare maximizing allocation, it must be the case that $c_t(y_t^*) \le c_t(y_{t'}^*)$ or else we could simply transfer flow from slot t to t' for this buyer and decrease the total cost. Therefore, $p_t^* \le p_{t'}^*$.
3) Suppose that $x_{it}^*, x_{it'}^* > 0$ and $c_t(y_t^*) < c_t(y_{t'}^*)$. Assume by contradiction that $x_{it}^* < \ell_{it}$: then clearly, one could increase the social welfare by transferring an $\epsilon > 0$ amount of flow corresponding to buyer i from slot t' to t.
4) Finally, suppose that $x_{it'}^* < \ell_{it'}$ for some slot t'. If $\lambda_i(x_i^*) > c_{t'}(y_{t'}^*)$, then we could increase the amount of energy that the buyer consumes from slot t' by some infinitesimal amount and increase the total social welfare. Therefore, by contradiction, it must be that $p_{t'}^* = c_{t'}(y_{t'}^*) \ge \lambda_i(x_i^*)$.

□

A. The Maximum Price Assumption

In private markets that supply an 'essential good', it is reasonable to assume that the seller is incentivized or in some cases even constrained by regulatory authorities to not employ extremely large prices. Along those lines, we assume that the price that the seller is allowed to choose for each item t is at most some global constant $\bar{P}$. An important assumption that we make in this work is that the global price cap $\bar{P}$ cannot be larger than any PEV's maximum valuation, i.e.,

Assumption For all $i \in \mathcal{B}$, $\bar{P} \le \lambda_i(0)$.

We refer to this as the *maximum price assumption*. The max-price assumption ensures that the prices are not 'so high' that PEV owners are forced to drop out of the market. At the same time, to ensure that the parameter is actually meaningful, we will also make the following assumption that the prices are not so low that the seller actually makes a loss, i.e.,

Assumption For every slot t, $p_t^* \le \bar{P}$.

Note that we only make this latter assumption for convenience of exposition. All of our results actually hold even without it.

B. Upper Bound on Optimum Revenue

In this section, we will provide a new upper bound for the optimum revenue in terms of the social welfare of a reduced instance that is closely associated with the original instance of our profit maximization problem. Given an instance $\mathcal{I}$ of the problem, define a new instance $\mathcal{I}'$ with the same bipartite graph, cost functions, and charging constraints on the time slots but new inverse demand functions defined as follows.

$$\bar{\lambda}_i(x) = \min(\lambda_i(x), \bar{P}) \quad \forall i \in \mathcal{B}$$

We now show how the social welfare and profit maximizing solutions for this new instance behave in terms of the old instance. To avoid excessively lengthy notation, we will use $\overline{SW}(\vec{p}, \vec{x}, \vec{y})$ to denote the social welfare of a pricing solution for the reduced instance and $\bar{\pi}(\vec{p}, \vec{x}, \vec{y})$ for the profit. Finally, for every buyer i, define $\bar{x}_i$ to be (one of) the values at which $\lambda_i(\bar{x}_i) = \bar{P}$.

Lemma III.3. *Suppose that $(\vec{p}, \vec{x}, \vec{y})$ is a valid pricing solution for the reduced instance satisfying $p_t \le \bar{P}$ for all $t \in \mathcal{T}$. Then, $(\vec{p}, \vec{x}, \vec{y})$ is also a valid pricing solution for the original instance. Moreover, we can represent the social welfare and profit of the pricing solution with respect to the reduced instance as :*

$$\begin{aligned} \overline{SW}(\vec{p}, \vec{x}, \vec{y}) &= SW(\vec{p}, \vec{x}, \vec{y}) - \sum_{i \in \mathcal{B}} (u_i(\bar{x}_i) - \bar{P}\bar{x}_i) \\ \bar{\pi}(\vec{p}, \vec{x}, \vec{y}) &= \pi(\vec{p}, \vec{x}, \vec{y}). \end{aligned}$$

For the rest of this work, we will also assume that $(\vec{p}^{opt}, \vec{x}^{opt}, \vec{y}^{opt})$ is the pricing solution with the maximum profit for the given instance that also obeys the maximum price condition. Let π^{opt} denote the value of the maximum profit.

Lemma III.4. *Given any instance $\mathcal{I}$, and the associated reduced instance $\mathcal{I}'$, the social welfare and profit maximizing solutions for the reduced instance coincide with those of the original instance. i.e.,*

1) *The social welfare maximizing solution for $\mathcal{I}'$ is $(\vec{p^*}, \vec{x^*}, \vec{y^*})$.*
2) *The profit-maximizing solution for $\mathcal{I}'$ is $(\vec{p}^{opt}, \vec{x}^{opt}, \vec{y}^{opt})$.*

Proof. We know that each element of the price vectors $\vec{p^*}$ and $\vec{p}^{opt}$ is smaller than or equal to $\bar{P}$. Therefore, it is not hard to verify that both $(\vec{p^*}, \vec{x^*}, \vec{y^*})$ and $(\vec{p}^{opt}, \vec{x}^{opt}, \vec{y}^{opt})$ are valid solutions for the reduced instance.

For a given instance, $\bar{x}_i$ is a constant. Therefore, for any given pricing solution that is valid for the reduced instance, its social welfare is the same as the social welfare of that solution for the original instance minus a constant term. Since $(\vec{p^*}, \vec{x^*}, \vec{y^*})$ is valid for the reduced instance, its optimality follows.

The same argument applies for profit except that in this case, the exact value of the profit actually coincides for both instances. □

We now provide a simple upper bound on the optimum revenue in terms of the social welfare of a carefully selected allocation.

Proposition III.5. *Given an instance of the problem $\mathcal{I}$, the optimum profit π^{opt} for this instance is no larger than $\overline{SW}(\vec{p^*}, \vec{x^*}, \vec{y^*})$.*

Proof. From Lemma III.4, we know that $\pi(\vec{p}^{opt}, \vec{x}^{opt}, \vec{y}^{opt}) = \bar{\pi}(\vec{p}^{opt}, \vec{x}^{opt}, \vec{y}^{opt})$. Moreover, $\bar{\pi}(\vec{p}^{opt}, \vec{x}^{opt}, \vec{y}^{opt}) \leq \overline{SW}(\vec{p}^{opt}, \vec{x}^{opt}, \vec{y}^{opt})$ since the profit of any pricing solution for the reduced instance is smaller than its social welfare. Finally, since $(\vec{p^*}, \vec{x^*}, \vec{y^*})$ denotes the welfare maximizing solution for the reduced instance, we also have that $\overline{SW}(\vec{p}^{opt}, \vec{x}^{opt}, \vec{y}^{opt}) \leq \overline{SW}(\vec{p^*}, \vec{x^*}, \vec{y^*})$. This completes the proof. □

IV. Profit Maximization Algorithm and Performance Bounds

The main problem studied in this paper is that of efficiently computing a per-slot pricing vector $\vec{\tilde{p}}$ in order to maximize the profit $\pi(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$, where $(\vec{\tilde{x}}, \vec{\tilde{y}})$ denote the valid buyer response vector corresponding to price $\vec{\tilde{p}}$. Unfortunately, the profit-maximization is NP-Hard to compute even for a special case of our problem where the cost functions are zero. Bearing this in mind, we present a simple polynomial-time algorithm for computing prices for each slot. Our main theoretical results are the following:

1) *(Theorem 1)* The profit guaranteed is within a $O(\frac{1}{1-\alpha})$-factor of the optimum profit.
2) *(Theorem 2)* In addition to the above profit, the algorithm simultaneously provides a $O(\frac{1}{1-\alpha})$-approximation for social welfare as well, with better constant factors than the profit bound. (See Figure V).
3) *(Theorem 3)* If the seller ignores computational considerations and employs the profit-maximizing prices, then we provide a lower bound on the resulting social welfare (i.e., system efficiency).

Recall that α is a convexity measure for the PEV utility functions and with increasing α, the derivative of the utility function becomes more volatile, and the utility function itself, more non-linear.

Main Algorithm for Computing Prices

For every slot t, set

$$\tilde{p}_t = \max(p^*_t, \bar{P}(1-\alpha)^{\frac{1}{\alpha}}).$$

Let $\vec{\tilde{x}}, \vec{\tilde{y}}$ denote a best-response demand and allocation vector corresponding to the pricing vector $(\tilde{p})_t$ such that $(\tilde{y})_t$ has minimum cost over all such vectors.

Before stating the theorem, we make the following simple observation. For every slot t, $\tilde{p}_t \leq \bar{P}$ since $p^*_t \leq \bar{P}$ and $(1-\alpha)^{\frac{1}{\alpha}} \leq 1$. Therefore, our pricing vector does not violate the hard constraint on the maximum price. We now state an obvious proposition without proof.

Proposition IV.1. *For any two time slots $t, t' \in \mathcal{T}$, if $p^*_t \leq p^*_{t'}$, then $\tilde{p}_t \leq \tilde{p}_{t'}$. In other words, our algorithm yields a pricing scheme without disturbing the relative ordering of prices between various time slots in $\vec{p^*}$.*

The main implication stemming from this proposition is that the proposed pricing strategy enjoys the same *demand management* properties of the social welfare maximizing solution, i.e., since peak-demand slots are priced higher, the demand is evenly distributed on the different intervals. We now prove that this deceptively simple strategy leads to a reasonable approximation with respect to the optimum profit.

Theorem 1. *For every given instance with α-SR demand functions, our algorithm returns a profit that is within a γ-factor (see below) of the optimum profit.*

$$\gamma = \left(2(\frac{1}{1-\alpha})^{\frac{1}{\alpha}} - 1 + \frac{1}{1-\alpha}\right) = \Theta(\frac{1}{1-\alpha})$$

In Figure V, we illustrate how the profit as well as the social welfare guaranteed by our the algorithm varies as we increase α along with the experimental results. The guarantee for social utility comes from Theorem 2 in the next section. We remark that although our guarantee worsens with increasing α, we get constant factor approximations even for reasonably large values of α: for example at $\alpha = 0.7$, we get a 13.5-approximation and a 4.33-approximation for welfare. Moreover, in our experiments, we obtain constant factor approximations even for large values of α.

Proof. We in fact prove a stronger result, namely that

$$\frac{\overline{SW}(\vec{p^*}, \vec{x^*}, \vec{y^*})}{\pi(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})} \leq \gamma.$$

First, define $\zeta = 2(\frac{1}{1-\alpha})^{\frac{1}{\alpha}}$. The proof proceeds in two parts. We first prove that $\overline{SW}(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}}) \leq (\zeta - 1)\pi(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$.

Following this, we show that $\overline{SW}(\vec{p^*}, \vec{x^*}, \vec{y^*}) - \overline{SW}(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}}) \leq \frac{1}{1-\alpha}\pi(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$. Adding up the two parts gives us a lower bound for the profit of the solution returned by our algorithm in terms of the optimum social welfare of the reduced instance, which we know to be larger than the optimum profit for the original instance (Lemma III.4).

General Characterization of Algorithm's Pricing Solution: We begin by characterizing the pricing solution computed by our algorithm in comparison to the social welfare maximizing pricing solution. Recall that for every slot $t \in \mathcal{T}$, $p^*_t \leq \tilde{p}$: we now prove a rather technical claim that highlights a simple property of our pricing solution, namely that for every $t \in \mathcal{T}$, $c_t(\tilde{y}_t) \leq c_t(y^*_t)$. Based on this, we can also conclude that $C(\vec{y}) \leq C(\vec{y^*})$.

Claim IV.2. *Consider the two pricing solutions $(\vec{p^*}, \vec{x^*}, \vec{y^*})$ and $(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$. For any given slot $t \in \mathcal{T}$, we have that $c_t(y^*_t) \geq c_t(\tilde{y}_t)$.*

Proof. The proof of this claim is somewhat involved, so we proceed carefully by contradiction. Suppose that for some item t, $c_t(\tilde{y}_t) > c_t(y^*_t)$. Since the marginal cost function is monotone non-decreasing, this must mean that $\tilde{y}_t > y^*_t$.

Now let us construct the following graph $G' = (\mathcal{T}, E')$ where $\mathcal{T}$ is the set of all slots. We say that there is a directed edge from slot t_1 to t_2 if there exists some buyer i such that

$$\tilde{x}_{it_1} > x^*_{it_1} \text{ and } \tilde{x}_{it_2} < x^*_{it_2}.$$

In simple terms, this means that i is receiving more amount of t_2 and less of t_1 in $\vec{x^*}$ than what she received in $\vec{\tilde{x}}$. This immediately implies the following set of inequalities:

$$\tilde{x}_{it_1} > 0, \quad x^*_{it_1} < \ell_{it_1}, \quad \tilde{x}_{it_2} < \ell_{it_2}, \quad x^*_{it_2} > 0.$$

Based on this, we present a simple sub-claim.

Lemma IV.3. *Consider any $(t_1, t_2) \in E'$: it must be the case that $p^*_{t_1} \geq p^*_{t_2}$ and therefore, $\tilde{p}_{t_1} \geq \tilde{p}_{t_2}$.*

Proof. The proof of $p^*_{t_1} \geq p^*_{t_2}$ comes from applying conditions (1), (4) from Proposition III.1 with respect to t_1 and t_2 respectively. The second inequality then follows from Proposition IV.1. □

Now, look at slot t. Since the total allocation from this slot is smaller in $\vec{y^*}$, this must mean that there is at least one buyer i who is consuming less of t in $\vec{y^*}$ as compared to $\vec{\tilde{y}}$. However, the total demand of i is only larger in $\vec{x^*}$, which means there must be some other slot t_1, which she is consuming more of in $\vec{x^*}$. This implies that $(t, t_1) \in E'$.

Suppose that S_t represents the set of vertices that are reachable from t in G' including t itself. We have already shown that S_t has at least one slot other than t. Our first claim is that all the nodes in S_t have a marginal cost in $\vec{y^*}$ that is no larger than the marginal cost of t in the same allocation. To show this consider an edge (t_1, t_2) where both the slots belong to S_t. By definition, there must be some buyer who has access to both these slots and is consuming non-zero amount of t_2 in $\vec{x^*}$. Since $(\vec{x^*}, \vec{y^*})$ is a welfare-maximizing allocation independent of $\vec{p^*}$, the cost of this allocation cannot decrease if we reduce buyer i's consumption from slot t_2 in $\vec{x^*}$ and increase consumption on t_1 by the same amount. Therefore, $c_{t_1}(y^*_{t_1}) \geq c_{t_2}(y^*_{t_2})$. Applying this transitively from t, all nodes reachable from t in G' must have a marginal cost in $\vec{y^*}$ smaller than or equal to $c_t(y^*_t)$.

Our second claim is that for every $t' \in S_t$, $c_{t'}(\tilde{y}_{t'}) \geq c_t(\tilde{y}_t)$. The proof proceeds in a similar way as our previous claim. Once again, consider $(t_1, t_2) \in E'$ both belonging to S_t, and let i be some buyer satisfying $\tilde{x}_{it_1} > x^*_{it_1}$ and $\tilde{x}_{it_2} < x^*_{it_2}$. We know from Lemma IV.3 that $\tilde{p}_{t_2} \leq \tilde{p}_{t_1}$. Among all valid pricing solutions of the form $(\vec{\tilde{p}}, \vec{x}, \vec{y})$, $(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$ denotes the one of minimal cost. Therefore, it must be the case that $c_{t_2}(\tilde{y}_{t_2}) \geq c_{t_1}(\tilde{y}_{t_1})$ or else we could transfer some of buyer i's purchases from t_1 to t_2. Using these inequalities regarding the marginal costs in $\vec{\tilde{y}}$ and $\vec{y^*}$, we get for all $t' \in S_t$,

$$c_{t'}(y^*_{t'}) \leq c_t(y^*_t) < c_t(\tilde{y}_t) \leq c_{t'}(\tilde{y}_{t'}). \tag{3}$$

What this means is that for all the slots in S_t, the total consumption is larger in $\vec{\tilde{y}}$ as compared to $\vec{y^*}$. Suppose that B^1_t is the complete set of buyers who consume non-zero amounts of the slots in S_t in $\vec{\tilde{x}}$. Our final claim is that every buyer in B^1_t consumes more or equal amount of the slots in S_t in $\vec{x^*}$ as compared to $\vec{\tilde{x}}$. That is for each buyer $i \in B^1_t$,

$$\sum_{t' \in S_t} \tilde{x}_{it'} \leq \sum_{t' \in S_t} x^*_{it'}.$$

To prove this, notice that for any buyer i, if this is not true, then there must exist at least one $t' \in S_t$ which she receives more of in $\vec{\tilde{x}}$ than $\vec{x^*}$. However, $x^*_i \geq \tilde{x}_i$, and so, there must be some t_3 outside of S_t such that buyer i consumes more of t_3 in $\vec{x^*}$ as compared to $\tilde{x}$. But this means that there must be an edge from t' to t_3 in G' and so $t_3 \in S_t$, a contradiction.

Now, we are ready to prove our main result. Recall that $\forall t' \in S_t$, $\tilde{y}_{t'} > y^*_{t'}$. Since for all $t' \in S_t$, the consumption in $\vec{\tilde{y}}$ can only come from the buyers in B^1_t (by definition), we have,

$$\begin{aligned}\sum_{t' \in S_t} \tilde{y}_{t'} &= \sum_{i \in B^1_t} \sum_{t' \in S_t} \tilde{x}_{it'} \\ &\leq \sum_{i \in B^1_t} \sum_{t' \in S_t} x^*_{it'} \\ &\leq \sum_{t' \in S_t} y^*_{t'} < \sum_{t' \in S_t} \tilde{y}_{t'}.\end{aligned}$$

This is a contradiction. □

Lemma IV.4. *The total cost incurred in $(\vec{p^*}, \vec{x^*}, \vec{y^*})$ is larger than or equal to the cost incurred in the solution $(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$, i.e., $C(\vec{y^*}) \geq C(\vec{\tilde{y}})$.*

Proof. From Claim IV.2, it is clear that for every $t \in \mathcal{T}$, we have $c_t(y^*_t) \geq c_t(\tilde{y}_t)$. The production cost function C is doubly convex and continuously differentiable, and therefore, we expect its derivative c to be convex and strictly increasing.

Therefore $c_t(y_t^*) \geq c_t(\tilde{y}_t)$ implies that $y_t^* \geq \tilde{y}_t$ and the lemma trivially follows. $\square$

Next, we prove an interesting dichotomy obeyed by our pricing solution $(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$, namely that either every PEV behaves exactly as it would at the social optimum (i.e., $\tilde{x}_i = x_i^*$) or that its derivative at $\tilde{x}_i$ is sufficient large (i.e., $|\lambda_i'(\tilde{x}_i)|\tilde{x}_i \geq \lambda_i(\tilde{x}_i)$). We will later use the second part of the dichotomy to show that since λ is non-increasing and its derivative is large enough, the loss in social welfare from $\tilde{x}_i$ to x_i^* can be bounded for each PEV.

Lemma IV.5. *Consider the pricing solution $(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$ computed by our algorithm. For every PEV $i \in \mathcal{B}$, one of the following is true:*

1) $\lambda_i(\tilde{x}_i) = \lambda_i(x_i^*)$.
2) $\frac{\lambda_i(\tilde{x}_i)}{|\lambda_i'(\tilde{x}_i)|} \leq \tilde{x}_i$.

Proof. Recall that as per the definition of the pricing scheme in our algorithm, for every slot t, $\tilde{p}_t \geq p_t^*$. Therefore, it is not hard to see that for every PEV i, $\lambda_i(\tilde{x}_i) \geq \lambda_i(x_i^*)$. Now, we need to show that for every i where $\lambda_i(\tilde{x}_i) > \lambda_i(x_i^*)$, the second condition must hold. Let i be some such buyer.

Since the inverse demand λ_i is continuous, it must be the case that there exists some slot $t \in B_i$ where $x_{it}^* > \tilde{x}_{it}$, and therefore, $\tilde{x}_{it} < \ell_{it}$. Applying property (4) of Proposition III.1 with respect to the solution $(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$, we get that $\lambda_i(\tilde{x}_i) \leq \tilde{p}_t$. However, $x_{it}^* > \tilde{x}_{it}$ also implies that $x_{it}^* > 0$ and using property (1) of the same proposition with respect to $(\vec{p^*}, \vec{x^*}, \vec{y^*})$, we have that $\lambda_i(x_i^*) \geq p_t^*$.

Combining the two insights along with our original assumption that $\lambda_i(x_i^*) < \lambda_i(\tilde{x}_i)$, we get that

$$p_t^* \leq \lambda_i(x_i^*) < \lambda_i(\tilde{x}_i) \leq \tilde{p}_t.$$

Since $\tilde{p}_t = \max(p_t^*, \bar{P}(1-\alpha)^{\frac{1}{\alpha}})$, we can safely conclude that $\tilde{p}_t = \bar{P}(1-\alpha)^{\frac{1}{\alpha}} \leq \lambda_i(0)(1-\alpha)^{\frac{1}{\alpha}}$ by the definition of $\bar{P}$.

The proof now directly follows from Lemma A.3. $\square$

Armed with our characterization results, we are ready to prove the upper bounds for $\overline{SW}(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$ and $\overline{SW}(\vec{p^*}, \vec{x^*}, \vec{y^*}) - \overline{SW}(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$ in terms of $\pi(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$.

Claim IV.6. *Suppose that $(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$ is the pricing solution as defined by our algorithm. Then, the total social welfare of this solution with respect to the reduced instance $\overline{SW}(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$ is at most a factor $\zeta - 1$ times the profit due to this solution $\pi(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$.*

The proof of this claim is adapted from that of a similar claim in [22], the main difference being that $\lambda_i(0)$ is no longer the same for all buyers i. *(Proof Sketch)*

Proof. The social welfare of the current solution with respect to the reduced instance is $\sum_{i\in B} \bar{u}_i(\tilde{x}_i) - C(\vec{\tilde{y}})$. The function $\bar{u}_i$ is concave for every $i \in B$ since its derivative $\bar{\lambda}_i(x)$ is non-increasing with x. Therefore $\bar{u}_i(\tilde{x}_i) \leq \bar{u}_i'(0) \cdot \tilde{x}_i = \bar{P} \cdot \tilde{x}_i$. Moreover, by definition, for every slot t, $\tilde{p}_t \geq \bar{P}(1-\alpha)^{\frac{1}{\alpha}}$.

So, our first inequality is the following,

$$\begin{aligned}\sum_{i\in B} \bar{u}_i(\tilde{x}_i) - C(\vec{\tilde{y}}) &\leq \sum_{i\in\mathcal{B}} \bar{P}\tilde{x}_i - C(\vec{\tilde{y}}) \\ &= \frac{\bar{P}}{\bar{P}(1-\alpha)^{\frac{1}{\alpha}}} \sum_{i\in\mathcal{B}} \bar{P}(1-\alpha)^{\frac{1}{\alpha}}\tilde{x}_i - C(\vec{\tilde{y}}) \\ &\leq \frac{1}{(1-\alpha)^{\frac{1}{\alpha}}} \sum_{t\in\mathcal{T}} \tilde{p}_t\tilde{y}_t - C(\vec{\tilde{y}}).\end{aligned}$$

The final inequality comes from the fact that $\bar{P}(1-\alpha)^{\frac{1}{\alpha}} \leq \tilde{p}_t$ for all $t \in T$ and from rearranging the allocation from the buyers to the goods. Now, the total profit that the seller makes at the given strategy $\pi(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})$ equals $\sum_{t\in T} \tilde{p}_t\tilde{y}_t - C(\vec{y})$. Using this, we get the following upper bound for the ratio of the welfare to profit

$$\begin{aligned}\frac{\sum_{i\in B} \bar{u}_i(\tilde{x}_i) - C(\vec{\tilde{y}})}{\pi(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}})} &\leq \frac{\frac{1}{(1-\alpha)^{\frac{1}{\alpha}}}\sum_{t\in T} \tilde{p}_t\tilde{y}_t - C(\vec{\tilde{y}})}{\sum_{t\in T} \tilde{p}_t\tilde{y}_t - C(\vec{y})} \\ &\leq \frac{\frac{1}{(1-\alpha)^{\frac{1}{\alpha}}}\sum_{t\in T} \tilde{p}_t\tilde{y}_t - \sum_{t\in T}\frac{1}{2}c_t(\tilde{y}_t)\tilde{y}_t}{\sum_{t\in T}[\tilde{p}_t - \frac{1}{2}c_t(\tilde{y}_t)]\tilde{y}_t} \\ &\leq \frac{\sum_{t\in T}[\frac{1}{(1-\alpha)^{\frac{1}{\alpha}}}\tilde{p}_t\tilde{y}_t - \frac{1}{2}\tilde{p}_t\tilde{y}_t]}{\sum_{t\in T}[\tilde{p}_t\tilde{y}_t - \frac{1}{2}\tilde{p}_t\tilde{y}_t]} \\ &= 2\frac{1}{(1-\alpha)^{\frac{1}{\alpha}}} - 1.\end{aligned}$$

The second inequality above comes from the definition of doubly convex cost functions according to which $C_t(\tilde{y}_t) \leq \frac{1}{2}c_t(\tilde{y}_t)\tilde{y}_t$. The third inequality comes from the fact that for every t, $c_t(\tilde{y}_t) \leq \tilde{p}_t$. To see why this is true, observe that for every $t \in \mathcal{T}$, $c_t(\tilde{y}_t) \leq c_t(y_t^*) = p_t^* \leq \tilde{p}_t$. This completes the proof of the first claim. $\square$

Claim IV.7. *The following inequality holds for the solution returned by our algorithm:*

$$\overline{SW}(\vec{p^*}, \vec{x^*}, \vec{y^*}) - \overline{SW}(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}}) \leq \frac{1}{1-\alpha}\pi(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}}).$$

Proof. We need to prove an upper bound on $SW_2 := \overline{SW}(\vec{p^*}, \vec{x^*}, \vec{y^*}) - \overline{SW}(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}}) = \sum_{i\in\mathcal{B}}[u_i(x_i^*) - u_i(\tilde{x}_i)] - [C(\vec{y^*}) - C(\vec{\tilde{y}})] \leq \sum_{i\in\mathcal{B}}[u_i(x_i^*) - u_i(\tilde{x}_i)]$. The inequality comes from the fact that the difference in costs is positive as shown in Lemma IV.4. Moreover, observe that for any two pricing solutions $(\vec{p_1}, \vec{x_1}, \vec{y_1})$ and $(\vec{p_2}, \vec{x_2}, \vec{y_2})$, we have that $\overline{SW}(\vec{p_1}, \vec{x_1}, \vec{y_1}) - \overline{SW}(\vec{p_2}, \vec{x_2}, \vec{y_2}) = SW(\vec{p_1}, \vec{x_1}, \vec{y_1}) - SW(\vec{p_2}, \vec{x_2}, \vec{y_2})$ due to Lemma III.3.

Therefore, the LHS can be simplified as

$$SW_2 \leq \sum_{i\in\mathcal{B}}[\int_{\tilde{x}_i}^{x_i^*} \lambda_i(x)dx$$

Since λ_i is α-SR for all $i \in \mathcal{B}$, we can use Lemma A.4 to bound the integral as follows $\int_{\tilde{x}_i}^{x_i^*} \lambda_i(x)dx \leq$

$\frac{1}{1-\alpha}(\frac{\lambda_i(\tilde{x}_i)}{|\lambda_i'(\tilde{x}_i)|})(\lambda_i(\tilde{x}_i)-\lambda_i(x_i^*))$. Summing this up, one obtains,

$$SW_2 \leq \frac{1}{1-\alpha}\sum_{i\in\mathcal{B}}\frac{\lambda_i(\tilde{x}_i)}{|\lambda_i'(\tilde{x}_i)|}(\lambda_i(\tilde{x}_i)-\lambda_i(x_i^*)).$$

Now, from Lemma IV.5, we get that for all i either $\lambda_i(\tilde{x}_i)=\lambda_i(x_i^*)$ or $\frac{\lambda_i(\tilde{x}_i)}{|\lambda_i'(\tilde{x}_i)|}\leq \tilde{x}_i$. Therefore, we get $SW_2 \leq \frac{1}{1-\alpha}\sum_{i\in\mathcal{B}}\tilde{x}_i(\lambda_i(\tilde{x}_i)-\lambda_i(x_i^*))$.

We complete the claim using the following simple lemma.

Lemma IV.8. *We have that* $\sum_{i\in\mathcal{B}}\tilde{x}_i(\lambda_i(\tilde{x}_i)-\lambda_i(x_i^*)) \leq \sum_{t\in\mathcal{T}}[\tilde{p}_t\tilde{y}_t - C_t(\tilde{y}_t)] = \pi(\vec{\tilde{p}},\vec{\tilde{x}},\vec{\tilde{y}})$.

Proof. Consider any buyer i for whom $\lambda_i(\tilde{x}_i)>\lambda_i(x_i^*)$, and hence $\tilde{x}_i < x_i^*$. We claim that for such a buyer $\lambda_i(\tilde{x}_i)=\bar{P}(1-\alpha)^{\frac{1}{\alpha}} := P_{TH}$. Assume by contradiction that $\lambda_i(\tilde{x}_i) > P_{TH}$ (the inequality cannot be in the other direction since $\tilde{p}_t \geq P_{TH}$ for all $t\in\mathcal{T}$).

Since $\tilde{x}_i < x_i^*$, there must exist at least one slot t such that $\tilde{x}_{it} < x_{it}^*$ and hence $\tilde{x}_{it} < \ell_{it}$. Applying condition (4) of Proposition III.1 to this solution and slot t, we have that $\tilde{p}_t \geq \lambda_i(\tilde{x}_i) > P_{TH}$. By definition of our algorithm, this can only mean that $\tilde{p}_t = p_t^*$.

Next, from condition (1) of Proposition III.1 applied to this slot with respect to the social welfare maximizing allocation, we have that $p_t^* \leq \lambda_i(x_i^*) < \lambda_i(\tilde{x}) \leq \tilde{p}_t = p_t^*$, which of course, is a contradiction. Therefore, $\lambda_i(\tilde{x}_i) = P_{TH}$.

Since the price on every slot t, $\tilde{p}_t$ is at least P_{TH}, and $\lambda_i(\tilde{x}_i) = P_{TH}$, we can conclude that this buyer purchases exclusively from slots whose price is P_{TH}. Therefore, let B^+ be the set of buyers for whom $\lambda_i(\tilde{x}_i)>\lambda_i(x_i^*)$. Then,

$$\sum_{i\in B^+}\lambda_i(\tilde{x}_i)\tilde{x}_i = \sum_{i\in B^+}\sum_{t\in B_i}\tilde{x}_{it}P_{TH} = \sum_{i\in B^+}\sum_{t\in B_i}\tilde{x}_{it}\tilde{p}_t.$$

Now, for the second part of the expression in the lemma. For every buyer i and slot t with $\tilde{x}_{it}>0$, we have (from Proposition III.1) that $\lambda_i(\tilde{x}_i)\geq \tilde{p}_t \geq p_t^* = c_t(y_t^*) \geq c_t(\tilde{y}_t)$. The last inequality comes from the crucial Claim IV.2. Therefore,

$$\sum_{i\in B^+}\lambda_i(x_i^*)\tilde{x}_i = \sum_{i\in B^+}\sum_{t\in B_i}\lambda_i(x_i^*)\tilde{x}_{it} \geq \sum_{i\in B^+}\sum_{t\in B_i}c_t(\tilde{y}_t)\tilde{x}_{it}.$$

Subtracting the two halves and adding the non-negative entity $\sum_{i\in\mathcal{B}\setminus B^+}\sum_{t\in B_i}\tilde{x}_{it}[\tilde{p}_t - c_t(\tilde{y}_t)]$, we get

$$\sum_{i\in\mathcal{B}}\sum_{t\in B_i}\tilde{x}_{it}[\tilde{p}_t - c_t(\tilde{y}_t)] = \sum_{t\in\mathcal{T}}[\tilde{p}_t\tilde{y}_t - \tilde{y}_t c_t(\tilde{y}_t)] \leq \sum_{t\in\mathcal{T}}\tilde{p}_t\tilde{y}_t - C_t(\vec{\tilde{y}}).$$

□

□

□

A. Bicriteria Approximation

Previously, we considered the NP-Hard problem of profit maximization and presented an algorithm whose approximation degrades gracefully as α increases. However, as argued previously, in such large markets with repeated engagement, it is important to characterize the efficiency of the proposed (approximately) profit maximizing solution. The main question that we consider here is the following: *does the proposed profit maximizing solution simultaneously result in good welfare?* The next theorem answers this question in the affirmative.

Theorem 2. *The social welfare of the solution returned by our algorithm is at most a factor* $1+\frac{1}{1-\alpha}$ *smaller than that of the optimum social welfare.*

Proof. Recall from Claim IV.7 that $\overline{SW}(\vec{p^*},\vec{x^*},\vec{y^*}) - \overline{SW}(\vec{\tilde{p}},\vec{\tilde{x}},\vec{\tilde{y}}) \leq \frac{1}{1-\alpha}\pi(\vec{\tilde{p}},\vec{\tilde{x}},\vec{\tilde{y}})$. However, note that $\overline{SW}(\vec{p^*},\vec{x^*},\vec{y^*}) - \overline{SW}(\vec{\tilde{p}},\vec{\tilde{x}},\vec{\tilde{y}}) = SW(\vec{p^*},\vec{x^*},\vec{y^*}) - SW(\vec{\tilde{p}},\vec{\tilde{x}},\vec{\tilde{y}})$. To conclude, we have that $SW(\vec{p^*},\vec{x^*},\vec{y^*}) \leq SW(\vec{\tilde{p}},\vec{\tilde{x}},\vec{\tilde{y}}) + \frac{1}{1-\alpha}\pi(\vec{\tilde{p}},\vec{\tilde{x}},\vec{\tilde{y}}) \leq SW(\vec{\tilde{p}},\vec{\tilde{x}},\vec{\tilde{y}}) + \frac{1}{1-\alpha}SW(\vec{\tilde{p}},\vec{\tilde{x}},\vec{\tilde{y}})$. □

Bounding the inefficiency at equilibrium

Although we presented an approximation algorithm with good profit and welfare properties, it is possible that the seller might employ other pricing strategies, particularly the profit-maximizing solution $(\vec{p}^{opt},\vec{x}^{opt},\vec{y}^{opt})$. Does this self-interested behavior lead to good system conditions? In this section, we investigate this question by presenting a bound on the following quantity, the quality of equilibrium solution compared to that of the system optimum.

Theorem IV.9. *The ratio of the social welfare of the welfare-maximizing allocation to that of the profit-maximizing pricing solution is at most* $2(\frac{1}{1-\alpha})^{\frac{1}{\alpha}}+\frac{\alpha}{1-\alpha}$, *i.e.,*

$$\frac{SW(\vec{p^*},\vec{x^*},\vec{y^*})}{SW(\vec{p}^{opt},\vec{x}^{opt},\vec{y}^{opt})} \leq \left(2(\frac{1}{1-\alpha})^{\frac{1}{\alpha}}+\frac{\alpha}{1-\alpha}\right).$$

Proof. Recall that for any pricing solution $(\vec{p},\vec{x},\vec{y})$, $\overline{SW}(\vec{p},\vec{x},\vec{y}) = SW(\vec{p},\vec{x},\vec{y}) - \kappa$, where $\kappa = \sum_{i\in\mathcal{B}}u_i(\bar{x}_i) - \bar{P}\bar{x}_i$ is a constant that is independent of the pricing solution (Lemma III.3). We begin by providing an upper bound for the desired ratio in terms of the reduced social welfare $\overline{SW}$:

$$\begin{aligned}\frac{SW(\vec{p^*},\vec{x^*},\vec{y^*})}{SW(\vec{p}^{opt},\vec{x}^{opt},\vec{y}^{opt})} &\leq \frac{SW(\vec{p^*},\vec{x^*},\vec{y^*})-\kappa}{SW(\vec{p}^{opt},\vec{x}^{opt},\vec{y}^{opt})-\kappa}\\ &\leq \frac{\overline{SW}(\vec{p^*},\vec{x^*},\vec{y^*})}{\overline{SW}(\vec{p}^{opt},\vec{x}^{opt},\vec{y}^{opt})}\end{aligned}$$

The first inequality comes from the simple fact that for any $a\geq b$ and $c<b$, $\frac{a}{b}\leq\frac{a-c}{b-c}$. In our case, κ is clearly smaller than $SW(\vec{p}^{opt},\vec{x}^{opt},\vec{y}^{opt})$. Of course, if we look at the reduced instance, we also get that $\overline{SW}(\vec{p}^{opt},\vec{x}^{opt},\vec{y}^{opt}) \leq \pi(\vec{\tilde{p}},\vec{\tilde{x}})$. Applying this, we have

$$\frac{\overline{SW}(\vec{p^*},\vec{x^*},\vec{y^*})}{\overline{SW}(\vec{p}^{opt},\vec{x}^{opt},\vec{y}^{opt})} \leq \frac{\overline{SW}(\vec{p^*},\vec{x^*},\vec{y^*})}{\pi(\vec{\tilde{p}},\vec{\tilde{x}},\vec{\tilde{y}})}$$

The required ratio follows from the statement of Theorem 1. □

V. Simulation Results

In this section, we test the efficacy of our pricing algorithms (in terms of profit and social welfare) via numerical experiments performed on simulated PEV charging markets with realistic parameters. Our goal is to understand the dependence of these objectives on the various market parameters, e.g., density of market, convexity of utility function, etc. Our results illustrate that the pricing algorithm presented in Section IV consistently and significantly outperforms the Walrasian prices in terms of profit while at the same time, having near-optimal social welfare. That is, the proposed technique is beneficial for both the seller and the system as a whole.

Setup: We consider a full day of 24 time slots, with the number of PEVs (in the distribution network under consideration) ranging from $N = 50$ to $N = 350$. Each PEV is constrained to charge over 6 time slots, which are distributed according to a Gaussian distribution centered around peak demand hours. The supply (energy procurement) cost that the aggregator incurs at time slot $t \in \mathcal{T}$ is given by $C_t(y) = a(D_t + y^2)$, where D_t is a parameter that reflects the external load on the supplier, estimated from actual demand data [24] with a being a suitable scaling parameter. The experimental results are quantified in terms of the following three performance measures, normalized by the optimum social welfare:

1) *Profit Guaranteed by Our Algorithm*: $\frac{SW(\vec{p^*},\vec{x^*},\vec{y^*})}{\pi(\vec{\tilde{p}}),\vec{\tilde{x}},\vec{\tilde{y}})}$.
2) *Social Welfare Guaranteed by Our Algorithm*: $\frac{SW(\vec{p^*},\vec{x^*},\vec{y^*})}{SW(\vec{\tilde{p}}),\vec{\tilde{x}},\vec{\tilde{y}})}$.
3) *Profit Guaranteed by the Walrasian Solution*: $\frac{SW(\vec{p^*},\vec{x^*},\vec{y^*})}{\pi(\vec{p^*},\vec{x^*},\vec{y^*})}$. While the Walrasian prices maximize social welfare, the profit guaranteed by such a solution can often be significantly sub-optimal.

Note that the optimum social welfare is larger than all of the above parameters, and therefore, the ratio of $SW(\vec{p^*},\vec{x^*},\vec{y^*})$ to any of the above quantities is at least one.

Effect of Increased Demand: First, we set out to understand how the social utility and profit resulting from our algorithm as well as the Walrasian profit vary as we increase the number of PEVs from $N = 50$ to $N = 350$. Our results indicate that for both polynomial ($\alpha = 0$, log-concave demand) and exponential utilities ($\alpha = 0.5$), the profit guaranteed by our algorithm is significantly better than the theoretical bound. The profit provided by the Walrasian solution improves as N increases since increased congestion leads higher marginal costs and hence Walrasian prices (Proposition III.2). However, the Walrasian profit is considerably smaller than that of our algorithm for both log-concave (factor of 2) and exponential (factor of 1.5) utilities. Simultaneously, the social welfare guaranteed by our algorithm is almost optimal (factor of 1.3 or smaller compared to the optimal social welfare).

For example, our theoretical result guarantees that our algorithm's profit is within a factor of $2e \approx 5.54$ from the optimum welfare for $\alpha = 0$ and factor of 9 for the $\alpha = 0.5$ case. Instead, and surprisingly, the profit is always within a factor of e or better from the optimum welfare.

Characterizing the effect of the Utility Concavity: Keeping the number of PEVs fixed at $N = 300$, we study how changing the utility function (whose volatility is given by α) affects the profit and social welfare. While our theoretical bounds predict poor profit as α becomes larger, our experiments reveal that even for sufficiently large values of α (for e.g., $\alpha = 0.8$), the performance is still optimal. Not surprisingly, at $\alpha = 1$, our algorithm's performance matches that the optimum social welfare: by definition, the prices computed by our algorithm approach the Walrasian prices when α becomes closer to 1 where the Walrasian prices are sufficiently large.

VI. Conclusion Remarks

References

[1] S.Deilami, A. S. Masoum, P. S. Moses, and M. A. S. Masoum, "Voltage profile and THD distortion of residential network with high penetration of plug-in electrical vehicles," *Proc. of IEEE PES Innov. Smart Grid Tech. Conf.*, October 2010.
[2] L. Fernandez, T. G. S. Roman, R. Cossent, C. M. Domingo, and P.Frias, "Assessment of the impact of plug-in electric vehicles on distribution networks," *IEEE Trans. on Power Systems*, vol. 26, February 2011.
[3] A. S. Masoum, S. Deilami, P. S. Moses, and A. Abu-Siada, "Impact of plug-in electrical vehicles on voltage profile and losses of residential system," *Proc. of AUPEC 2010*, December 2010.
[4] S. Babaei, D. Steen, L. A. Tuan, O. Carlson, and L. Bertling, "Effects of plug-in electric vehicles on distribution systems: A real case of Gothenburg," *Proc. of IEEE ISGT 2010*, October 2010.
[5] J. Waddell, M. Rylander, A. Maitra, and J. A. Taylor, "Impact of plug in electric vehicles on Manitoba Hydro's distribution system," *Proc. of IEEE Electrical Power and Energy Conf.*, October 2011.
[6] N. Melo, F. Mira, A. de Almeida, and J. Delgado, "Integration of PEV in Portuguese distribution grid analysis of harmonic current emissions in charging points," *Proc. of Intl. Conf. on Electrical Power Quality and Utilisation*, October 2011.
[7] M. D. Galus and G. Andersson, "Demand management of grid connected plug-in hybrid electric vehicles (phev)," in *2008 IEEE Energy 2030 Conference*, pp. 1–8, Nov 2008.
[8] Z. Ma, D. S. Callaway, and I. A. Hiskens, "Decentralized charging control of large populations of plug-in electric vehicles," *IEEE Trans. on Control System Technology*, vol. 21, pp. 67–78, January 2013.
[9] W. Tushar, W. Saad, V. Poor, and D. Smith, "Economics of electric vehicle charging: a game theoretic approach," *IEEE Trans. Smart Grid*, vol. 3, pp. 1767–1778, December 2012.
[10] O. Beaude, S. Lasaulce, and M. Hennebel, "Charging games in networks of electrical vehicles," in *Proc. of NetGCoop 2014*, pp. 96–103, Nov 2012.
[11] A. Ghavami, K. Kar, S. Bhattacharya, and A. Gupta, "Price-driven charging of plug-in electric vehicles: Nash equilibrium and best-response convergence analysis," *Proc. of CISS 2013*, March 2013.
[12] L. Gan, U. Topcu, and S. H. Low, "Optimal decentralized protocol for electric vehicle charging," *IEEE Trans. Power Systems*, vol. 28, pp. 940–951, May 2013.
[13] K. Clement-Nyns, E. Haesen, and J. Driesen, "The impact of charging plug-in hybrid electric vehicles on a residential distribution grid," *IEEE Trans. Power Systems*, vol. 25, pp. 371–380, Feb 2010.
[14] N. Rotering and M. Ilic, "Optimal charge control of plug-in hybrid electric vehicles in deregulated electricity markets," *IEEE Trans. Power Systems*, vol. 26, pp. 1021–1029, Aug 2011.
[15] I. S. Bayram, A. Tajer, M. Abdallah, and K. Qaraqe, "Capacity planning frameworks for electric vehicle charging stations with multiclass customers," *IEEE Trans. Smart Grid*, vol. 6, pp. 1934–1943, July 2015.
[16] Z. Fan, "A distributed demand response algorithm and its application to phev charging in smart grids," *IEEE Trans. Smart Grid*, vol. 3, pp. 1280–1290, Sept 2012.
[17] N. Li, L. Chen, and S. H. Low, "Optimal demand response based on utility maximization in power networks," in *2011 IEEE Power and Energy Society General Meeting*, pp. 1–8, July 2011.

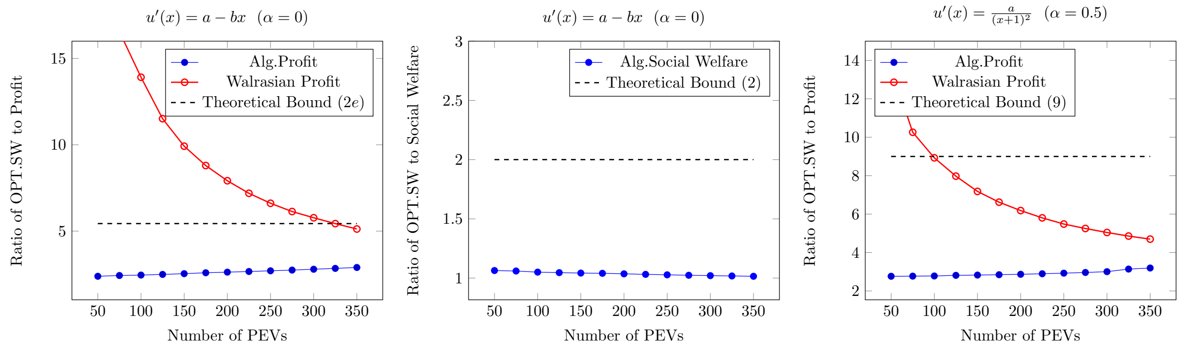

Figure 1. Ratio of the optimum social utility to the profit and welfare obtained by our algorithm as a function of N in comparison to the respective theoretical bounds and the profit obtained by the Walrasian prices for utility functions corresponding to $\alpha = 0$ and $\alpha = 0.5$. For both profit and welfare, smaller values are better since a ratio of 1 implies optimality.

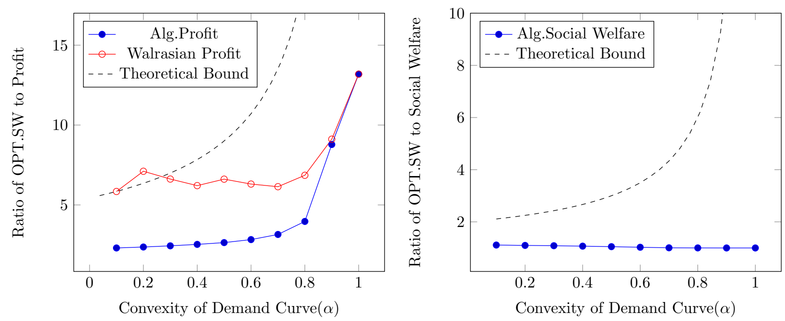

Figure 2. Profit and Social Welfare of our algorithm as a function of α (the convexity of λ) in comparison to the theoretical bounds and the Walrasian profit. For both profit and welfare, smaller values are better since a ratio of 1 implies optimality.

[18] A. Ghavami, K. Kar, S. Bhattacharya, and A. Gupta, "Price-driven charging of plug-in electric vehicles: Nash equilibrium, social optimality and best-response convergence," in *Proc. of CISS 2013*, March 2013.

[19] S. Bhattacharya, K. Kar, J. H. Chow, and A. Gupta, "Extended second price auctions with elastic supply for pev charging in the smart grid," *IEEE Trans. Smart Grid*, vol. 7, pp. 2082–2093, July 2016.

[20] A. Ghavami and K. Kar, "Nonlinear pricing for social optimality of pev charging under uncertain user preferences," *Proc. of 48th Annual Conf. on Information Sciences and Systems (CISS)*, February 2014.

[21] A. Ghosh and V. Aggarwal, "Control of charging of electric vehicles through menu-based pricing," *IEEE Trans. Smart Grid*, 2017. To appear.

[22] E. Anshelevich, K. Kar, and S. Sekar, "Pricing to maximize revenue and welfare simultaneously in large markets," in *Proc. 12th Int. Conf. Web Internet Econ.*, pp. 145–159, 2016.

[23] R. Cole and T. Roughgarden, "The sample complexity of revenue maximization," in *Proc. 46th Symp. Theory Comp.*, pp. 243–252, 2014.

[24] S. Shao, T. Zhang, M. Pipattanasomporn, and S. Rahman, "Impact of tou rates on distribution load shapes in a smart grid with phev penetration," in *Proc. IEEE T&D PES 2010*, pp. 1–6, 2010.

[25] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry, "On profit-maximizing envy-free pricing," in *Proc. Sixteenth Annual ACM-SIAM Symposium Discrete Algorithms*, pp. 1164–1173, 2005.

Appendix

Proposition A.1. *The problem of computing per-slot prices in order to maximize the seller's profit is NP-Hard even when all production cost functions are uniformly zero, and there are no charging constraints (i.e., $\ell_{it} = \infty$ for all $i \in \mathcal{B}$, $t \in \mathcal{T}$).*

Proof. We just sketch the proof here since the general idea is the same as the hardness proof in [25]. Consider an instance of vertex cover specified by a graph $G' = (V', E')$ with $|V'| = n, |E'| = m$. Reducing this to our problem, there is one time slot for each vertex in V', i.e., the set of time slots is given by $(t_i)_{i \in V'}$. Next, there are a total of $n + nm$ vertices that are described as follows:

1) $\mathcal{B}_1$: There is one PEV v_i for each $i \in V$. The PEV v_i only has an edge to time slot t_i. All of the PEVs in $\mathcal{B}_1$ have a utility function that is appropriately defined so that $\lambda(x) = u'(x) = \max(2 - x^d, 0)$ for some sufficiently large $d > 1$. It is not hard to deduce that as long as the exponent d is sufficiently large, this function can be interpreted as $\lambda(x) \approx 2$ for $x \leq 1$ and $\lambda(x) = 0$ for $x \geq 1$. In our proof, we will treat it as such.
2) $\mathcal{B}_2$: For every edge $(i, j) \in E'$, there are n PEVs in $\mathcal{B}_2$, $(w_{ij}^{(r)})_{(i,j) \in E', 1 \leq r \leq n}$. Moreover, every PEV $w_{ij}^{(r)}$ consists of edges to the time slots t_i and t_j, i.e., the vertices that form its endpoints. For every PEV in $\mathcal{B}_2$, we can define their utility functions appropriately such that $\lambda(x) = \max(2 - x, 0)$.

Since there is one unique time slot for each vertex in V', we use the terms time slots and vertices interchangeably. Our first claim is that in the optimum (profit-maximizing) pricing policy $(\vec{p}^{opt})$, the set of time slots whose price is strictly smaller than $2 - \epsilon$ (for some sufficiently small $\epsilon > 0$) must form a vertex cover in G'. Indeed, assume that this is not the case and let $(i,j) \in E'$ such that $p_{t_i}^{opt}, p_{t_j}^{opt} \geq 2-\epsilon$. Then, the contribution of the n PEVs (from the set $\mathcal{B}_2$) in the set $(w_{ij}^{(r)})_{1\leq r\leq n}$ to the optimum profit is close to zero. Therefore, if we reduce the price of (say) time slot t_i to $p_{t_i}=1$, the profit due to each of the n PEVs in the set $(w_{ij}^{(r)})_{1\leq r\leq n}$ becomes $n \times ((2-1)\times 1) = n$. Moreover, there is at most one node connected to t_i belonging to $\mathcal{B}_1$, and the profit due to this node reduces at most from $2-\epsilon$ to $1-\epsilon$. Therefore, the net profit strictly increases, which contradicts the optimality of $(\vec{p}^{opt})$.

In a similar manner, we can argue that the set of time slots whose price is strictly smaller than $2-\epsilon$ (call this set $\mathcal{T}_1$) is a minimum vertex cover in G'. If this were not a minimum vertex cover, we could obtain a contradiction by identifying an alternative pricing policy that leads to increased profit: let $VC \subseteq V'$ be a minimum vertex cover in G'. Then, we can simply price every time slot in VC at $p=1$ and price all of the other time slots at $p=2-\epsilon$, and show that this leads to a larger profit. Therefore, for any given instance of the vertex cover problem, we can form an instance of our PEV pricing problem such that the optimal solution to the latter could be used to efficiently identify a minimum vertex cover. So, the NP-Hardness claim extends to our problem. □

Proposition A.2. *For any given $\alpha \in [0,1]$, there exists an instance with $\alpha - SR$ demand functions where the profit maximizing solution has a social welfare that is at least a factor of $(1+\frac{\alpha}{1-\alpha})^{\frac{1}{\alpha}}$ smaller than the optimum social welfare.*

Lemmas A.1 and A.2

The following claims are borrowed from [22]

Lemma A.3. *Let $f(x)$ be a non-increasing, non-negative α-Strongly regular function and x_1 be some point satisfying $f(0) \geq (\frac{1}{1-\alpha})^{\frac{1}{\alpha}} f(x_1)$. Then, $\frac{f(x_1)}{|f'(x_1)|} \leq x_1$.*

Lemma A.4. *Let $f(x)$ be any non-increasing, non-negative α-Strongly regular function. Then for any $x_2 \geq x_1$, the following inequality is true,*

$$\int_{x_1}^{x_2} f(x)dx \leq \frac{1}{1-\alpha}\frac{f(x_1)}{|f'(x_1)|}(f(x_1)-f(x_2)).$$