Journal of Artificial Intelligence Research 50 (2014) 885-922 Submitted 4/14; published 8/14

Demand Side Energy Management via Multiagent Coordination in Consumer Cooperatives

Andreas Veit ANDREAS@CS.CORNELL.EDU
Department of Computer Science
Cornell University
Ithaca, NY 14853 USA

Ying Xu YINGX1@ANDREW.CMU.EDU
Ronghuo Zheng RONGHUOZ@ANDREW.CMU.EDU
Tepper School of Business
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213 USA

Nilanjan Chakraborty NILANJAN.CHAKRABORTY@STONYBROOK.EDU
Department of Mechanical Engineering
Stony Brook University
Stony Brook, NY 11794 USA

Katia Sycara KATIA@CS.CMU.EDU
Robotics Institute, School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213 USA

Abstract

A key challenge in creating a sustainable and energy-efficient society is to make consumer demand adaptive to the supply of energy, especially to the renewable supply. In this article, we propose a partially-centralized organization of consumers (or agents), namely, a consumer cooperative that purchases electricity from the market. In the cooperative, a central coordinator buys the electricity for the whole group. The technical challenge is that consumers make their own demand decisions, based on their private demand constraints and preferences, which they do not share with the coordinator or other agents. We propose a novel multiagent coordination algorithm, to shape the energy demand of the cooperative. To coordinate individual consumers under incomplete information, the coordinator determines virtual price signals that it sends to the consumers to induce them to shift their demands when required. We prove that this algorithm converges to the central optimal solution and minimizes the electric energy cost of the cooperative. Additionally, we present results on the time complexity of the iterative algorithm and its implications for agents' incentive compatibility. Furthermore, we perform simulations based on real world consumption data to (a) characterize the convergence properties of our algorithm and (b) understand the effect of differing demand characteristics of participants as well as of different price functions on the cost reduction. The results show that the convergence time scales linearly with the agent population size and length of the optimization horizon. Finally, we observe that as participants' flexibility of shifting their demands increases, cost reduction increases and that the cost reduction is not sensitive to variation in consumption patterns of the consumers.

1. Introduction

Two key issues in creating a sustainable and energy-efficient society are to increase the penetration of renewable sources, and to manage supply and demand so as to reduce demand peaks while maintaining supply and demand balance. One way, which is most commonly used, to achieve the demand supply balance is to supply all the requested demand whenever it occurs. However, attempting to achieve demand supply balance by adjusting only the supply side leads to the use of flexible (usually diesel operated) power plants that can be expensive, inefficient, and emit large amount of carbon. An alternative to adjusting the supply side only, is to also adjust the demand of the consumers (Palensky & Dietrich, 2011) via demand response programs. Demand Response is defined as "the changes in electricity consumption by end users from their normal consumption patterns in response to changes in the price of electricity over time" (Albadi & El-Saadany, 2007).

Several different forms of demand response programs have been developed (for an overview see Albadi & El-Saadany, 2007). A typical example of an incentive based program, where customers receive payment for their participation, are Direct Load Control programs, where utilities remotely control the power consumption of consumers' appliances by switching them on/off. In small scale pilot studies, direct load control has been successful in reducing peak energy consumption, however consumers were uncomfortable with yielding control of their appliances to utility companies (Rahimi & Ipakchi, 2010; Medina, Muller, & Roytelman, 2010). Another type of demand management program is price based, where energy rates are variable and follow the real cost of electricity. The objective of this indirect method is to control the overall demand by incentivizing consumers to flatten the demand curve through shifting energy from peak to off-peak times. A typical example of these programs is Time of Use pricing, where the price during peak times is higher than the price during off-peak times. Recent technological advances in smart meters and smart appliances have created the potential to enable direct and real time participation of individual consumers in the energy market and thus make real-time price based demand management programs a reality. However, there are two key problems in realizing this potential. First, despite the presence of small pilot programs, utilities consider individual consumers of insufficient size to be considered for demand response services. Second, if consumers participate in the market directly, rather than through the utilities, the stability of the system may be compromised (e.g., herding) (Ramchurn, Vytelingum, Rogers, & Jennings, 2012). Considering these challenges, Mohsenian-Rad, Wong, Jatskevich, Schober, and Leon-Garcia (2010) argue that a good demand side management program should focus on controlling the *aggregate demand* (which is also important for economic load dispatching, Wood & Wollenberg, 1996) of a group of consumers instead of individual consumers. In this paper, the problem of coordinating a group of consumers called *consumer cooperative* is introduced and studied.

A consumer cooperative, or collective, allows *partial centralization* of consumers represented by a group coordinator (mediator) agent, who purchases electricity from utilities or the market on their behalf. Such consumer configurations can potentially increase energy efficiency via aggregation of demand to reduce peak power demand. The coordinator is neither a market maker nor a traditional demand response aggregator (Jellings & Chamberlin, 1993), since it does not set energy prices or aims to incur profits by selling to the market. Rather, its role is akin to a social planner's, in the sense that it manages the demand of its associated consumer group so that (a) the electricity cost to the group is minimized, and (b) individual group members autonomously and in a decentralized manner decide how to shift demands, while maintaining privacy of individual demand preferences and constraints. The members of such a cooperative are typically geographically co-located in close proximity to one another, for example a small neighborhood of households and/or enterprises. In practice, close proximity is required due to limitations in the distribution infrastructure, geographically separated system operators, markets, and electricity suppliers. The proposed coordination algorithm could also benefit already existing structures such as universities, malls, industrial parks, commercial estates, and large residential complexes. Although such organizations may already purchase electricity in a centralized manner, their constituent members (e.g. firms in an industrial park) are not currently coordinated so as to keep their constraints private.

Consumer cooperatives offer advantages to both energy utility companies and consumers. From a utility company's perspective, the consumer groups are large enough to be useful in demand response programs and have more predictable demand shifts compared to individual consumers. Barbose, Goldman, and Neenan (2004) give a survey of utility experience with price based programs and conclude that most participants were large industrial customers. For individual consumers, participation in such energy groups allows them to retain control of their own appliances. In addition, the consumers can obtain electricity at better prices than they would have, if they had purchased electricity individually. The price advantage is due to three reasons: First, the group's size allows the group to enter into more flexible purchase contracts, so that the price paid by the consumers reflects the actual cost of production more accurately. This is not the case in current long term fixed contract structures (Kirschen, 2003). Second, by buying collectively, the group can benefit from volume discounts analogous to group insurance programs. Third, in negotiated electricity contracts, the price usually consists of two components, one reflecting the actual energy production cost and the other as a premium against volatility in the energy demand and/or supply. Buying as a group can help reduce the premium against volatility, provided that the demands of the group members are coordinated, making their total demand more stable.

Our aim in this paper is to design effective schemes for coordinating the electricity demand of agents who are purchasing electricity as a consumer cooperative. The technical challenge of this endeavor is the fact that *the central coordinator does not know the constraints of the individual consumers*, and thus cannot compute the optimal demand schedule on its own. Furthermore, the actual cost of electricity consumption depends on the aggregate consumption profile of all agents. However, the agents may not want to share their demand patterns or constraints with other agents or the coordinator. Therefore, we present an algorithm designed to enable the central agent to coordinate the consumers to achieve the optimal centralized load, while the individual agents decide their demand shifting autonomously and retain their private knowledge about their demand constraints.[1]

The paper's contributions are as follows. First, we present an iterative coordination algorithm to minimize the energy cost of a consumer cooperative that preserves the privacy of individual demand constraints and costs of the consumers. Second, we prove that the algorithm converges to the centralized optimal solution and provide computational complexity results. Third, we provide formal arguments on the incentive compatibility of the coordination scheme. Fourth, we present and discuss extensive simulation results based on real world data. A preliminary version of this work appeared in the work of Veit, Xu, Zheng, Chakraborty, and Sycara (2013).

This paper is organized as follows: In Section 2 we give an overview of the related work and point out the differences to the approach in this paper. In Section 3 we formulate the cost optimization problem of the consumer cooperative. Then, in Section 4 we introduce the demand scheduling algorithms for the consumer cooperative. In particular, in Section 4.1 we introduce the basic iterative algorithm and in Section 4.2 we prove its convergence to the optimal solution. In Section 5, we introduce the general iterative algorithm and in Section 5.3 we prove its convergence in general

1. This problem has some surface resemblance to problems where a centralized coordinator determines resource allocations for agents with private preferences, which the coordinator tries to elicit. Such problems are typically addressed via Vickrey-Clark-Groves mechanisms. Our problem differs from those in fundamental ways that will be discussed in the Related Work section.

settings. In Section 5.4, we provide complexity analysis of the algorithm's convergence and in Section 5.5 we provide formal arguments on the incentive compatibility of the consumers. In Section 6 we evaluate the coordination algorithms using simulations based on real world consumption data. Finally, in Section 7, we summarize the main contributions of this paper and offer a perspective on future work.

2. Related Work

As mentioned in the introduction, the demand response programs vary from classical direct load control to price based programs with real time prices. In this paper, we introduce an algorithm that uses variable price signals to coordinate the energy consumption of a consumer cooperative. Therefore, we restrict this discussion to demand management using variable price signals.

Current literature on price-based demand shaping mostly operates under the assumption that it is desirable to have an automated or autonomous system (e.g. smart meter) that receives price signals and uses them to help a consumer schedule its demand so as to minimize its electricity cost, while satisfying its demand preferences and constraints. There are two types of price signals that can be used: dynamic and deterministic prices. Most literature on price-based demand response programs considers the case of deterministic prices, in which the electricity prices of all time slots are known before consumption. This case applies to all long-term contracts and day-ahead markets, if the planning horizon is sufficiently short (see Vytelingum, Ramchurn, Rogers, & Jennings, 2010; Ramchurn, Vytelingum, Rogers, & Jennings, 2011). With known electricity prices, consumers can compute their consumption schedule ahead of time, e.g., a daily consumption plan, or an annual production plan. Our paper falls in this category of demand scheduling under deterministic prices.

The approaches for demand scheduling proposed in the literature differ in some important characteristics. First, they differ in the *level at which the problem is studied* and secondly they differ in the *objective of the demand scheduling*. The different levels at which the problem has been studied include the level of the single consumer, the market maker, or the grid operator. The different objectives include minimizing the cost of a single consumer, minimizing the total cost of power generation, reducing the peak-to-average ratio in demand, and optimizing grid stability. There are only a few papers studying the demand scheduling problem at the level of grid operators, where the major concern is grid stability. In particular, objectives include the minimization of power flow fluctuations (Tanaka et al., 2011), and the minimization of power losses and voltage deviations (Clement-Nyns, Haesen, & Driesen, 2010).

Most work on demand scheduling is done at the level of the consumers. An important characteristic of this regime is that the electricity prices are often exogenously fixed and not influenced by the demand scheduling; i.e., the consumer is a price taker. Almost all papers in this regime focus on the (micro) demand scheduling of one or multiple appliances in a single residential household or a commercial building, with the typical objective to minimize the incurred electricity cost. For example, Chu and Jong (2008) study air-conditioning load control; Pedrasa, Spooner, and MacGill (2010) optimize the operation schedule of various distributed energy resources including space heater and pool pump, etc. There are other papers studying the problem from different perspectives. For example, Philpott and Pettersen (2006) and Samadi, Mohsenian-Rad, Wong, and Schober (2013) address the challenge of the uncertainty in the loads of the consumer's energy consumption. This stream of research is close to our work, since we as well study the demand management at the level of consumers. However, instead of studying a single consumer, we consider a group of consumers in a consumer cooperative, who buys the electricity from a utility company under a known price scheme. Our work differs from the above literature in the following aspects: size of shift, system stability, indirect control, and, most crucially, information privacy. First, while all the above literature only considers the control of a single consumer, in practice, it may be difficult for utility companies or grid operators to deal with individual consumers in such demand response programs. The demand shift of an individual consumer might be too small in magnitude compared to the aggregated necessary shift. Thus, it is unclear whether such a scheme will induce a shift of sufficient size. In contrast, from the utility's perspective, the consumer cooperative studied in our work is large enough to be useful in demand response programs. Second, most of the above literature assumes that the consumers can participate in the markets directly. However, there have been concerns voiced that, without control by the utility companies, the stability of the system may be compromised with such uncontrolled distributed interactions (Kirschen, 2003). In contrast, the demands of consumers in our consumer cooperative are coordinated by a central coordinator to minimize the electricity procurement cost of the cooperative. Buying as a group helps to reduce the demand volatility, and the coordinated demand management can achieve higher stability of demand and reduce demand peaks. Third, the common assumption in the above literature is that the demand manager has direct control over the appliances in the household and perfect knowledge about their loads and operation constraints. However, in our cooperative, the coordinator lacks such control on demand scheduling for individual consumers. Moreover, the demand constraints and preferences are private knowledge to the individual consumer and not known either to the coordinator or other consumers.

Another stream of research related to our work is the demand scheduling at the level of the market maker. In contrast to the previous approaches, in this regime, the coordinator (i.e., the market maker) can set the electricity prices for consumers, and often use the price lever to influence the demand of a number of consumers. The studies in this stream can be categorized into two groups: centralized and decentralized. Dietrich, Latorre, Olmos, and Ramos (2012) compare demand response programs in an electric system with high wind penetration under two different settings (i.e., centralized vs. decentralized) and show that the centralized approach often reaches higher overall cost savings, but has the disadvantage that central knowledge of consumers' constraints and preferences is necessary. Our partially decentralized approach avoids this limitation and keeps consumers constraints and preferences private while also achieving proven optimality of the solution.

When the coordinator uses prices to incentivize consumers to shift demand, the coordinator needs to pay close attention to the possibility of a *herding phenomenon*, whereby agents move their demand towards the low price times simultaneously and thus cause a spike in demand and bring instability to the system. To address this issue, in addition to a price signal, some papers adopt auxiliary methods to make agents gradually change their loads. Voice, Vytelingum, Ramchurn, Rogers, and Jennings (2011) charge agents an additional fee based on how much they change their demand profile from one period to the next. Ramchurn et al. (2011) introduce an adaptive mechanism controlling the rate and frequency at which the agents are allowed to adapt their loads and to readjust their demand profile. Vytelingum et al. (2010) introduce a compensation signal that is sent to the agents, providing an estimate of how much they should aim to change their behavior. These approaches may perform well, but, unlike our work, they do not provide any formal guarantees that their proposed algorithms will converge to the optimal solution. Moreover, these approaches require that the coordinator be a market maker who can charge consumers arbitrarily, e.g., by imposing additional fees, a feature that may not be implementable and acceptable in practice. In contrast, our coordinator is akin to a social planner who must ensure that the total charge for the energy demand of agents should equal the actual amount to be paid to the energy supplier.

Other papers address the decentralized demand-side management problem from a game theoretic perspective. They mainly focus on deriving a charging mechanism that can make all consumers reach a stable demand equilibrium, which achieves a central objective. These papers often share an underlying assumption that an individual consumer's unit electricity cost increases in other consumers aggregate consumption. This assumption holds under the setting that the energy cost in each period increases with the difference between demand and supply (Wu, Mohsenian-Rad, & Huang, 2012) or with the Peak-To-Average ratio (Mohsenian-Rad et al., 2010; Nguyen, Song, & Han, 2012), or the unit market price per period increases with the aggregate demand (Vytelingum, Voice, Ramchurn, Rogers, & Jennings, 2011; Atzeni, Ordonez, Scutari, Palomar, & Fonollosa, 2013). In the work of Mohsenian-Rad et al. (2010), a Nash equilibrium of demand achieves optimality via a designed billing strategy under which individual cost is proportional to the total energy cost. However, under this billing strategy, a consumer must have an estimate of other consumers demand before making its own decision. Therefore, to implement the proposed algorithms, the consumers needs to coordinate in an iterative manner and exchange their demand profiles with each other in each iteration, which may cause invasions of the consumers' privacy. In other works of Wu et al. (2012), Nguyen et al. (2012), Vytelingum et al. (2011) and Atzeni et al. (2013), individual consumers don't need to interact with each other, but only communicate with the central coordinator. In these papers, the coordinator maintains the equilibrium by designing the unit price, based on the aggregate demand. The algorithm is often implemented in a distributed and iterative scheme: in each iteration the agents report their loads to the coordinator and the coordinator updates the price accordingly. Although these papers ensure a minimum information exchange, none of them has theoretically proven that their algorithms would converge to the centralized optimal solution. In contrast, our paper not only safeguards the privacy of consumers, but also proves that the algorithm converges to the centralized optimal solution.

Our problem has surface resemblance to problems, where a central coordinator determines resource allocations for agents with private preferences, which the coordinator tries to elicit. Such problems are typically addressed via Vickrey-Clark-Groves mechanisms. The reason for not considering VCG-type mechanisms for our problem is that one of the key requirements of our application is that the coordinator has to pay the supplier for the group's electricity demand at the prices given by the supplier. Thus, the revenue, namely the sum of the payments individual agents make should equal the actual amount to be paid to the supplier. In other words, in the jargon of mechanism design, *budget balance* is a key requirement for our problem. It is well-documented in the literature that VCG-type mechanisms do not guarantee budget balance (Green & Laffont, 1977; Hurwicz, 1975)[2] and it is impossible to design a mechanism that achieves all the three properties, namely, efficiency, budget balance and strategy-proofness. For our problem, in addition to budget balance, *allocation efficiency* is also a desired social goal. Our algorithm achieves both budget balance (Lemma 2) and allocation efficiency (Theorem 3). Although the algorithm cannot guarantee strategy-proofness, due to the impossibility results, we prove that *no manipulation strategy exists that dominates truth reporting* (Theorem 5).

3. Problem Formulation

In our model, the consumer group consists of N members with the planning period divided into M discrete time slots. The number of discrete time slots depends on the market price structure, which can differ depending on the utility companies. For example, $M = 2$ for time-of-use pricing with different prices during day and night, whereas $M = 24$ for hourly time of use pricing. Let $\mathbf{R}$ be an $N \times M$ matrix where each row of the matrix, $\mathbf{r}_i$ is the electricity demand of the agent i,

2. In fact, as has been shown by Ausubel and Milgrom (2006), in the context of broadcast spectrum allocation, the revenue obtained from VCG-type mechanisms can even be zero! In our case, this would mean the agents would not pay any money to the coordinator, making it impossible to pay the supplier.

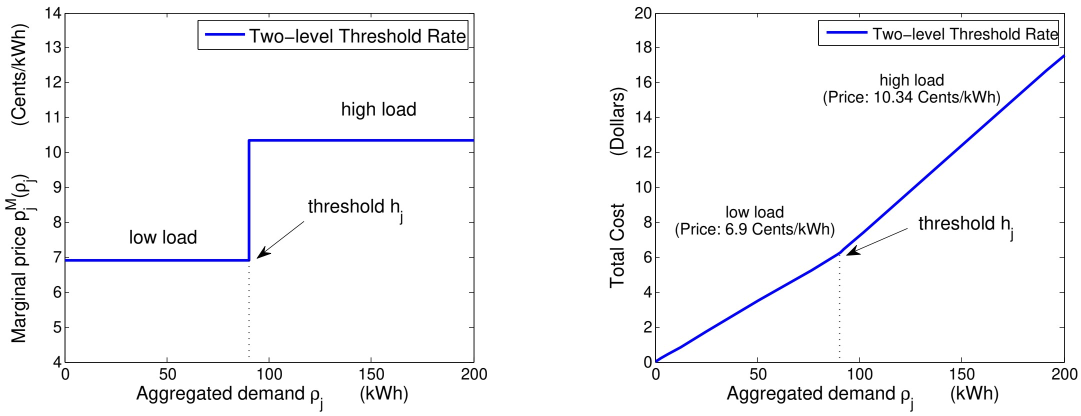

(a) Marginal electricity prices.

(b) Total electricity cost.

Figure 1: Sample two-level increasing threshold pricing model as used by BC Hydro

$i \in \{1, 2, \ldots, N\}$. We call $\mathbf{r}_i$ the *demand profile* of agent i. Each entry r_{ij} is the electricity demand of agent i for time slot j. The total aggregated demand in time slot j is $\rho_j = \sum_{i=1}^{N} r_{ij}$. The average market price of a unit of electricity for the consumer group at time slot j is defined as $\mathrm{p}_j(\rho_j)$.

We assume a typical market price function, where the prices are different in each time slot and the price has a threshold structure. This means that the marginal electricity prices differ among different demand levels. For each time slot, every unit of electricity consumed below a specified *price threshold* is charged at a lower price, while any additional unit exceeding that threshold is charged at a higher price. Thus, the marginal electricity price in a time slot, denoted by $\mathrm{p}_j^m(\rho_j)$, is a non-decreasing function of the total demand. The *marginal price* at a given demand level is the payment increment (decrement) for adding (reducing) one unit of electricity. Figure 1a shows an example of a two-level increasing threshold pricing model adopted from BC Hydro.[3] The marginal price of a two-level threshold structure can formally be written as follows: $\mathrm{p}_j^m(\rho_j) = \begin{cases} p_j^H & \rho_j > h_j \\ p_j^L & \rho_j \leq h_j \end{cases}$ with $p_j^H > p_j^L$, where h_j is the price threshold in time slot j. Let x^+ denote the positive value of a term x, i.e., $x^+ = \max\{0, x\}$, and x^- denote the negative value respectively, i.e., $x^- = \min\{0, x\}$. The total energy cost for time slot j is thus the integral of the marginal prices. Figure 1b shows the total electricity cost for the aggregated demand based on the two-level threshold pricing model. The total electricity cost can be computed as:

$$\mathrm{p}_j(\rho_j)\,\rho_j = p_j^H(\rho_j - h_j)^+ + p_j^L(\rho_j - h_j)^- + p_j^L h_j \tag{1}$$

The demand profile of each agent $\mathbf{r}_i$ must satisfy its individual constraints. The overall demand consists of two types of loads: shiftable loads and non-shiftable loads. Mohsenian-Rad et al. (2010), Mohsenian-Rad and Leon-Garcia (2010) and Wu et al. (2012) model the demand constraints of shiftable and non-shiftable loads. An example for non-shiftable loads is typically a refrigerator and

3. BC Hydro is a Canadian utility company. This pricing model is obtained from www.bchydro.com.

shiftable loads include dishwasher, electric vehicles, washer/dryer, etc. The appliances are modeled by a total demand that is required over the course of the planning horizon, upper and lower bounds on the demand in each time slot as well as earliest start and latest end times. These constraints form a convex set. For our problem, we also assume that the total demand of each agent during the whole planning period is fixed, i.e., $\sum_{j=1}^{M} r_{ij} = \tau_i$, where τ_i is the total demand for agent i. Further, we also consider loads, where the demand constraints form a convex set and denote it as $\mathcal{X}_i$. This constraint set is private knowledge and an agent does not share it, neither with other firms nor with the coordinator. In some application scenarios, when an agent determines its energy demand profile, it has to consider additional costs associated with the demand schedule. For example, in any given factory the energy is most commonly used for production. Changing the energy demand schedule, therefore, may mean changing the production process and, thereby, the production cost. For agent i, this cost is denoted by $\mathrm{g}_i(\mathbf{r}_i)$. We assume this cost function to be convex. The overall cost function of each agent is then $\sum_{j=1}^{M} \mathrm{p}_j\left(\rho_j\right) r_{ij} + \mathrm{g}_i(\mathbf{r}_i)$.

With the objective to minimize the sum of all agents' costs, the *overall energy allocation problem* can be written as:

$$\begin{aligned} \min \mathrm{C}\left(\mathbf{R}\right) := \; & \sum_{i=1}^{N} \sum_{j=1}^{M} \mathrm{p}_j(\rho_j) r_{ij} + \sum_{i=1}^{N} \mathrm{g}_i(\mathbf{r}_i) \\ \text{s.t. } & \mathbf{r}_i \in \mathcal{X}_i, \; \sum_{j=1}^{M} r_{ij} = \tau_i. \end{aligned} \tag{2}$$

where the energy allocations r_{ij} are the optimization variables and $\mathbf{R}$ is the matrix of the demand profiles of all the agents. Note that the above problem is defined on a convex set $\mathcal{X}_i$. Although the objective function is non-linear, it is convex because of the following. First, $\sum_{i=1}^{N} \sum_{j=1}^{M} \mathrm{p}_j\left(\rho_j\right) r_{ij} = \sum_{j=1}^{M} \mathrm{p}_j\left(\rho_j\right) \rho_j$ is convex and non-decreasing in ρ_j as indicated by Equation 1. Together with $\rho_j = \sum_{i=1}^{N} r_{ij}$, we can conclude that $\sum_{j=1}^{M} \mathrm{p}_j\left(\rho_j\right) \rho_j$ is convex in r_{ij}, $\forall i, j$, (Boyd & Vandenberghe, 2004). Since $\mathrm{g}_i(\mathbf{r}_i)$ is also convex, the total cost function $\mathrm{C}\left(\mathbf{R}\right)$ is a summation of convex functions and so also convex. Thus, Problem 2 is a convex minimization problem.

4. Solution Approach

Although the Problem 2 is a convex optimization problem, since the constraints and preferences of the agents are private knowledge, the optimal demand profiles cannot be computed directly by the central coordinator. The objective function, although a sum of the individual costs of each agent, is coupled, because the price of electricity in any time slot j depends on the aggregated demand of all agents ρ_j. However, since the constraints in Problem 2 are agent-specific, they are naturally separable. Therefore, a primal decomposition approach (Bertsekas & Tsitsiklis, 1989) is used to solve the problem in which the sub-problems correspond to each agent optimizing its own energy cost subject to its individual constraints. The central coordinator has to compute the appropriate information to be sent to the agents so as to guide the demand pattern towards time slots with lower prices (this corresponds to the master problem in primal decomposition methods).

Since the agents know the electricity market prices, they individually optimize their demand according to those prices. Let the resulting demand profile be called the *uncoordinated demand profile*. Figure 2a depicts such an aggregated uncoordinated demand profile in a setting with three time slots. It can be seen that the aggregated demand in time slot 2 is above the threshold, $\rho_2 > h_2$, and the aggregated demand in time slots 1 and 3 is below, $\rho_1 < h_1, \rho_3 < h_3$. Thus, a shift of demand from time slot 2 to the other time slots would reduce the total cost for the group. However, since the agents don't know the demand of the other agents, they cannot shift their demand. An intuitive solution approach, for the coordinator, to coordinate the demand would be to inform the

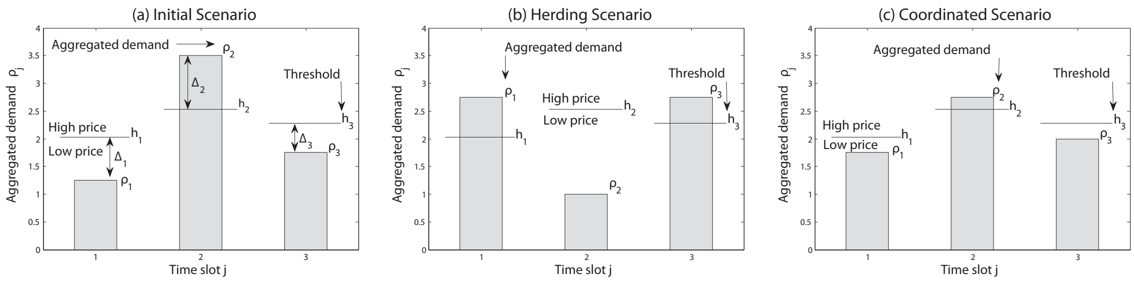

Figure 2: A comparison of demand profiles in Uncoordinated, Herding and Coordinated scenarios.

agents about the aggregated demand in each time slot. Knowing the market price, the agents could then solve their individual optimization problems. This approach is problematic, because the agents don't know the demand constraints and preferences of the other agents in the group, while their costs strongly depend on the demand of the other agents. For example, all agents knowing the market price and the current aggregated demand could shift as much demand as possible to a supposedly cheap time slot. This would lead to load synchronization, *herding phenomenon*, where all agents shift demand to the supposedly cheap time slot, resulting in a new demand peak in that time slot and thus increasing the total cost. The effect of this herding phenomenon is shown in Figure 2b, where too much demand was shifted from time slot 2 resulting in demand above the threshold in time slots 1 and 3. Thus, the key challenge is to design the information that the coordinator sends to the agents, the virtual price signal, to enable the system to minimize the overall cost.

A virtual price signal is not the final price the agents have to pay, but information about what they would have to pay, given the current aggregated demand. The virtual price signal enables the agents to foresee the possible price increment/reduction caused by their demand shifting. Therefore, the virtual price signal for agent i in time slot j, $\mathrm{s}_{ij}^{v}\left(r_{ij}|\mathbf{R}\right)$, is a function of the variable r_{ij}, denoting the new demand of agent i in time slot j. The price signal is computed based on the previous aggregate demand profile $\mathbf{R}$, which is therefore included in the price function. For ease of readability, time is not made explicit in this notation and will only be used in the proofs. The superscript v indicates the *virtual* price, in contrast to the real market prices $\mathrm{p}_j(\rho_j)$. To design the virtual price signal, the coordinator first computes the amount of demand that should be ideally shifted in each time slot. As shown in Figure (2a), this amount, denoted by Δ_j, $j = 1, 2, 3$, is the difference between the total aggregated demand and the price threshold in each time slot. For readability, we will refer to Δ_j as *delta increment*, noting that the delta could have negative values, i.e., it could be a decrement. To avoid herding, the delta increment needs to be divided among the agents and a threshold price signal needs to be designed for each agent, so that the price below the threshold is lower than the price above the threshold. This serves to penalize the total demand in a time slot going above the threshold. Thus, the agents know the maximum amount of demand they could shift at what prices and can solve their individual optimization problem. The exact calculation of the price signal $\mathrm{s}_{ij}^{v}\left(r_{ij}|\mathbf{R}\right)$ is shown in Section 4.1.2. Given the price signal, the *virtual cost optimization problem* each agent solves is

$$\begin{aligned} \min \mathrm{C}_i^v\left(\mathbf{r}_i|\mathbf{R}\right) := \ \min \ \textstyle\sum_{j=1}^{M} \mathrm{s}_{ij}^{v}\left(r_{ij}|\mathbf{R}\right) r_{ij} + \mathrm{g}_i(\mathbf{r}_i) \\ \text{s.t. } \mathbf{r}_i \in \mathcal{X}_i,\ \textstyle\sum_{j=1}^{M} r_{ij} = \tau_i. \end{aligned} \tag{3}$$

Note that this problem, like the overall problem, is a convex optimization problem and thus solvable. However, because of their individual constraints and cost functions, some agents might not be able to shift as much demand as was assigned to them by means of the virtual price signal. This implies that the aggregated demand shift can be less than the amount that could have been achieved. Figure 2c shows this case, where the total demand in the second time slot remains above the threshold, because not the whole Δ_2 could be shifted. In order to shift the remaining demand, another price signal, dividing the remaining amount, would be necessary. This motivates us to design an iterative algorithm for the coordinator to update the virtual price signal based on the consumers' feedback and thus gradually adjust the individual demands to the central optimal solution.

4.1 Basic Coordination Algorithm

We present the basic coordination algorithm and the details of the virtual price signal design for the energy allocation in the setting with a planning horizon of only two time slots, $M = 2$, and no cost for shifting demand, $\mathrm{g}(\cdot) = 0$. Although called basic setting in this paper, this setting has practical relevance, because it represents the commonly used Time of Use pricing schemes that divide the planning horizon in two time slots (Albadi & El-Saadany, 2007). One time slot with typically high load has high prices and one time slot with typically low load has low prices for electricity.

4.1.1 Overview of Algorithm

Recall that $\mathbf{r}_i$ denotes the demand profile of agent i and that $\mathbf{R}$ is the matrix of the demand profiles of all agents. Let $\mathbf{r}'_i$ be the updated demand profile of agent i after an iteration and $\mathbf{R}'$ be the new demand profile of all agents.

Initialization: Each agent computes an initial uncoordinated electricity demand profile $\mathbf{r}_i$ by solving Problem 3 based on the market prices and sends it to the coordinator.

1. The coordinator adds up the individual demands to determine the aggregated demand ρ_j for each time slot j and then calculates the delta increment Δ_j of demand to be shifted in or out of each time slot. Finally, the coordinator divides that demand among all agents and computes the virtual price signals $\mathrm{s}^v_{ij}(r_{ij}|\mathbf{R})$ for each agent i and time slot j.

2. The coordinator sends the virtual price signals to all agents.

3. After receiving the virtual price signal, each agent individually calculate its new demand profiles $\mathbf{r}'_i$ according to the optimization Problem 3.

4. The agents send their new demand profiles back to the coordinator.

5. The coordinator compares the new demand profiles to the old profiles. If no agent changed its demand profile, i.e., $\mathbf{R} = \mathbf{R}'$, the coordinator stops the algorithm. Otherwise, it sets $\mathbf{R} = \mathbf{R}'$ and goes to step (1).

4.1.2 Coordination with Virtual Price Signal

The virtual price signal for one agent in one time slot is a threshold price function of the demand in that time slot. The demand up to a specified threshold is charged at a low price and the demand above the threshold is charged at a higher price. The virtual price signal is therefore parameterized by the low marginal price, p^L_j, the high marginal price, p^H_j, and the price threshold, $\mathrm{h}_{ij}(\mathbf{R})$, that specifies the demand levels at which the prices apply. The virtual price, s^v_{ij}, for agent i in time slot j is computed based on these parameters as follows

$$\mathrm{s}^v_{ij}(r_{ij}|\mathbf{R}) = \begin{cases} \frac{p^L_j \, \mathrm{h}_{ij}(\mathbf{R}) + p^H_j (r_{ij} - \mathrm{h}_{ij}(\mathbf{R}))}{r_{ij}} & r_{ij} > \mathrm{h}_{ij}(\mathbf{R}) \\ p^L_j & r_{ij} \leq \mathrm{h}_{ij}(\mathbf{R}) \end{cases} \tag{4}$$

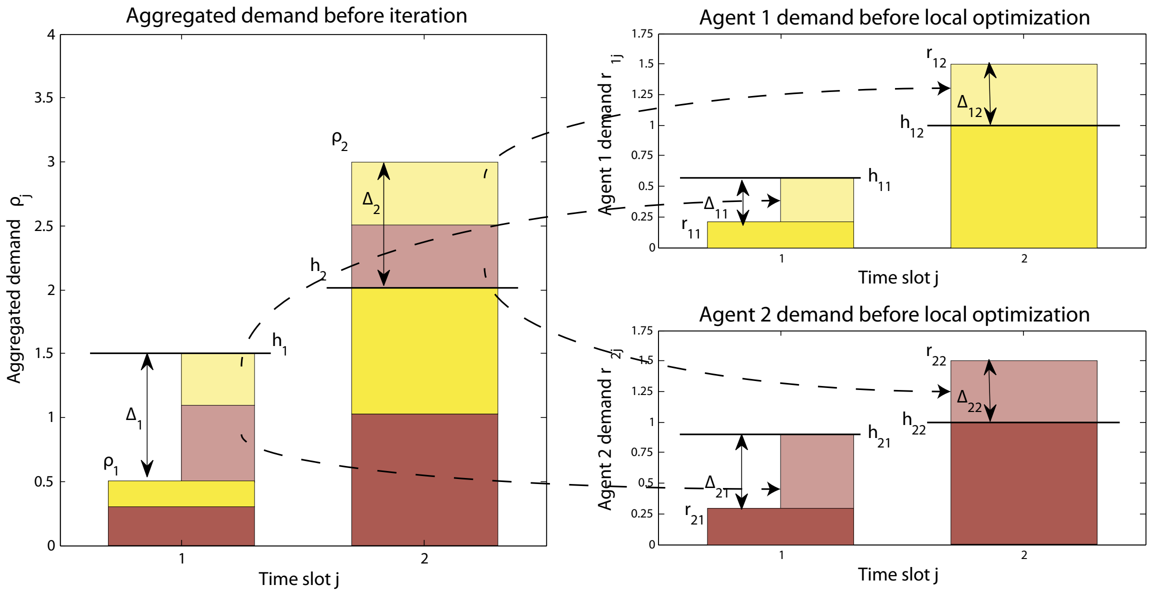

Figure 3: Division of Δ_j among agents for the virtual price signal. The yellow demand belongs to agent 1 and the red demand to agent 2. Here Δ_1 is the amount of demand that can be shifted into time slot 1, while Δ_2 indicates the amount of demand that needs to be shifted from time slot 2.

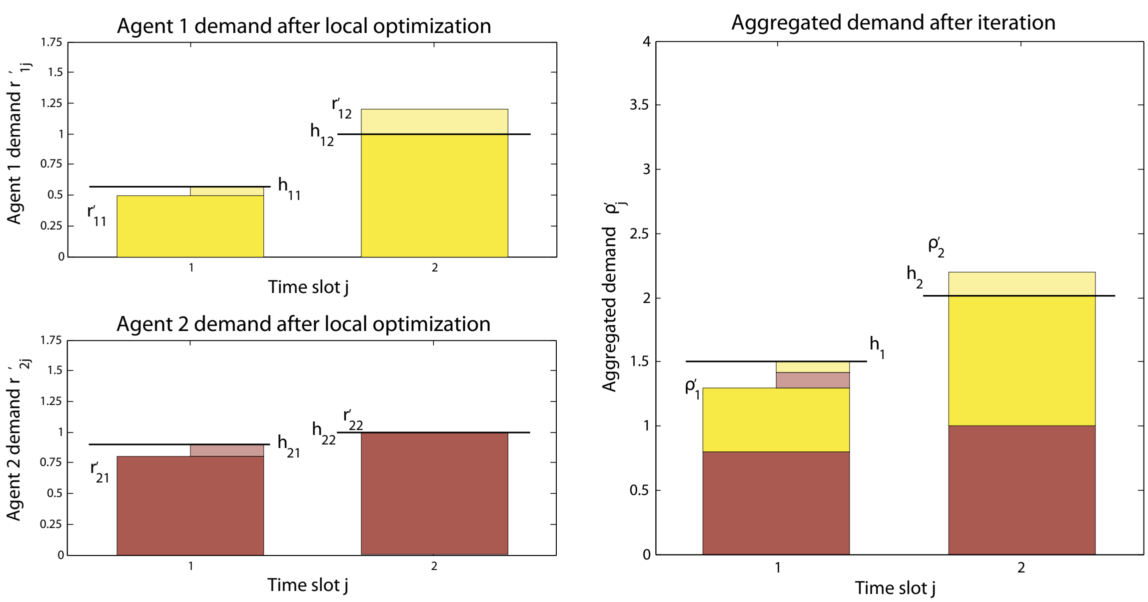

Figure 4: Demand after agent's local optimization. Agent 1 cannot reduce its entire demand from time slot 2 (i.e., $r_{12} > h_{12}$), because of its own demand constraints (left top plot). Thus, the aggregated demand in time slot 1 is still below the threshold (i.e., $\rho'_1 < h_1$). Therefore, the central coordinator again sends a modified price signal to the agents and the algorithm continues until agents stop shifting their demands.

Although the prices p_j^H and p_j^L are derived from the market prices, s_{ij}^v is called a virtual price signal, because the threshold $\mathrm{h}_{ij}(\mathbf{R})$ changes, if the agents change their demand profiles. The coordinator chooses $\mathrm{h}_{ij}(\mathbf{R}) < r_{ij}$ to induce the agents to reduce demand or $\mathrm{h}_{ij}(\mathbf{R}) > r_{ij}$ to increase demand in one time slot. With Δ_{ij} as the amount the coordinator wants agent i to change demand in time slot j, the threshold $\mathrm{h}_{ij}(\mathbf{R})$ is updated based on the demand profiles submitted in the last iteration:

$$\mathrm{h}_{ij}(\mathbf{R}) = r_{ij} + \Delta_{ij} \tag{5}$$

Thus, the agents know that, at the current market price, they can at most change their demand in time slot j by Δ_{ij}. For the demand exceeding $\mathrm{h}_{ij}(\mathbf{R})$, they need to pay a higher price. The demand, Δ_j, the coordinator wants to change in time slot j is calculated as the difference between the current aggregated demand and the threshold of the market price:

$$\Delta_j = h_j - \rho_j \tag{6}$$

Since the coordinator wants the total demand change of the agents to be less than Δ_j, it has to ensure that $\sum_i \Delta_{ij} \leq \Delta_j$. The allowable shift for an agent is proportional to the agent's share of the aggregated demand in that time slot, i.e., $\Delta_{ij} = \frac{\Delta_j r_{ij}}{\sum_i r_{ij}}$. Figure 3 shows how the total demand change is divided among the agents in order to create the individual virtual price signals. The left side shows the initial aggregated demand of two agents. The yellow demand belongs to agent 1 and the red demand to agent 2. Since the aggregated demand is below the threshold in time slot 1 (i.e., $\rho_1 < h_1$) and above the threshold in time slot 2 (i.e., $\rho_2 > h_2$), the coordinator wants the agents to shift demand from time slot 2 to time slot 1. The amount of demand that can be shifted into time slot 1 is Δ_1, while Δ_2 is the amount of demand that should to be shifted out of time slot 2. The right side shows the individual thresholds of the agents as determined by the central coordinator using the procedure described above. The allocation of the demand change to the agents is illustrated by the dashed arrows. The current demand of agent 1 in time slot 1 is r_{11} and its threshold for time slot 1 is h_{11} (the other notations can be interpreted similarly).

4.1.3 THE AGENT'S RESPONSE TO THE VIRTUAL PRICE SIGNAL

Having received the virtual price signal, the agents will independently optimize their demand profiles in order to minimize their cost according to Problem 3. The agents' objective function $\mathrm{C}_i^v(\mathbf{r}_i|\mathbf{R}) = \sum_{j=1}^{M} \mathrm{s}_{ij}^v(r_{ij}|\mathbf{R})\, r_{ij} + \mathrm{g}_i(\mathbf{r}_i)$ can be written as:

$$\mathrm{C}_i^v(\mathbf{r}_i|\mathbf{R}) = \sum_{j=1}^{M} \left[p_j^H (r_{ij} - \mathrm{h}_{ij}(\mathbf{R}))^+ + p_j^L (r_{ij} - \mathrm{h}_{ij}(\mathbf{R}))^- + p_j^L \mathrm{h}_{ij}(\mathbf{R}) \right] + \mathrm{g}_i(\mathbf{r}_i) \tag{7}$$

Since the agents have to pay the high price p_j^H for demand exceeding their individual threshold, no agent will shift too much demand, based on a false impression of possible cost reduction. Figure 4 shows the agents' demand profiles after their individual optimization. The left side shows the individual problems of the agents, after they have optimized their demand profile. In comparison to the right side in Figure 3, it can be seen that agent 2 shifted the whole allocated amount, but agent 1 only shifted some part of it (e.g., due to its constraints). The right side in Figure 4 shows the central problem after the agents' individual optimization. It can be seen that there is still demand left to be shifted from time slot 2 to time slot 1. This remaining demand would again be divided among the agents in the subsequent iteration.

4.1.4 PAYMENT RULE

After the algorithm has converged, the agents have to pay the virtual unit price for their demand in each time slot. In the last iteration the virtual unit price equals the average unit price for the aggregate demand of the whole cooperative, because no agent changed its demand profile, Equation 9. With $\mathbf{R}$ as final demand profiles of the whole cooperative and $\mathbf{r}_i$ as final demand profile of agent i, the payment for agent i is given as

$$payment_i = \sum_{j=1}^{M} \mathrm{s}_{ij}^{v}\left(r_{ij}|\mathbf{R}\right) r_{ij} \tag{8}$$

4.2 Convergence of the Basic Algorithm

In this section we prove that the basic iterative procedure always converges to an optimal solution in the basic setting with $M = 2$, $\mathrm{g}_i(\cdot) = 0$ and $p_j^H > p_k^L,\ \forall j,k$. In Lemma 1, we show that the algorithm strictly reduces cost in every iteration. This fact will be used in Theorem 1 to show that the algorithm always converges. Then, in Theorem 2 we show that, when $M = 2$, $\mathrm{g}_i(\cdot) = 0$ and $p_j^H > p_k^L,\ \forall j,k$, the converged solution is an optimal solution. Subsequently, we show that the algorithm can get stuck in a suboptimal solution in general settings if $M > 2$ (Lemma 3) or if $\mathrm{g}_i(\cdot) \neq 0$ (Lemma 4).

Lemma 1. *The algorithm strictly reduces the total cost in every iteration:* $\mathrm{C}\left(\mathbf{R}'\right) < \mathrm{C}\left(\mathbf{R}\right)$.

Proof. Let's first introduce some notation that will be used throughout this proof. Let $\mathbf{R}$ be the total demand profile at the end of iteration round t and $\mathbf{R}'$ be the total demand profile at the end of round $t+1$. Similarly, let $\mathrm{C}(\mathbf{R})$ be the total cost at the end of iteration t and $\mathrm{C}(\mathbf{R}')$ be the total cost at the end of iteration $t+1$. The virtual prices for round $t+1$ are computed by the central coordinator using $\mathbf{R}$. Let $\mathrm{C}_i^v(\mathbf{r}_i|\mathbf{R})$ be the cost for agent i computed according to the virtual price signal for demands at the beginning of round $t+1$ and $\mathrm{C}_i^v(\mathbf{r}_i'|\mathbf{R})$ at the end of round $t+1$.

At the beginning of each iteration the total cost for the consumer group based on market prices (given by the objective function in Problem 2) equals the sum of the individual cost of the agents based on the virtual price signals (given by Problem 3), i.e., $\sum_{i=1}^{N} \mathrm{C}_i^v\left(\mathbf{r}_i|\mathbf{R}\right) = \mathrm{C}\left(\mathbf{R}\right)$:

$$\begin{aligned}
&\sum_{i=1}^{N}\sum_{j=1}^{M} \mathrm{s}_{ij}^{v}\left(r_{ij}|\mathbf{R}\right) r_{ij} + \sum_{i=1}^{N} \mathrm{g}_i(\mathbf{r}_i) \\
&= \sum_{i=1}^{N}\sum_{j=1}^{M} \left[p_j^H \left(r_{ij} - \left(r_{ij} + \frac{(h_j - \rho_j)\, r_{ij}}{\sum_i r_{ij}} \right)\right)^{+} \right.\\
&+ p_j^L \left(r_{ij} - \left(r_{ij} + \frac{(h_j - \rho_j)\, r_{ij}}{\sum_i r_{ij}} \right)\right)^{-} \\
&\left. + p_j^L \left(r_{ij} + \frac{(h_j - \rho_j)\, r_{ij}}{\sum_i r_{ij}} \right) \right] + \sum_{i=1}^{N} \mathrm{g}_i(\mathbf{r}_i) \\
&= \sum_{j=1}^{M} \left[p_j^H \left(\rho_j - h_j\right)^{+} + p_j^L \left(\rho_j - h_j\right)^{-} + p_j^L h_j \right] + \sum_{i=1}^{N} \mathrm{g}_i(\mathbf{r}_i) \\
&= \sum_{i=1}^{N}\sum_{j=1}^{M} \mathrm{p}_j\left(\rho_j\right) r_{ij} + \sum_{i=1}^{N} \mathrm{g}_i(\mathbf{r}_i)
\end{aligned} \tag{9}$$

If the algorithm has not stopped, at least one agent has changed its demand profile, i.e., $\exists i$ with $\mathbf{r_i'} \neq \mathbf{r_i}$. Agents only change their demand profile, if that reduces their cost according to Problem 3. Thus, given the virtual price, signal for agent i the cost of the new demand profile $\mathbf{r_i'}$ is strictly lower than of its previous demand profile $\mathbf{r_i}$:

$$\mathrm{C}_i^v\left(\mathbf{r}_i'|\mathbf{R}\right) < \mathrm{C}_i^v\left(\mathbf{r}_i|\mathbf{R}\right) \tag{10}$$

After all agents have submitted their new demand profile, the new aggregated demand is computed as: $\rho_j' = \sum_{i=1}^N r_{ij}'$. Next, we show that the sum of the agents' individual cost according to the virtual price signals is an upper bound on the total central cost at market prices. Thus, the total cost, based on the new aggregated demand, is lower or equal to the sum of the agents' individual cost, based on their new demand. This fact is very important, because it prevents the herding behavior: *No solution that reduces the cost for the agents' individual problems can lead to a worse solution in the central problem.* This is proved by showing with Equations 1 and 7 that for every time slot j the difference between the central cost for the aggregated demand and the sum of the agents' individual cost is less than or equal to 0. The following is a sketch of the proof omitting some algebraic steps for ease of readability. Consult Appendix A for the complete proof. For any time slot, j, we have

$$\begin{aligned}
&\sum_{i=1}^{N}\mathrm{p}_j\left(\rho_j'\right)r_{ij}' - \sum_{i=1}^{N}\mathrm{s}_{ij}^v\left(r_{ij}'|\mathbf{R}\right)r_{ij}'\\
&=\begin{cases}\sum_{i:\ r_{ij}'\le h_{ij}}^{N}\left(p_j^H-p_j^L\right)\left(r_{ij}'-\mathrm{h}_{ij}\left(\mathbf{R}\right)\right) & \rho_j'>h_j\\ \sum_{i:\ r_{ij}'>h_{ij}}^{N}\left(p_j^L-p_j^H\right)\left(r_{ij}'-\mathrm{h}_{ij}\left(\mathbf{R}\right)\right) & \rho_j'\le h_j\end{cases}\\
&=\begin{cases}\left[\sum_{i=1}^{N}\left(p_j^H-p_j^L\right)\left(r_{ij}'-\mathrm{h}_{ij}\left(\mathbf{R}\right)\right)^-\right]\le 0 & \rho_j'>h_j\\ \left[\sum_{i=1}^{N}\left(p_j^L-p_j^H\right)\left(r_{ij}'-\mathrm{h}_{ij}\left(\mathbf{R}\right)\right)^+\right]\le 0 & \rho_j'\le h_j\end{cases}
\end{aligned}\tag{11}$$

Since $\sum_{i=1}^N \mathrm{p}_j\left(\rho_j'\right) r_{ij}' \le \sum_{i=1}^N \mathrm{s}_{ij}^v\left(r_{ij}'|\mathbf{R}\right) r_{ij}'\ \forall j$, it also holds for the sum of all time slots: $\mathrm{C}\left(\mathbf{R}'\right) \le \sum_{i=1}^N \mathrm{C}_i^v\left(\mathbf{r}_i'|\mathbf{R}\right)$.

From Equations 9, 10 and 11 we can conclude that:

$$\mathrm{C}\left(\mathbf{R}'\right) \le \sum_{i=1}^{N}\mathrm{C}_i^v\left(\mathbf{r}_i'|\mathbf{R}\right) < \sum_{i=1}^{N}\mathrm{C}_i^v\left(\mathbf{r}_i|\mathbf{R}\right) = \mathrm{C}\left(\mathbf{R}\right) \tag{12}$$

Thus, the total cost is strictly reduced in each iteration. □

Theorem 1. *The basic iterative algorithm for solving Problem 2 always converges.*

Proof. From the definition we have that Problem 2 is convex and a lower bound on the total cost can be obtained by the sum of the individual initial demand profile costs at market prices. From Lemma 1 we have that the algorithm reduces the total cost in each iteration. Thus, it can be concluded that the algorithm converges. □

Lemma 2. *The real market cost, $\sum_{i=1}^N \mathrm{p}_j\left(\Sigma_{i=1}^N r_{ij}\right) r_{ij}$, equals the virtual cost $\sum_{i=1}^N \mathrm{s}_{ij}^v\left(r_{ij}|\mathbf{R}\right) r_{ij}$ in a time slot j, if either $\forall i,\ r_{ij} \ge h_{ij}\left(\mathbf{R}\right)$ or $\forall i,\ r_{ij} \le h_{ij}\left(\mathbf{R}\right)$.*

Proof. If $\forall i,\ r_{ij} \ge h_{ij}\left(\mathbf{R}\right)$, then $\rho_j \ge h_j$ and from the first case in Equation 11 follows the cost are equal. If $\forall i,\ r_{ij} \le h_{ij}\left(\mathbf{R}\right)$, then $\rho_j \le h_j$ and from the second case in Equation 11 follows the cost are equal. □

Theorem 2. *The solution* $\mathbf{R}$ *of the basic algorithm is optimal, in either of the following cases*

A) $\forall j$, $\mathbf{R}$ *satisfies* $\rho_j \neq h_j$ *with no restrictions on* M *and* $\mathrm{g}_i(\cdot)$

B) $\mathbf{R}$ *satisfies* $\exists j,$ *s.t.* $\rho_j = h_j$ *and* $M = 2$, $\mathrm{g}_i(\cdot) = 0$ *and the high price in any time slot is greater than the low price in any other time slot, i.e.,* $p_j^H > p_k^L,\ \forall j, k$

Proof. We prove by contradiction that when the algorithm has converged to the solution $\mathbf{R}$ that satisfies Case A, there is no other solution $\mathbf{R}'$ with lower cost. Suppose there exists such a solution $\mathbf{R}'$, i.e., $\mathrm{C}(\mathbf{R}') < \mathrm{C}(\mathbf{R})$. We will show that this contradicts the convergence conditions of the algorithm that $\forall i, \mathbf{r}_i$ is the solution of the agent's individual problem, $\mathbf{r}_i = \arg\min_{\mathbf{x}_i \in \mathcal{X}_i} \mathrm{C}_i^v(\mathbf{x}_i|\mathbf{R})$. Denote $\mathbf{R}^{\#}(\alpha) = \alpha\mathbf{R} + (1-\alpha)\mathbf{R}'$, $\alpha \in (0,1)$, a linear combination of $\mathbf{R}$ and $\mathbf{R}'$. Since the central cost function is convex and $\mathrm{C}(\mathbf{R}') < \mathrm{C}(\mathbf{R})$, we also have $\mathrm{C}(\mathbf{R}^{\#}(\alpha)) < \mathrm{C}(\mathbf{R})$. If $\forall j$, $\mathbf{R}$ satisfies $\rho_j \neq h_j$, then for the demand in the beginning of the iteration r_{ij} we have $\forall j$ either $\forall i, r_{ij} > \mathrm{h}_{ij}(\mathbf{R})$ or $\forall i, r_{ij} < \mathrm{h}_{ij}(\mathbf{R})$, because if $\Delta_j > 0$ then $\forall i, \Delta_{ij} > 0$ and similarly if $\Delta_j < 0$ then $\forall i, \Delta_{ij} < 0$. It follows, $\exists \alpha \in (0,1)$, s.t. $\forall j$, either $\forall i, r_{ij}^{\#}(\alpha) \geq \mathrm{h}_{ij}(\mathbf{R})$ or $\forall i, r_{ij}^{\#}(\alpha_1) \leq \mathrm{h}_{ij}(\mathbf{R})$.

Therefore, by Lemma 2, $\sum_{i=1}^{N} \mathrm{C}_i^v\left(\mathbf{r}_i^{\#}(\alpha)|\mathbf{R}\right) = \mathrm{C}(\mathbf{R}^{\#}(\alpha))$. Moreover, we have $\mathrm{C}(\mathbf{R}^{\#}(\alpha)) < \mathrm{C}(\mathbf{R})$ and from Equation 9 we also have $\sum_{i=1}^{N} \mathrm{C}_i^v(\mathbf{r}_i|\mathbf{R}) = \mathrm{C}(\mathbf{R})$. It follows, $\exists i$, s.t. $\mathrm{C}_i^v\left(\mathbf{r}_i^{\#}(\alpha)|\mathbf{R}\right) < \mathrm{C}_i^v(\mathbf{r}_i|\mathbf{R})$, which conflicts with $\mathbf{r}_i = \arg\min_{\mathbf{x}_i \in \mathcal{X}_i} \mathrm{C}_i^v(\mathbf{x}_i|\mathbf{R})$, i.e., $\mathbf{r}_i$ is not the solution of the agent's individual Problem.

We prove by contradiction that when the algorithm has converged to the solution $\mathbf{R}$ that satisfies Case B, there is no other solution $\mathbf{R}'$ with lower cost. Assume there exists a solution $\mathbf{R}'$ with $\mathrm{C}(\mathbf{R}') < \mathrm{C}(\mathbf{R})$. Let the 2 time slots be $\{j, k\}$, where $\rho_j' < \rho_j$ and $\rho_k' > \rho_k$ (without loss of generality). Note that $\mathrm{p}_j^m(\rho_j) > \mathrm{p}_k^m(\rho_k)$, because otherwise the new cost would not be lower. If $\rho_j = h_j$, $\rho_k = h_k$ then $\mathrm{p}_j^m(\rho_j) = p_j^L$ (as ρ_j decreases) and $\mathrm{p}_k^m(\rho_k) = p_k^H$ (as ρ_k increases). From the market price structure we have $p_k^H > p_j^L$. It follows $\mathrm{p}_j^m(\rho_j) < \mathrm{p}_k^m(\rho_k)$, which leads to a contradiction. If $\rho_j \neq h_j$, $\rho_k = h_k$ and if $\rho_j < h_j$ then $\mathrm{p}_j^m(\rho_j) = p_j^L$ and $\mathrm{p}_k^m(\rho_k) = p_k^H$ thus $\mathrm{p}_j^m(\rho_j) < \mathrm{p}_k^m(\rho_k)$. If $\rho_j > h_j$ then $\mathrm{p}_j^m(\rho_j) = p_j^H$ and $\mathrm{p}_k^m(\rho_k) = p_k^H$. We have $p_k^H > p_j^H$, because otherwise a shift from j to k would be beneficial for at least one agent and the algorithm would not have stopped. It follows $\mathrm{p}_j^m(\rho_j) < \mathrm{p}_k^m(\rho_k)$, which again leads to a contradiction. The case of $\rho_j = h_j$, $\rho_k \neq h_k$ works similarly.

It follows that $\mathbf{R}$ is the optimal solution, as no solution with lower cost exists. Thus, the proposed iterative algorithm converges to the optimal solution. Since the problem is convex that solution is also the global optimal solution (Boyd & Vandenberghe, 2004). □

Lemma 3. *The basic algorithm could converge to a suboptimal solution* $\mathbf{R}$ *in settings, where* $\mathbf{R}$ *satisfies* $\exists j,$ *s.t.* $\rho_j = h_j$, $M > 2$ *and* $\mathrm{g}_i(\cdot) = 0$.

Proof. We now prove by presenting a counterexample that the basic algorithm can converge to a suboptimal solution in settings, where $\mathbf{R}$ satisfies $\exists j,$ s.t. $\rho_j = h_j$, $M > 2$ and $\mathrm{g}_i(\cdot) = 0$. Consider a population of 2 agents, $N = 2$, and a planning horizon of 3 time slots, $M = 3$. The agents' constraints are in a form so that in the converged solution the aggregated demand is below the threshold in one time slot, directly at the threshold in another time slot and in one time slot above the threshold.

Let the price function for the three time slots be given as:

$$(p_1^L, p_1^H) = (3, 6),\ h_1 = 10;\ \ (p_2^L, p_2^H) = (2, 5),\ h_2 = 10;\ \ (p_3^L, p_3^H) = (1, 4),\ h_3 = 10;$$

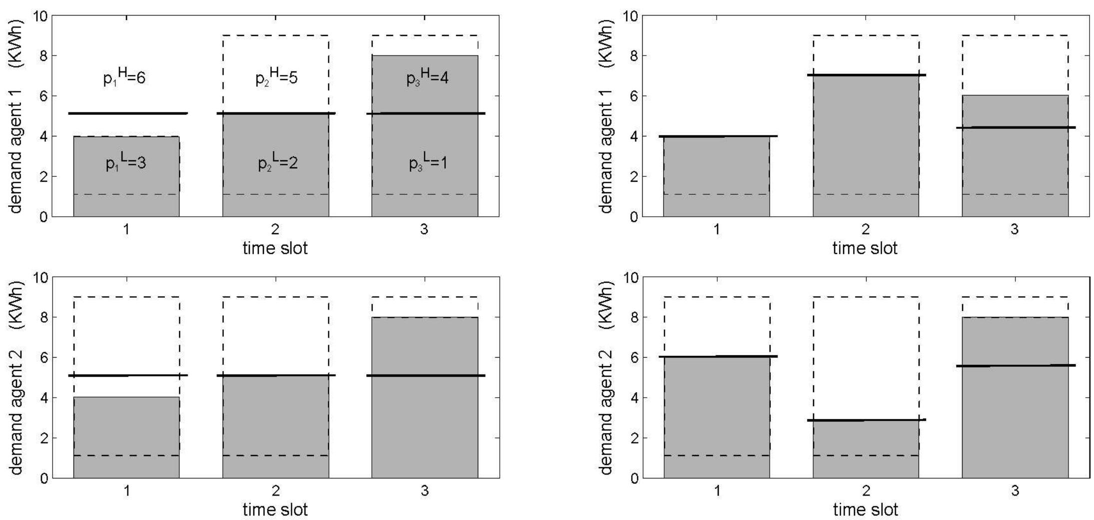

(a) Converged Solution $\mathbf{R}^{(T)}$ (b) Optimal Solution $\mathbf{R}'$

Figure 5: Converged and Optimal solution for the example in Lemma 3. The top row shows the demand of agent 1 and the bottom row the demand of agent 2. The dashed lines indicate the upper and lower bounds of the demands of the agents in a time slot. The solid black lines show the price thresholds for each time slot.

The individual constraints on the agents' demand are: (a) upper and lower bounds on the demand in each time slot and (b) constant total demand over all time slots. Specifically,

$$r_{11} \in [1,4],\ r_{12} \in [1,9],\ r_{13} \in [1,9],\ r_{11} + r_{12} + r_{13} = \tau_1 = 17$$
$$r_{21} \in [1,9],\ r_{22} \in [1,9],\ r_{23} \in [8,9],\ r_{21} + r_{22} + r_{23} = \tau_2 = 17$$

Let $\mathbf{R}^{(t)}$ denote the demand profile of the agents in the tth iteration and let $t = 1$ be the initial iteration and $t = T$ be the final iteration at convergence. At the beginning the agents compute their initial demand profiles based on the market prices $\mathbf{r}_1^{(1)} = (1, 7, 9)$, $\mathbf{r}_2^{(1)} = (1, 7, 9)$. The cost based on the initial demand profiles is $\mathrm{C}\left(\mathbf{R}^{(1)}\right) = 88$. At convergence, the profiles of the two agents are $\mathbf{r}_1^{(T)} = (4, 5, 8)$, $\mathbf{r}_2^{(T)} = (4, 5, 8)$ (see Figure 5a for a graphical representation). The cost based on these demand profiles is $\mathrm{C}\left(\mathbf{R}^{(T)}\right) = 78$. However, a different demand profile of the agents exists with $\mathbf{r}'_1 = (4, 7, 6)$, $\mathbf{r}'_2 = (6, 3, 8)$ (see Figure 5b for a graphical representation). This profile is feasible and leads to lower total cost, i.e., $\mathrm{C}\left(\mathbf{R}'\right) = 76$. It follows that the algorithm has stopped in a suboptimal solution. □

Lemma 4. *The basic algorithm could converge to a suboptimal solution* $\mathbf{R}$ *in settings, where* $\mathbf{R}$ *satisfies* $\exists j,$ *s.t.* $\rho_j = h_j$, $M = 2$ *and* $\mathrm{g}_i(\cdot) \neq 0$.

Proof. In Appendix B a counterexample is presented, proving that the basic algorithm can converge to a suboptimal solution in settings, where $\mathbf{R}$ satisfies $\exists j,$ s.t. $\rho_j = h_j$, $M = 2$ and $\mathrm{g}_i(\cdot) \neq 0$. □

5. General Coordination Algorithm

In the previous sections we introduced a basic iterative coordination algorithm for optimal energy allocation in the basic setting with a planning horizon of only two time slots, $M = 2$, and no cost for shifting demand, $\mathrm{g}(\cdot) = 0$. Moreover, we presented counter examples showing that, in general settings with $M > 2$ or $\mathrm{g}(\cdot) \neq 0$, the algorithm can get stuck in a suboptimal solution. In this section, we present a general coordination algorithm for the energy allocation in general settings with more than two time slots and non-zero individual costs for shifting demand.

5.1 The Additional Phase

The reason for convergence to a suboptimal solution is that the coordinator cannot determine the agents' marginal valuation for their demand, if the aggregated demand in some time slot is at a threshold. The agents might have different marginal valuations, so that a shift from an agent with a lower valuation to an agent with a higher valuation might be beneficial. However, this opportunity for improvement cannot be realized, because the coordinator does not know the marginal valuations of the agents.

We now introduce an additional phase to the algorithm to address this problem. When the basic algorithm has converged, the coordinator checks whether the aggregated demand is at the threshold for at least one time slot. If that is not the case, the algorithm stops and the solution is optimal. However, if the demand is at the threshold in some time slot, the coordinator initiates another phase. The basic idea is that, besides the virtual price signal, the coordinator sends an additional query to the agents, requesting their valuations for an ε-increase and ε-decrease of threshold in the time slots, where the aggregated demand is at the threshold. Then, the coordinator uses this information to adjust the virtual price signals.

Let v_{ij}^{+} denote the value for agent i of an ε-increase of its threshold in time slot j. Let further $\mathbf{R}^{ij+}$ be the demand profile, where in the resulting price signal for agent i the threshold in time slot j is increased by ε, i.e., $h_{ij}(\mathbf{R}^{ij+}) = h_{ij}(\mathbf{R}) + \varepsilon$. Moreover, let $\mathrm{C}_i^v\left(\mathbf{r}_i^{j+}|\mathbf{R}^{ij+}\right)$ be its virtual cost, where the threshold in time slot j is increased by ε.

Definition 1. *Agent i's* ***marginal valuation*** *for an ε-increase of the threshold in time slot j is the difference in its virtual cost between the original virtual problem and the virtual problem with an ε-increased threshold in time slot j. The valuation for an ε-decrease is defined analogously.*

$$v_{ij}^{+}(\mathbf{R}) = \mathrm{C}_i^v\left(\mathbf{r}_i^{j+}|\mathbf{R}^{ij+}\right) - \mathrm{C}_i\left(\mathbf{r}_i'|\mathbf{R}\right);\;\; v_{ij}^{-}(\mathbf{R}) = \mathrm{C}_i^v\left(\mathbf{r}_i^{j-}|\mathbf{R}^{ij-}\right) - \mathrm{C}_i\left(\mathbf{r}_i'|\mathbf{R}\right). \quad (13)$$

When the agents receive their virtual price signal and the query for their marginal valuations, they first compute their optimal demand profile $\mathbf{r}_i'$, based on the price signal. Second, they compute their marginal valuations $v_{ij}^{+}(\mathbf{R})$ and $v_{ij}^{-}(\mathbf{R})$ for every time slot, with the aggregated demand at the threshold. Let the vector $\mathbf{v}_i^{+}$ consist of the valuations for ε-increments and $\mathbf{v}_i^{-}$ for decrements respectively. Finally, the agents send $\mathbf{r}_i$, $\mathbf{v}_i^{+}$ and $\mathbf{v}_i^{-}$ back to the coordinator.

Having received the agents' information, the coordinator computes the virtual price signals for the next iteration. If the aggregated demand is at the threshold, it finds the agents with the lowest cost for an ε-increase of the threshold: $l = \arg\min_i \left\{v_{ij}^{+}\right\}$ and the lowest cost for an ε-decrease of the threshold: $k = \arg\min_i \left\{v_{ij}^{-}\right\}$. If the combined cost is negative $v_{kj}^{-} + v_{lj}^{+} < 0$, a beneficial shift from agent l to agent k exists. In that case, the coordinator updates the thresholds of agents l, k as

$$\Delta_{lj} = \varepsilon, \ \ \Delta_{kj} = -\varepsilon.$$

The algorithm stops, if the demand is not at the threshold in any time slot or if for all time slots at the thresholds no pair $\{l, k\}$ exists with $v^-_{kj} + v^+_{lj} < 0$. It follows, when the algorithm converges, in every time slot with the aggregated demand at the threshold, $\rho_j = h_j$, for all agents, the cost reduction of an ε-increase of the threshold is less than the additional cost from a ε-decrease of the threshold for all other agents, i.e.,

$$v^+_{lj} \geq -v^-_{kj}, \ \forall l, k \in \{1, ..., N\} \text{ and } \forall j \text{ s.t. } \rho_j = h_j. \tag{14}$$

Otherwise, the coordinator would change the price signals and the algorithm would not stop.

5.2 General Algorithm

The general algorithm is designed as an iterative algorithm, where the coordinator updates the virtual price signals based on the consumers' feedback and thus gradually adjusts the individual demands to the central optimal solution. The overall algorithm is shown in Algorithm 1. Each iteration consists of two steps: First, the central coordinator aggregates the demand submitted by the agents and computes virtual price signals for each agent. The virtual price signals are computed in the subroutine $\mathrm{CalculateVirtualPriceSignals}$ (Algorithm 2). Second, the individual agents use the virtual price signal to solve their individual cost optimization problem and compute their marginal valuation for their electricity demand and report them to the coordinator. The agents' response is computed in $\mathrm{CalculateDemandProfileAndValuation}$ (Algorithm 3).

The general algorithm first only runs the basic algorithm until convergence and subsequently performs the additional phase, if necessary. In particular, the additional phase is performed, if the basic algorithm has converged (Algorithm 1 line 8) and in at least one time slot the aggregated demand is not at the threshold (Algorithm 1 line 9). The algorithm reaches the optimal solution when $\mathbf{R} = \mathbf{R}'$ and either $\rho_j \neq h_j, \forall j$ or in the additional phase $v^+_{lj} \geq -v^-_{kj}$, $\forall l, k \in \{1, ..., N\}, \forall j$ s.t. $\rho_j = h_j$.

5.3 Convergence of the General Algorithm to an ε-optimal Solution

Please note that Theorem 1 still holds for the general algorithm, because the algorithm still reduces the total cost in each iteration. Thus, the extended algorithm converges. In Theorem 3 we prove that the solution of the general algorithm lies within an ε-neighborhood of the optimal solution.

Theorem 3. *The converged solution $\mathbf{R}$ of the general algorithm lies within an ε neighborhood of the optimal solution $\mathbf{R}^*$, where ε is the amount for which the agents compute their marginal valuations.*

Proof. We now prove by contradiction that when the general algorithm has converged to the solution $\mathbf{R}$, then no other solution $\mathbf{R}'$ exists with lower cost with respect to the central Problem 2 outside of an ε-neighborhood around $\mathbf{R}$. Suppose there exists a solution $\mathbf{R}'$ having a lower total energy cost than $\mathbf{R}$, i.e., $\mathrm{C}(\mathbf{R}') < \mathrm{C}(\mathbf{R})$. We will show that this contradicts the convergence conditions of the general algorithm that $\forall i, \mathbf{r}_i$ is the solution of the agent's individual problem, $\mathbf{r}_i = \arg\min_{\mathbf{x}_i \in \mathcal{X}_i} \mathrm{C}^v_i\left(\mathbf{x}_i|\mathbf{R}\right)$, and that $v^+_{lj} \geq -v^-_{kj}$, $\forall l, k \in \{1, ..., N\}$ and $\forall j$ s.t. $\rho_j = h_j$.

The only case left from Theorem 2 is that $\exists j$, $\rho_j = h_j$, which implies $\forall i$, $r_{ij} = \mathrm{h}_{ij}(\mathbf{R})$, and that $\exists L, K \neq \emptyset$, s.t. $\forall l \in L$, $r'_{lj} > \mathrm{h}_{lj}(\mathbf{R})$ and $\forall k \in K$, $r'_{kj} < \mathrm{h}_{kj}(\mathbf{R})$, where $\exists l, k$ s.t. $\left|r'_{lj} - r_{lj}\right| \geq \varepsilon$, $\left|r'_{kj} - r_{kj}\right| \geq \varepsilon$. We will show that this case conflicts with the convergence conditions, too.

Algorithm 1: Overall algorithm.

Data: Scenario with electricity contract and agent definitions.
Result: Optimal demand schedule $\mathbf{R}^*$.

Initialization: All agents compute an initial energy demand profile $\mathbf{r}_i$ by solving Problem 3 based on the market prices and send it to the coordinator.;
$addPhase = false$;
while *demand schedule* $\mathbf{R}$ *is not optimal* **do**
  Coordinator calculates the price signals using Algorithm 2:
  $\mathrm{s}^v_{ij}\left(r_{ij}|\mathbf{R}\right), \forall i, j \leftarrow$ CalculateVirtualPriceSignals $(\mathbf{r}_i, \mathbf{v}^+_i, \mathbf{v}^-_i, addPhase)$;
  The coordinator sends the virtual price signals to all agents;
  Agents calculate demand profiles and valuations for all time slots, j, using Algorithm 3:
  $\forall i$: $\mathbf{r}'_i, \mathbf{v}^+_i, \mathbf{v}^-_i \leftarrow$ CalculateDemandProfileAndValuation $(\mathrm{s}^v_{ij}\left(r_{ij}|\mathbf{R}\right), addPhase)$;
  The agents send their new demand profiles and valuations back to the coordinator;
  if *algorithm has converged, i.e.,* $\mathbf{R} = \mathbf{R}'$ **then**
    if $\rho_j \neq h_j, \forall j$ **then**
      the demand schedule $\mathbf{R}'$ is optimal;
    else
      if $addPhase$ *and* $v^+_{lj} \geq -v^-_{kj}$, $\forall l, k \in \{1, ..., N\}, \forall j$ *s.t.* $\rho_j = h_j$ **then**
        the demand schedule $\mathbf{R}'$ is optimal;
      else
        start additional phase, i.e., $addPhase = true$;
      end
    end
  else
    $\mathbf{R} \leftarrow \mathbf{R}'$;
  end
end

We start with the case in which $|L| = |K| = 1$. Then we have $\exists \alpha \in (0, 1)$, s.t. $\left| r^{\#}_{lj}(\alpha) - r_{lj} \right| = \left| r^{\#}_{kj}(\alpha) - r_{kj} \right| = \varepsilon$. Moreover, denote $\mathbf{R}^{\#\#}(\alpha)$ s.t.,

$$r^{\#\#}_{it}(\alpha) = \begin{cases} r^{\#}_{it}(\alpha) & \text{if } i \in L \cup K, t = j \\ r_{it} & \text{otherwise} \end{cases}$$

a demand profile reflecting the changes in demand in time slot j so that in the resulting individual problems again all agents have their demand on the threshold, $r^{\#}_{lj} = \mathrm{h}_{lj}(\mathbf{R}^{\#\#})$ and $r^{\#}_{kj} = \mathrm{h}_{kj}(\mathbf{R}^{\#\#})$. Thus, with Lemma 2 we get that the sum of the virtual cost is equal to the central cost.

$$\begin{aligned} & \sum_{i=1, i\neq l,k}^{N} \mathrm{C}^v_i\left(\mathbf{r}^{\#}_i(\alpha)\,|\mathbf{R}\right) + \mathrm{C}^v_l\left(\mathbf{r}^{\#}_l(\alpha)\,|\mathbf{R}^{\#\#}(\alpha)\right) + \mathrm{C}^v_k\left(\mathbf{r}^{\#}_k(\alpha)\,|\mathbf{R}^{\#\#}(\alpha)\right) \\ = \quad & \mathrm{C}(\mathbf{R}^{\#}(\alpha)) < \mathrm{C}(\mathbf{R}) = \sum_{i=1}^{N} \mathrm{C}^v_i\left(\mathbf{r}_i|\mathbf{R}\right) \end{aligned}$$

Algorithm 2: Calculation of the virtual price signal by the coordinator.

$\mathrm{CalculateVirtualPriceSignals}$ ($\mathbf{r}_i, \mathbf{v}_i^+, \mathbf{v}_i^-, addPhase$);
Data: demand profiles and valuations of all agents, $\mathbf{r}_i \mathbf{v}_i^+, \mathbf{v}_i^-, \forall i$.
Result: virtual price signals for all agents, $\mathrm{s}_{ij}^v \left(r_{ij}|\mathbf{R}\right), \forall i, j$.
Compute the aggregated demand: $\rho_j \leftarrow \sum_{i=1}^N r_{ij}$;
Compute demand to be shifted in each time slot: $\Delta_j \leftarrow h_j - \rho_j$;
Divide that demand among all agents: $\Delta_{ij} \leftarrow \frac{\Delta_j r_{ij}}{\sum_i r_{ij}}$;
if $addPhase$ **then**
  Find time slot j and agents l, k s.t. $\min\left\{v_{lj}^+ - v_{kj}^-\right\}$;
  if $v_{lj}^+ + v_{kj}^- < 0$ **then**
    Adapt demand to be shifted for agent l: $\Delta_{lj} \leftarrow \varepsilon$ and agent k: $\Delta_{kj} \leftarrow -\varepsilon$;
  end
end
Compute thresholds based on demands to be shifted: $h_{ij} \leftarrow r_{ij} + \Delta_{ij}$;

Algorithm 3: Calculation of the individual demand profiles and marginal valuation of energy for agent i.

$\mathrm{CalculateDemandProfileAndValuation}$ ($\mathrm{s}_{ij}^v \left(r_{ij}|\mathbf{R}\right), addPhase$);
Data: virtual price signal $\mathrm{s}_{ij}^v \left(r_{ij}|\mathbf{R}\right)$, $j = 1, \ldots, M$.
Result: new demand profile $\mathbf{r}_i'$ and valuations $\mathbf{v}_i^+, \mathbf{v}_i^-$.
Compute new demand profile $\mathbf{r}_i'$ according to virtual price signal:
$\mathbf{r}_i = \arg\min_{\mathbf{x}_i \in \mathcal{X}_i} \mathrm{C}_i^v \left(\mathbf{x}_i|\mathbf{R}\right)$;
if $addPhase$ **then**
  foreach *time slot j at the threshold, $r_{ij} = h_{ij}$* **do**
    Compute valuation of increased threshold: $v_{ij}^+ \leftarrow \mathrm{C}_i^v\left(\mathbf{r}_i^{j+}|\mathbf{R}^{ij+}\right) - \mathrm{C}_i^v(\mathbf{r}_i'|\mathbf{R})$;
    Compute valuation of decreased threshold: $v_{ij}^- \leftarrow \mathrm{C}_i^v\left(\mathbf{r}_i^{j-}|\mathbf{R}^{ij-}\right) - \mathrm{C}_i^v(\mathbf{r}_i'|\mathbf{R})$;
  end
end

which implies either

$$\begin{aligned} \sum_{i=1, i\neq l,k}^{N} \mathrm{C}_i^v \left(\mathbf{r}_i^{\#}\left(\alpha\right)|\mathbf{R}\right) &< \sum_{i=1, i\neq l,k}^{N} \mathrm{C}_i^v\left(\mathbf{r}_i|\mathbf{R}\right) \\ \Rightarrow \quad \exists i, s.t.\, \mathrm{C}_i^v \left(\mathbf{r}_i^{\#}\left(\alpha\right)|\mathbf{R}\right) &< \mathrm{C}_i^v\left(\mathbf{r}_i|\mathbf{R}\right) \end{aligned}$$

which conflicts with the convergence condition $\mathbf{r}_i = \arg\min_{\mathbf{x}_i \in \mathcal{X}_i} \mathrm{C}_i^v\left(\mathbf{x}_i|\mathbf{R}\right)$;
or

$$\begin{aligned} \mathrm{C}_l^v \left(\mathbf{r}_l^{\#}\left(\alpha\right)|\mathbf{R}^{\#\#}\left(\alpha\right)\right) + \mathrm{C}_k^v \left(\mathbf{r}_k^{\#}\left(\alpha\right)|\mathbf{R}^{\#\#}\left(\alpha\right)\right) &< \mathrm{C}_l^v\left(\mathbf{r}_l|\mathbf{R}\right) + \mathrm{C}_k^v\left(\mathbf{r}_k|\mathbf{R}\right) \\ \Rightarrow \quad \mathrm{C}_l^v \left(\mathbf{r}_l^{\#}\left(\alpha\right)|\mathbf{R}^{\#\#}\left(\alpha\right)\right) - \mathrm{C}_l^v\left(\mathbf{r}_l|\mathbf{R}\right) &< -\left[\mathrm{C}_k^v \left(\mathbf{r}_k^{\#}\left(\alpha\right)|\mathbf{R}^{\#\#}\left(\alpha\right)\right) - \mathrm{C}_k^v\left(\mathbf{r}_k|\mathbf{R}\right)\right] \end{aligned}$$

For $\left|r_{kj}^{\#}(\alpha) - r_{kj}\right| = \varepsilon$ we get that $\mathrm{C}_l^v\left(\mathbf{r}_l^{\#}(\alpha)\,|\mathbf{R}^{\#\#}(\alpha)\right)$ is equal to $\mathrm{C}_l^v\left(\mathbf{r}_l^{j+}|\mathbf{R}^{ij+}\right)$ and $\mathrm{C}_k^v\left(\mathbf{r}_k^{\#}(\alpha)\,|\mathbf{R}^{\#\#}(\alpha)\right)$ is equal to $\mathrm{C}_k^v\left(\mathbf{r}_k^{j-}|\mathbf{R}^{ij-}\right)$. Thus, with Equation 13 follows

$$v_{lj}^{+}(\mathbf{R}) < -v_{kj}^{-}(\mathbf{R})$$

which conflicts with the convergence condition, $\forall l,k \in \{1,2,\cdots,N\}$, $v_{lj}^{+}(\mathbf{R}) \geq -v_{kj}^{-}(\mathbf{R})$.

For the general case of $|L|, |K| \geq 1$, we can get

$$\begin{aligned}
&\sum_{i=1, i\notin L\cup K}^{N} \mathrm{C}_i^v\left(\mathbf{r}_i^{\#}(\alpha)\,|\mathbf{R}\right) + \sum_{l\in L} \mathrm{C}_l^v\left(\mathbf{r}_l^{\#}(\alpha)\,|\mathbf{R}^{\#\#}(\alpha)\right) + \sum_{k\in K} \mathrm{C}_k^v\left(\mathbf{r}_k^{\#}(\alpha)\,|\mathbf{R}^{\#\#}(\alpha)\right) \\
= \;& \mathrm{C}(\mathbf{R}^{\#}(\alpha)) < \mathrm{C}(\mathbf{R}) = \sum_{i=1}^{N} \mathrm{C}_i^v(\mathbf{r}_i|\mathbf{R})
\end{aligned}$$

which implies either

$$\begin{aligned}
&\sum_{i=1, i\notin L\cup K}^{N} \mathrm{C}_i^v\left(\mathbf{r}_i^{\#}(\alpha)\,|\mathbf{R}\right) < \sum_{i=1, i\notin L\cup K}^{N} \mathrm{C}_i^v(\mathbf{r}_i|\mathbf{R}) \\
\Rightarrow\;& \exists i, s.t.\, \mathrm{C}_i^v\left(\mathbf{r}_i^{\#}(\alpha)\,|\mathbf{R}\right) < \mathrm{C}_i^v(\mathbf{r}_i|\mathbf{R})
\end{aligned}$$

which again conflicts with the convergence condition $\mathbf{r}_i = \arg\min_{\mathbf{x}_i \in \mathcal{X}_i} \mathrm{C}_i^v(\mathbf{x}_i|\mathbf{R})$;

or

$$\begin{aligned}
&\sum_{l\in L} \mathrm{C}_l^v\left(\mathbf{r}_l^{\#}(\alpha)\,|\mathbf{R}^{\#\#}(\alpha)\right) + \sum_{k\in K} \mathrm{C}_k^v\left(\mathbf{r}_k^{\#}(\alpha)\,|\mathbf{R}^{\#\#}(\alpha)\right) < \sum_{l\in L} \mathrm{C}_l^v(\mathbf{r}_l|\mathbf{R}) + \sum_{k\in K} \mathrm{C}_k^v(\mathbf{r}_k|\mathbf{R}) \\
\Rightarrow\;& \sum_{l\in L} \mathrm{C}_l^v\left(\mathbf{r}_l^{\#}(\alpha)\,|\mathbf{R}^{\#\#}(\alpha)\right) - \sum_{l\in L} \mathrm{C}_l^v(\mathbf{r}_l|\mathbf{R}) < -\left[\sum_{k\in K} \mathrm{C}_k^v\left(\mathbf{r}_k^{\#}(\alpha)\,|\mathbf{R}^{\#\#}(\alpha)\right) - \sum_{k\in K} \mathrm{C}_k^v(\mathbf{r}_k|\mathbf{R})\right]
\end{aligned}$$

Since the amount of demand increased in time slot j is equal to the demand decreased, and the cost decrease on the LHS is greater than the increase in cost on the RHS, there exists a pair of agents, l,k such that the marginal cost decrease of l is larger than the marginal cost increase of k

$$\Rightarrow \exists l \in L, k \in K, v_{lj}^{+}(\mathbf{R}) < -v_{kj}^{-}(\mathbf{R})$$

which conflicts with the convergence condition, $\forall l,k \in \{1,2,\cdots,N\}$, $v_{lj}^{+}(\mathbf{R}) \geq -v_{kj}^{-}(\mathbf{R})$. □

The solution $\mathbf{R}$ lies within an ε neighborhood of the optimal solution $\mathbf{R}^*$, where ε is the amount for which the agents compute their marginal valuations. When no beneficial shift of the size greater or equal to ε exists any more, the algorithm stops. However, a beneficial shift smaller than ε might still exist.

5.4 Convergence Time of the Algorithm

In this section, we give an upper bound on the number of steps the algorithm needs to converge. First, we give a bound for the basic algorithm in Lemma 5 and the extended phase in Lemma 6. Second, we give a bound for the general algorithm in Theorem 4.

Lemma 5. *The basic algorithm converges within a maximum of $\lceil -\log \delta \cdot N \rceil$ steps, where N is the number of agents and δ is the distance of the converged solution to the optimal solution of the basic algorithm.*

Proof. We prove this by showing an upper bound on the number of iterations needed in one time slot, say, time slot j, to get from the initial demand level into the neighborhood of the optimal demand level of the basic algorithm. This is sufficient, because when the demand in the last time slot has reached its optimal level, the demand in all other time slots also reached their optimal level. In this proof we will consider a time slot with its initial demand below the optimal demand. This is sufficient, because every shift consists of a demand decrease in at least one time slot and also an increase in at least one time slot. Thus, when all time slots that were below their optimal demand level, reached their optimal level, the remaining time slots also reached their optimal levels.

Since all variables in this proof are for time slot j, we omit the subscript j for increased readability. Let D_t be the difference between the optimal aggregated demand and the aggregated demand after iteration t. Since the demand is assumed to be below the optimal demand, $D_t > 0$. Without loss of generality, assume the initial value of D_t is 1 ($D_0 = 1$) and the optimum is 0; for each agent i, its initial demand $r_{ij} = 1$. The worst case is that only one agent (e.g., agent 1) can increase its demand. In this case, only agent 1 can increase its demand in each iteration, and its increment equals the amount of demand shift assigned by the coordinator.

We denote a_t as the increment of agent 1's demand in iteration t and c_t as agent 1's demand after iteration t ($c_0 = 1$). With N denoting the total number of agents, the demand shift assigned to agent i is given based on the definition of the price signal by $\Delta_{ij} = \frac{\Delta_j r_{ij}}{\sum_{i=1}^{N} r_{ij}}$. Then, we have for agent 1's increment

$$a_t = \frac{c_{t-1} D_{t-1}}{N - 1 + c_{t-1}}. \tag{15}$$

Agent 1's demand after iteration t is its initial demand plus its increments in the first t iterations. Similarly, the difference D_t is the initial difference minus agent 1's increments.

$$c_t = 1 + \sum_{l=1}^{t} a_l, \quad D_t = 1 - \sum_{l=1}^{t} a_l \Rightarrow a_t = c_t - c_{t-1} = -D_t + D_{t-1}, \quad c_t = 2 - D_t.$$

By substituting the above into Equation 15, we have

$$\begin{aligned} -D_t + D_{t-1} &= \frac{(2 - D_{t-1}) D_{t-1}}{N + 1 - D_{t-1}} \\ (N+1)(-D_t + D_{t-1}) + D_t D_{t-1} - D_{t-1}^2 &= 2D_{t-1} - D_{t-1}^2 \\ D_t &= \frac{(N-1) D_{t-1}}{N + 1 - D_{t-1}} \end{aligned}$$

Since $\forall n \geq 1$, $D_t \in (0, 1)$, we get

$$D_t < \frac{(N-1) D_{t-1}}{N} \Rightarrow D_t < \left(\frac{N-1}{N}\right)^n$$

Thus, $\forall \delta > 0$, $\exists t^* = \left\lceil \frac{\log \delta}{\log \frac{N-1}{N}} \right\rceil = \left\lceil \frac{-\log \delta}{\log N - \log(N-1)} \right\rceil$, s.t., $\forall t > t^*$, we have $D_t < \delta$. Moreover, since logarithm is strictly increasing and concave, we have $\log N - \log(N-1) > \frac{1}{N}$. Thus, $\forall \delta > 0$, $\exists t^* = \lceil -\log \delta \cdot N \rceil$, s.t., $\forall t > t^*$, we have $D_t < \delta$. $\square$

Lemma 6. *The extended phase converges at maximum within a number of* $\left\lceil \frac{(\mathrm{C}' - \mathrm{C}^*)}{\Delta_v} \right\rceil$ *steps, where* C' *is the total cost in the solution from the basic algorithm,* C^* *is the total cost in the optimal solution and* Δ_v *is the minimum value for the marginal valuations of an* ε*-increase and decrease of the threshold in one time slot.*

Proof. We will prove this by showing an upper bound on the number of iterations needed to reduce the total cost from the solution of the basic algorithm, i.e., C', to the total cost of the optimal solution, i.e., C^*. In each iteration of the extended phase, the thresholds are adapted for two agents, say k, j. Thus, the total cost is reduced by $v_{kj}^{-} + v_{lj}^{+}$ (Equation 13). Let Δ_v be the minimum value for $v_{kj}^{-} + v_{lj}^{+}$. Then, as long as the algorithm does not converge, after each iteration, the total cost is reduced at least by Δ_v. Thus, the algorithm stops at most within $\left\lceil \frac{(\mathrm{C}' - \mathrm{C}^*)}{\Delta_v} \right\rceil$ steps. $\square$

Theorem 4. *The time complexity of the general algorithm lies is* $\mathrm{O}(\mathrm{N} + (\mathrm{C}'\text{-}\mathrm{C}^*))$*, i.e., the algorithm is linear in the number of agents and pseudopolynomial in* $(C' - C^*)$.

Proof. From Lemma 5 we get the upper bound on the number of iterations of the basic algorithm as $\lceil -N \log \delta \rceil$. With $-\log \delta$ being a constant, which is independent of N, the time complexity of the basic algorithm is $\mathrm{O}(N)$. From Lemma 6 we get the upper bound on the number of iterations of the extended phase as $\left\lceil \frac{(\mathrm{C}' - \mathrm{C}^*)}{\Delta_v} \right\rceil$. The cost difference $(\mathrm{C}' - \mathrm{C}^*)$ does not depend on N or M, but on the specific characteristics of the agents. The minimum valuation Δ_v is proportional to the size of the threshold increment/decrement ε. Thus Δ_v is a constant for any constant choice of ε. Thus, the bound of the extended phase is pseudopolynomial in $(\mathrm{C}' - \mathrm{C}^*)$. $\square$

5.5 Incentive Compatibility

In general, for allocation problems, there are three key properties of interest, namely, efficiency, strategy-proofness, and budget-balance. In the context of our problem, efficiency implies that the overall electricity consumption cost is minimized. Strategy-proofness implies that telling the truth is the dominant strategy for the agents and hence they report truthfully. Budget-balance implies that the total amount paid by the members of the cooperative is equal to the actual electricity consumption cost. There are two well known facts about these three properties in the literature: First, if the allocation is efficient, the *only way* to implement payments that guarantee strategy-proofness is to use VCG-type mechanisms, see the work of Green and Laffont (1977), Hurwicz (1975). Second, as Green and Laffont (1977) and Hurwicz (1975) also prove, payments obtained by VCG-type mechanisms cannot achieve budget balance. Thus, it is impossible to design a mechanism that achieves all the three properties, namely, efficiency, budget balance and strategy-proofness.

As we have stated in our discussion in the related work section, budget balance is a key requirement of our problem from the application point of view. Furthermore, allocation efficiency is a desired social goal. Our algorithm achieves both budget balance (Lemma 2) and allocation efficiency (Theorem 3). Therefore, from the impossibility results, we cannot achieve strategy-proofness, i.e., truth telling cannot be a dominant strategy. However, we believe each agent may choose to truthfully respond, because agents cannot anticipate the future development of the algorithm and, as we prove below, *no manipulation strategy exists that dominates truth reporting*. Therefore, we define a notion of weak incentive compatibility.

Definition 2. *An algorithm satisfies* ***weak incentive compatibility*** *or is weakly incentive compatible, if and only if no strategy dominates telling the truth in the algorithm.*

Theorem 5. *The general algorithm is weakly incentive compatible.*

Proof. In our algorithm the agents have two possible ways to deviate from truthful reporting: First, they could report a demand profile that does not minimize their individual cost according to the current virtual price signals with the intent that this misreporting would benefit them for the final payment. Second, they could report marginal valuations that differ from their true marginal valuations. For our analysis we assume rational agents, i.e., agents always prefer lower cost.

Note that in the payment rule (Equation 8) each agent pays for its electricity demand based on its final price signal. In each iteration, a truthful demand report minimizes an agent's cost given the price signal. Therefore, if an agent does not report truthfully at the final price signal (i.e., final iteration) it will incur a higher cost. Furthermore, agents do not know whether any given price signal is the final price signal or not. This is, because agents have limited knowledge and do not know the demand profiles of other agents or their constraints and preferences. Thus, manipulation at any iteration may result in higher cost than a truthful report. Therefore, no manipulation strategy involving misreported demand profiles dominates truthful reporting.

Regarding the reporting of the marginal valuations (Definition 1), it is obvious that deflating its valuation hurts the agent. We now demonstrate with an extension of the example from Lemma 3, that deviating from truth revelation through inflating the reported marginal valuation can also hurt an agent. The aggregated demand is at the threshold in time slot 2 so that the coordinator asks for the agents' valuations for an ε-increase of the threshold. Without loss of generality assume $\varepsilon = 1$. Agent 1's true valuation is (reduce in time slot 3 and increase in 2) $v_{12}^{+} = p_3^H - p_2^L = 2$ and agent 2's is (reduce in time slot 1 and increase in 2) $v_{22}^{+} = p_1^H - p_2^L = 1$. Let's assume agent 2 misreports its valuation so that $v_{22}^{+} > 2$. Then, the coordinator increases the threshold in time slot 2 for agent 2 and decreases the threshold for agent 1. It follows that agent 1 has to shift demand from time slot 2 to time slot 3 and that agent 2 can now shift demand from time slot 1 to time slot 2. This leads to the demand profiles: $\mathbf{r}_1^{deviate} = (4, 4, 9)$, $\mathbf{r}_2^{deviate} = (3, 6, 8)$. However, this profile does not only lead to an increase in total cost, but also to an increase of the individual cost for the deviating agent 2. Table 1 shows the individual cost.[4] Since misreporting the valuations may result in higher electricity cost than a truthful report for the misreporting agent, no manipulation strategy involving misreported marginal valuations dominates truthful reporting. $\square$

cost	agent 1's cost	agent 2's cost	total cost
truthful reports	37.143	38,857	76
agent 2 deviates	40.118	38.882	79

Table 1: Cost of different demand profiles for individual agents.

6. Simulation

We perform simulations based on real world consumption data to (a) characterize the convergence properties of our algorithm and (b) understand the effect of different parameters that characterize the electricity demand profile on the cost reduction through coordination by our algorithm. The results show that the algorithm scales linearly with the number of agents and time slots. Further, we observe that as the participants' flexibility of shifting their demands increases, cost reduction

4. Since Agent 2's deviation increases both agent 1's as well as agent 2's cost, one might argue that increasing the other agents' cost could also be an incentive for deviation. However, note that the cost of an agent is not visible to others. Therefore, we believe that each agent may choose to report truthfully.

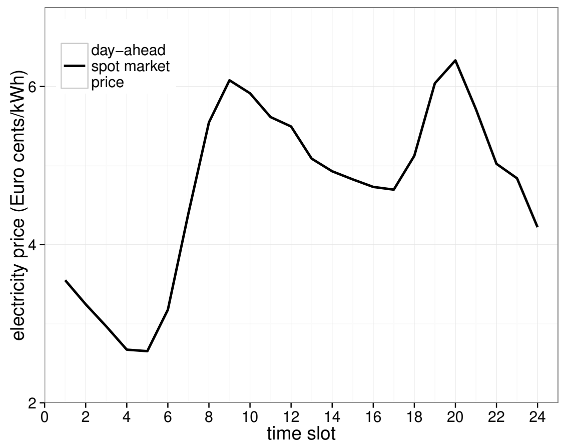

Figure 6: Day-ahead spot market prices from EPEX Spot for an average Tuesday in 2013

increases as well and that the cost reduction is not very sensitive to whether the consumers use more electricity during the evening or during the day.

6.1 Datasets

CER Electricity Consumption Data: We use real electricity consumption data to generate consumer demand profiles. The data was gathered by the Irish Commission for Energy Regulation (CER) in the context of a smart metering study. In this study, the electricity consumption data of 485 small and medium enterprises and $4,225$ private households was collected over a period of about 1.5 years. The observation period is divided into time intervals of 30 minutes. One data point represents the average electricity consumption in kilowatts of one participant during a 30 minute interval. In particular, we used the data from 46 enterprises. The participating enterprises also answered questionnaires about, among others, their number of employees and typical hours of operation. We will refer to this data set as the `CER_data_set`[5].

EEX Electricity Prices Market Data: We generate the cooperative's electricity prices using real day-ahead market electricity prices, gathered from the European Energy Exchange (EEX), which is the leading energy exchange in Europe. In particular, we use the average hourly day-ahead prices from the EPEX Spot market from the 20 Tuesdays from January 1st 2013 to May 14th 2013. Figure 6 shows the average day-ahead spot market prices of the observation period. We will refer to the hourly market price data as the `EEX_data_set`[6].

6.2 Simulation Parameters

As of the model in the work of Mohsenian-Rad et al. (2010), we define each agent by its predetermined total electricity demand over the planning horizon, $\tau_i = \sum_{j=1}^{M} r_{ij}$, and by upper and lower bounds on the electricity demand in each time slot, i.e., $r_{ij} \in \left[\underline{r_{ij}}, \overline{r_{ij}}\right]$. Regarding the individual cost functions g_i, we assume $\mathrm{g}_i = 0$, because we do not have data on industry specific cost functions.

5. The data set is available at http://www.ucd.ie/issda/data/commissionforenergyregulation/.
6. The data is available at http://www.eex.com/en/Market%20Data.

6.2.1 MODELING THE COMPOSITION OF COOPERATIVES:

Agents can have a variety of different demand characteristics. To reflect this diversity, we use the `CER_data_set` to identify two classes of consumers with similar characteristics. To group consumers with similar characteristics, we clustered the participants of the study using k-means clustering based on the information from the questionnaires and the electricity consumption data. In particular, we used the consumption profiles of Tuesdays, because Tuesday is the most average work day. Here, Class 1 represents consumers consuming most of their electricity during the day, whereas Class 2 represents consumers having a stable consumption during the day, but have a higher consumption at night. Based on these classes, cooperatives with varying compositions can be simulated. The parameter $\mathbf{p\_fracAgent}$ defines the composition of the cooperative by specifying the fraction of agents from Class 1 (and Class 2 by $1 - p\_fracAgent$).

6.2.2 GENERATING THE NOMINAL DEMAND OF AGENTS:

An agent's nominal demand is the value around which the agent can vary its demand. We simulate scenarios with $M \in \{12, 24, 48\}$ time slots and use time intervals of 30, 60 and 120 minutes, which we generate by interpolating and averaging the data points from the two data sets. Then, $j = 1$ represents the first and $j = M$ the last time slot of the day. To get different agents, the nominal demand for each agent and time slot is the outcome x_{ij} of the random variable X_{Cj}. Let $\overline{\overline{y_{Cj}}}$ be the mean of the average demands of all participants in class C in time slot j and let $s\left(\overline{y}_{Cj}\right)$ be the corrected sample standard deviation of their average demands. The nominal demand is drawn from the uniform distribution $X_{Cj} \sim U\left[\overline{\overline{y_{Cj}}} - s\left(\overline{y}_{Cj}\right), \overline{\overline{y_{Cj}}} + s\left(\overline{y}_{Cj}\right)\right]$. The distributions of X_{Cj} for the two Classes are shown in Figure 7. The total demand of an agent for the whole day, is computed by $\tau_i = \sum_{j=1}^{M} x_{ij}$.

6.2.3 MODELING THE FLEXIBILITY OF CONSUMERS:

Agents can have different flexibilities of changing their demand profiles. Here, this ability is expressed through the agents' upper and lower bound constraints on their electricity demand in each time slot. The parameter $\mathbf{p\_flexShift}$ defines the flexibility of the agents by specifying the percentage, by which agents can vary their demands above and below their nominal demand. Thus, the larger $p\_flexShift$, the further apart are the upper and lower bounds. The parameter is the same for all agents and fixed for one scenario. Based on the flexibility and the nominal demand, the upper and lower bounds can be computed as

$$
\begin{aligned}
\underline{r_{ij}} &= x_{ij}\left(1 - p\_flexShift\right) \\
\overline{r_{ij}} &= x_{ij}\left(1 + p\_flexShift\right)
\end{aligned}
\tag{16}
$$

6.2.4 MODELING THE ELECTRICITY PRICE:

The cooperatives' electricity price function in each time slot is defined by the marginal price for the low load, p_j^L, and for the high load, p_j^H, and the price threshold, h_j, specifying at which demand levels the marginal prices apply. We use the prices from the `EEX_data_set` as exemplar prices to generate p_j^L and p_j^H. Let mp_j be the average spot market price from EPEX Spot in time slot j. We

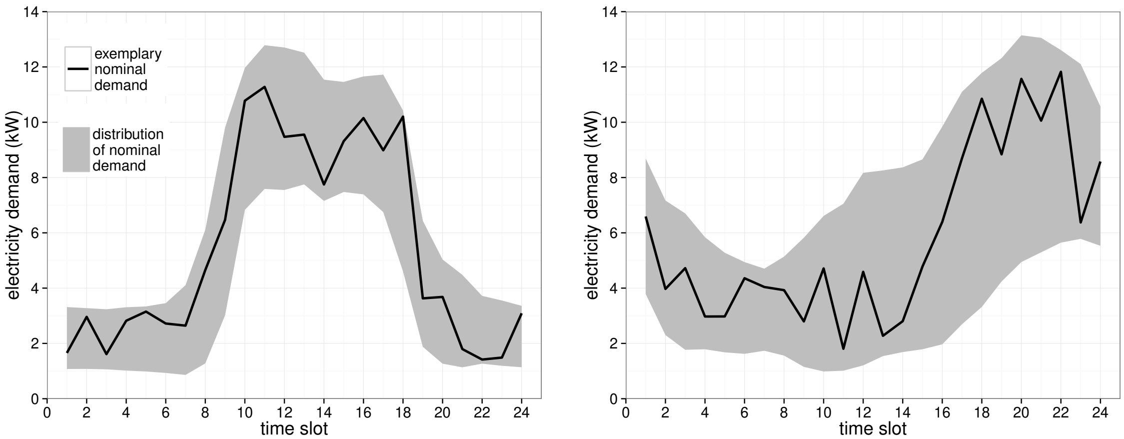

(a) Class 1 - day consumer

(b) Class 2 - night consumer

Figure 7: Distributions of the nominal demand for the two different consumer classes.

compute the marginal electricity prices as

$$\begin{aligned} p_j^L &= mp_j \\ p_j^H &= mp_j + \max_{x\in\{1,\ldots,M\}} \{mp_x\} - \min_{x\in\{1,\ldots,M\}} \{mp_x\} \end{aligned} \tag{17}$$

We keep these marginal prices fixed across all simulations. However, the price thresholds may vary across the different simulation scenarios. For example, the price threshold could be above the aggregated nominal demand so that the demand can be increased or it could be below so that the demand should be reduced. The parameter $\mathbf{p\_distThresh}$ defines the distance between the aggregated nominal demand and the price thresholds by specifying the percentage the thresholds are above or below the aggregated nominal demand. With a negative value of $p\_distThresh$, the thresholds lie below and with a positive value they lie above the aggregated nominal demand. For each scenario there is only one fixed value for all time slots. Figure 8a illustrates scenarios that have different values of $p\_distThresh$. In addition, the price thresholds could either be different in each time slot (e.g., following the profile of the aggregated nominal demand) or could be the same for each time slot (flat thresholds). Here, we compute flat thresholds by defining them as the moving average of the aggregated nominal demand. The parameter $\mathbf{p\_flatThresh}$ defines the flatness of the threshold by specifying the width of the interval over which the thresholds are flattened by a moving average. For $p\_flatThresh = 0$ the thresholds follow exactly the aggregated nominal demand, whereas for $p\_flatThresh = M$ the thresholds are the same in every time slot. Figure 8b illustrates scenarios with different values $p\_flatThresh$. The thresholds are computed by

$$h_j = \frac{1}{1 + 2p\_flatThresh} \sum_{j-p\_flatThresh}^{j+p\_flatThresh} \left((1 + p\_distThresh) \sum_i x_{ij} \right) \tag{18}$$

6.2.5 SIMULATION SCENARIOS:

To create different scenarios, we vary the input parameters as follows: $p\_fracAgent \in \{0, 0.25, 0.5, 0.75, 1\}$, $p\_flexShift \in \{0.1, 0.2, 0.3\}$, $p\_flatThresh \in \{0, 12, 24\}$, $p\_distThresh \in \{-0.2,$

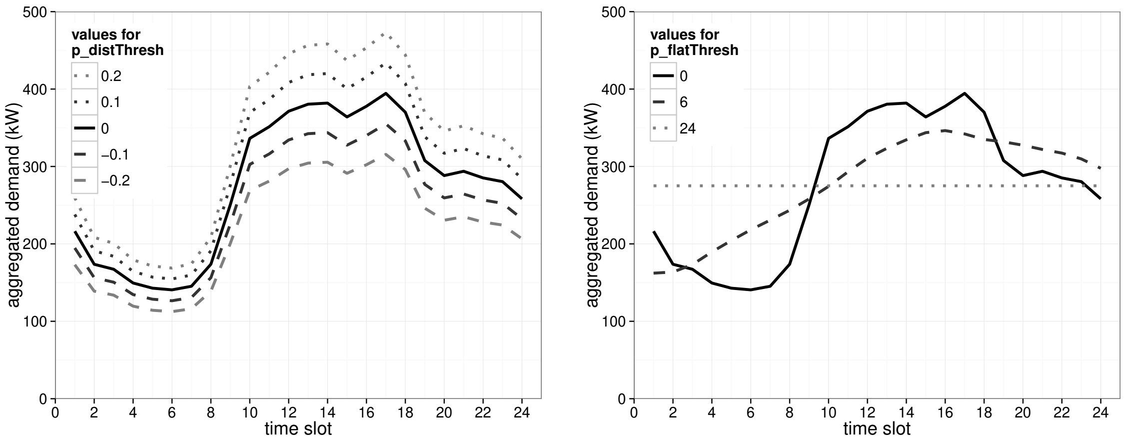

(a) Price thresholds for different values of $p\_distThresh$. (b) Price thresholds for different values of $p\_flatThresh$.

Figure 8: Price thresholds for different values of the parameters $p\_distThresh$ and $p\_flatThresh$.

$-0.1, 0, 0.1, 0.2\}$, $\#agents \in \{20, 40, 60, 80, 100\}$, $M \in \{12, 24, 48\}$ and $\varepsilon \in \{0.5, 1, 2\}$. This leads to 3375 different scenarios. We define the algorithm to be converged, when the cost reduction in one iteration gets less than 0.00001%, i.e., $\mathrm{C}(\mathbf{R}) / \mathrm{C}(\mathbf{R}') < 1.0000001$.

6.3 Simulation Results

We first present the effects of the different parameters on the cost reduction through the algorithm and subsequently discuss the convergence time.

Definition 3. *The **cost reduction** is the difference between the cost of the uncoordinated profile* C_0 *(chosen by each agent when optimizing its cost according to the market price) and the coordinated profile* C^* *as a percentage of the cost of the uncoordinated profile, i.e.,* $\mathrm{CR} = \frac{\mathrm{C}_0 - \mathrm{C}^*}{\mathrm{C}_0} \cdot 100$.

Figure 9 shows a subset of the results for $p\_fracAgent \in \{0, 0.5, 1\}$, $M = 24$, $\#agents = 40$ and $\varepsilon = 1$. Figure 9a shows the cost reduction in cooperatives only consisting of consumers mainly consuming electricity at night. Figure 9b shows the cost reductions in cooperatives that consist of both consumer classes in equal proportions. Finally, Figure 9c shows the cost reductions in cooperatives consisting of only consumers with their main demand during the day. The x-axis represents the flexibility of shifting demand, the y-axis the flatness of the thresholds and the z-axis the height of the thresholds. In all figures are scenarios with high cost reduction (white) and scenarios with less cost reduction (dark red). Over all scenarios, we observed a mean cost reduction of 2.57%, with results varying from 0% to 7.44%. However, the sampled cases represent all combinations of the input parameters and we cannot say how realistic individual settings are. For example, in many scenarios no optimization is possible and we put in a cost reduction of 0%. No optimization is possible, if the price threshold provided by the supplier is higher than the sum of the agents upper bound constraints or if the price threshold is lower than the sum of the lower bound constraints. Not taking into account those scenarios, we observe a mean of 2.96%. Moreover, looking at the cost reduction from the initial nominal demand profile, we get a mean of 3.7%, with results up to 11.1%.

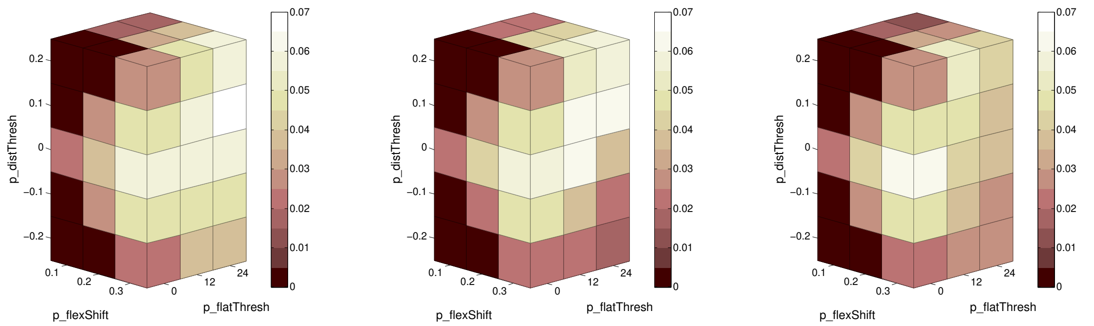

(a) Only consumers who mainly consume at night, $p\_fracAgent = 0$.

(b) Consumers from both classes have equal fractions, $p\_fracAgent = 0.5$

(c) Only consumers who mainly consume during day, $p\_fracAgent = 1$.

Figure 9: The cost reduction through the algorithm as a function of the four input parameters: $p\_fracAgent$, $p\_flatThresh$, $p\_distThresh$ and $p\_flexShift$. Scenarios with high cost reduction are white and scenarios with less cost reduction are dark red.

6.3.1 SENSITIVITY ANALYSIS

The purpose of this sensitivity analysis is to understand the effect of the different input parameters on the cost reduction through our algorithm. For the sensitivity analysis we perform a multiple linear regression. As dependent variable (criterion) we choose the cost reduction (Definition 3). As independent variables (predictors) we choose the smallest interpretable model that provides a good adjusted R-squared. The resulting model is shown in Table 2. The model explains 76.36% of the variance with an adjusted R-squared of 0.7571. The remaining variance cannot be explained by the model, because the agents' demand constraints are generated randomly.

predictors	coefficient	stand. error	p-value
$p\_fracAgent$	0.0019	0.0011	0.0910
$p\_flatThresh$	0.0077	0.0017	1.66E-05
$p\_flexShift$	0.0114	0.0010	1.35E-25
$\lvert p\_distThresh \rvert$	-0.0072	0.0008	2.08E-16
$p\_distThresh * p\_flexShift$	0.0018	0.0002	6.67E-17
$p\_fracAgent * p\_flatThresh$	-0.0016	0.0005	0.0035

Table 2: Multiple linear regression model; adjusted $R^2 = 0.7571$

Result 1: As agents' flexibility of shifting their demands increases, cost reduction increases, i.e., $p\_flexShift \nearrow \Rightarrow G \nearrow$.

The multiple linear regression model summarized in Table 2 shows that the effect of $p\_flexShift$ is highly significant with $p < 0.01$. The positive value of the coefficient (0.0114) shows that an increase of the flexibility leads to increased cost reduction.

Intuitively, with higher flexibility, more demand from time slots with high prices can be shifted to those with lower prices. Consequently, higher cost reduction can be achieved. For example, if $p\_flexShift = 0$, no demand can be shifted and no cost reduced. However, another reason for the high cost reduction is that the uncoordinated profile before the algorithm has higher cost. Recall that for the uncoordinated demand profile the agents optimize their demand according to the hourly

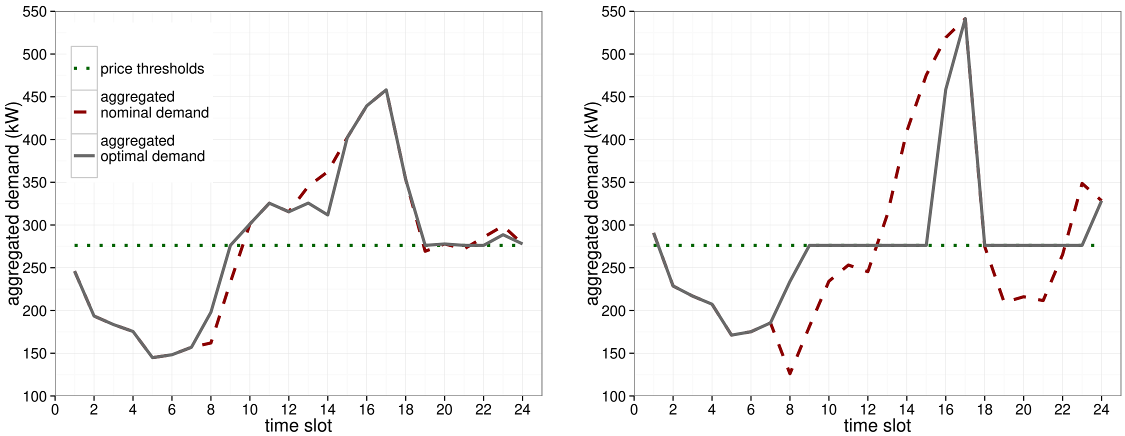

(a) Lower flexibility, $p\_flexShift = 0.1$.

(b) Higher flexibility, $p\_flexShift = 0.3$.

Figure 10: Example scenarios with different flexibilities for the agents of shifting their demands.

market prices. Consequently, the higher the agents' flexibility, the more they adapt their demand schedules to the hourly prices. However, since not being coordinated, all agents shift as much demand as possible to cheap time slots. This leads to load synchronization, so that the total demand in the former cheap slots exceeds the thresholds, leading to higher costs (*herding behavior*). Thus, high flexibility of uncoordinated consumers leads to highly synchronized demand. These high initial costs allow for more cost reduction through the algorithm. Figure 10 illustrates settings with low ($p\_flexShift = 0.1$) and high ($p\_flexShift = 0.3$) flexibility. The figure shows clearly that the initial demand peak is higher in the scenario with more flexibility. Recall from Figure 6 that slots 14, 15, 16 and 17 are relatively low priced slots with 17 being the lowest. The area between the red dashed curve of the uncoordinated demand profile and the coordinated demand profile in gray multiplied by the respective marginal prices shows the cost reduction through coordination. With the increasing freedom in Figure 10b it can be seen that the area between the two curves increases. The reason the height of the peak stays the same, is that at first the agents shift as much demand as possible from time slots with aggregated demands above the thresholds (peaks) to time slots with demand levels below the thresholds. When all agents reach their upper bounds in those time slots, they shift the remaining demand to the time slots with the lowest prices above the threshold. According to the hourly prices, the cheapest remaining time slot is 17. Since the freedom allows demanding most of the remaining electricity in that time slot, the height of the peak does not change.

Result 2: As the absolute distance of the demand thresholds to the aggregated load profile decreases, cost reduction increases, i.e., $|p\_distThresh| \searrow \Rightarrow G \nearrow$.

The multiple linear regression model summarized in Table 2 shows that the effect of $|p\_distThresh|$ is highly significant with $p < 0.01$. The negative coefficient (-0.0072) shows that the further away the thresholds are from the aggregated load profile the less cost reduction can be achieved.

Figure 11 illustrates settings with low ($p\_distThresh = -0.2$), normal ($p\_distThresh = 0$) and high ($p\_distThresh = 0.2$) thresholds. The cost reduction is limited by how much the agents can decrease demand above and increase below the thresholds. The demand above the thresholds is represented by the area between the initial demand profile (red dashed line) and the thresholds (yellow flat line) above the thresholds, and the unused demand below the thresholds by the area between the two lines below the thresholds. These amounts are also affected by the flatness of

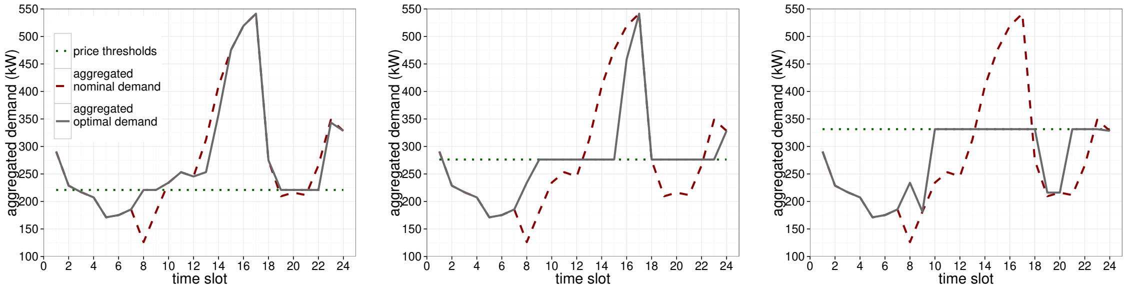

(a) Thresholds below the average load, $p\_distThresh = -0.2$.
(b) Thresholds at the average load, $p\_distThresh = 0$.
(c) Thresholds above the average load, $p\_distThresh = 0.2$.

Figure 11: Example scenarios with different heights of the price thresholds.

the thresholds, i.e., the flatter the thresholds, the larger the two areas become. Thus, with flatter thresholds the cost reduction increases as well. The regression model supports this, because the parameter $p\_flatThresh$ has a positive coefficient (0.0077) and is highly significant with $p < 0.01$.

However, the cost reduction in the scenario with high thresholds in Figure 11c is greater than in the case of normal thresholds in Figure 11b. That happens, because the flexibility of shifting demand is a limiting factor. Since the absolute distance between the upper and lower bounds is smaller in time slots with low demand, the demand hits the upper bounds in those time slots before it reaches the lower bounds in the time slots with high demand. In this situation, an increment of the thresholds leads to more time slots below the thresholds and thus more possibilities to shift demand. Consequently, if the flexibility is a limiting factor, the cost reduction increases, if the thresholds are slightly increased. This is supported by the regression model, because the interaction term $p\_distThresh * p\_flexShift$ is highly significant with $p < 0.01$.

Result 3: We observe similar cost reductions in groups with mostly consumers from one class and groups with agents from both classes in similar fractions.

The multiple linear regression model in Table 2 shows that the parameter $p\_fracAgent$ is not a significant predictor for the cost reduction with $p > 0.05$. This supports the observation that we can see cost reductions from the coordinated behavior in each of the two classes (day consumers and night consumers) as well as mixed groups. Figure 12 illustrates settings with only night consumers (Figure 12a), a mixed group (Figure 12b), and only day consumers (Figure 12c). However, in these examples, with an increasing fraction of day consumers, the cost reduction decreases to a larger extent than indicated by the parameter $p\_fracAgent$ in the regression model. This effect is due to an interaction effect between the composition of the cooperative and the flatness of the thresholds. The multiple linear regression model shows that the interaction term $p\_fracAgent * p\_flatThresh$ is a significant predictor with $p < 0.01$. The negative value of the coefficient (-5.219e-04) shows that as the flatness of the thresholds increases and the fraction of day consumers also increases, cost reduction decreases. In Figure 12 it can be seen that the thresholds are lying outside of the agents' demand bounds in many time slots, because the thresholds are the same in each time slot. In this situation not the thresholds, but the demand bounds limit the potential cost reduction through coordination. Consequently, as the thresholds get flatter, they get out of reach of the agents' demand bounds. Since the load profiles of day consumers have a more significant difference between peak and low load, this effect is stronger for groups with a larger fraction of day consumers.

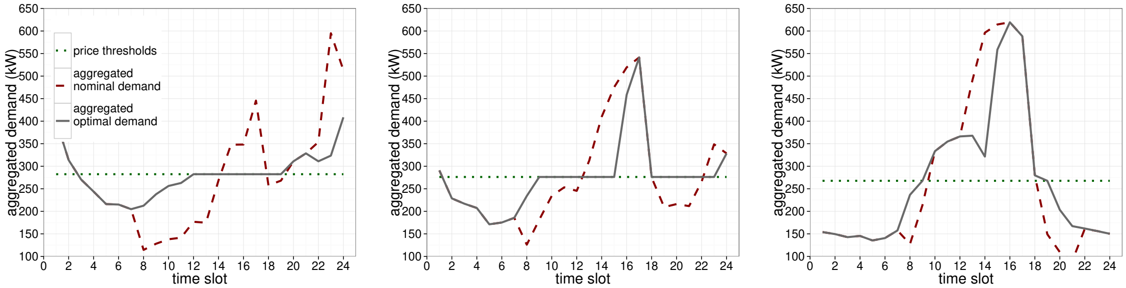

(a) Only consumers who mainly consume at night, $p\_fracAgent = 0$.

(b) Consumers from both classes have equal fractions, $p\_fracAgent = 0.5$.

(c) Only consumers mainly consuming during the day, $p\_fracAgent = 1$.

Figure 12: Example scenarios with different compositions of the consumer cooperative.

6.3.2 CONVERGENCE PROPERTIES OF THE ALGORITHM

Let the first iteration be the submission of the agents' uncoordinated initial demand profiles. Then, in the second iteration the coordinator sends the first virtual price signals to the agents. Recall that in the beginning the basic algorithm is performed and after the basic algorithm converged, the additional phase is started, if necessary. Also recall that we defined the algorithm to have converged, when the cost reduction in one iteration gets less than 0.00001%, i.e., $\mathrm{C}\left(\mathbf{R}\right) / \mathrm{C}\left(\mathbf{R}'\right) < 1.0000001$. In our simulations, the basic algorithm converges in average within 9.57 iterations. In the cases that need the additional phase, the general algorithm converges in average in additional 35.10 iterations. The experiment was setup using one HP Pavilion dmt4-1000 with an Intel Core i5 520M with two cores operating at 2.4GHz and 8 Gigabyte of DDR3 physical memory. Using this setup, one iteration of the basic algorithm took on average 25.12 seconds and one iteration of the extended phase took on average 52.13 seconds. Looking at the agents demand profiles at convergence, we observe that for each agent in average in 67.6% of the time slots the demand bounds are tight. In all sampled scenarios, there is at least one time slot, where the demand bound is not tight, because the agents have to satisfy their constant total demand over the whole planning horizon.

Definition 4. *The **convergence accuracy** is the difference between the cost of the converged solution and the optimal solution as percentage of the difference between the cost of the uncoordinated profile and the optimal solution, i.e.,* $\frac{\mathrm{C}(\mathbf{R}^T)-\mathrm{C}(\mathbf{R}^*)}{\mathrm{C}(\mathbf{R}^1)-\mathrm{C}(\mathbf{R}^*)} \cdot 100$.

Result 4: In our simulations, the basic algorithm achieves the optimal solution in 44.9% of all sampled scenarios.

We found that in 44.9% of all sampled cases the basic algorithm achieves the optimal solution. In the remaining scenarios the basic algorithm comes in average 1.06% close to the optimal solution. This indicates that in most cases the basic algorithm can provide good results.

Result 5: Our simulations indicate that larger step sizes, ε, lead to faster convergence at the cost of small reductions in accuracy.

The parameter ε defines the step size of the extended phase. Intuitively, larger steps lead to a faster convergence, at a cost of reduced accuracy. Table 3 shows the convergence accuracy and the average number of iterations until convergence for varying $\varepsilon \in \{0.5, 1, 2\}$. The number of agents and time slots are fixed as $N = 40$ and $M = 24$. The results indicate that as ε increases, the number of iterations until convergence as well as the accuracy decreases. However, the reduction in accuracy is very small and the algorithm achieves good results for all ε. Our results indicate that in order to achieve accurate solutions within few iterations, *the step size should be adapted at run-time*. In particular, the value of ε should be reduced over time.

ε	Convergence accuracy	Number of iterations until convergence
0.5	0.22%	36.5
1	0.33%	26.1
2	0.55%	17.9

Table 3: Convergence accuracy and time for varying step sizes

Result 6: Our simulations indicate that the convergence time scales linearly with the agent population size and number of time slots.
To analyze the influence on the convergence time of the algorithm, we vary the number of agents from 20 to 100 and time slots from 12 to 48. Table 4a shows the convergence accuracy and Table 4b the number of iterations until convergence. The step size is fixed at $\varepsilon = 1$. The results indicate that the number of iterations until convergence increases linearly with the number of agents and time slots. The accuracy does not change much, as we increase the agent population and time slots.

agents \ time slots	12	24	48
20	0.21%	0.19%	0.18%
40	0.24%	0.29%	0.20%
60	0.23%	0.26%	0.28%
80	0.26%	0.26%	0.30%
100	0.25%	0.28%	0.38%

(a) Convergence accuracy

agents \ time slots	12	24	48
20	9.8	16	22.4
40	12.4	20.1	30.2
60	13.5	23.7	35.4
80	15.2	27.1	40.3
100	16.9	28.2	43.2

(b) Number of iterations until convergece

Table 4: Convergence accuracy and time for varying number of agents and time slots

7. Conclusion and Future Work

In this paper, we presented an iterative coordination algorithm to minimize the energy cost of a consumer cooperative, given that information about the agents' individual demand constraints and preferences remains private. We proved that the algorithm converges to the optimal demand schedule and presented results on the time complexity of the iterative algorithm and the agents' incentive compatibility. Additionally, we conducted evaluations on the algorithm using multiagent simulation based on real world consumption data. Through simulations, we characterized the convergence properties of our algorithm and the effect of differing demand characteristics in the cooperative and price functions on the cost reduction through our algorithm. The results show that the convergence time scales linearly with the population size and length of the optimization horizon. Finally, we observe that as participants' flexibility of shifting their demands increases, cost reduction increases and that the cost reduction is not very sensitive to variation in consumption patterns of the consumers.

This work can be extended in several directions. Future work can investigate settings in which the agents might not be able to compute a guaranteed optimal solution of their individual problem, but only a provably good approximation. This could apply to settings with more detailed load models for the agents. The overall demand can come from two types of loads: shiftable loads and non-shiftable loads. These loads can be divided further into interruptible and non-interruptible loads. In addition, these loads can be subject to temporal constraints. This can lead to the situation where the individual problems for the agents are no longer convex, and thus no agent can solve its individual problem optimally. Luo, Chakraborty, and Sycara (2013) present a distributed iterative algorithm for the generalized task assignment problem in the context of a multirobot system (MR-GAP). Based on the (approximate) best responses from the agents, this algorithm has a provable approximate ratio. It would be interesting to investigate such a distributed algorithm in the context of this problem.

In this paper, the demand of the cooperative in each time slot solely consists of the aggregated demand of the agents. Future work can consider problems with generation and/ or storage (that can be centralized, i.e., owned by the cooperative, or distributed, i.e., owned by an individual agent). Another avenue of future work is to consider a problem formulation where the cooperative faces uncertainty in electricity prices. For example, consider a 24-hour planning horizon and instead of a long term contract the electricity is bought from an hourly spot market. Here, for scheduling demand, one only knows the price for the next hour and the prices for the future hourly time slots are uncertain. The spot market electricity price depends on many factors that are not controlled by the coordinator. Hence, for planning purposes, the prices can be assumed to be an externally specified stochastic process. Under this assumption, the goal would be to design algorithms for (1) determining policies (for generation, storage, and price signals to be sent to the firms) for the central coordinator and (2) determine the schedules for the individual firms, such that the expected cost of buying electricity is minimized.

Acknowledgements

This work was partially supported by NSF award IIS-1218542.

Appendix A. Proof: Virtual Cost is an Upper Bound on the Total Cost

Proof. Here, we give the full proof showing that the sum of the agents' individual cost according to their virtual price signals is an upper bound on the total central cost at market prices. We prove this, by showing with Equations 1 and 7 that for every time slot j the difference between the total cost for the aggregated demand and the sum of the agents' individual cost is lower or equal to 0. Since it is true for every time slot, it is also true for the sum over all time slots.

The central cost for the aggregated demand is given in Equation 1 as

$$
\begin{aligned}
&\sum_{i=1}^{N} \mathrm{p}_j\left(\rho'_j\right) r'_{ij},\ \forall j \\
&= p_j^H\left(\rho'_j - h_j\right)^+ + p_j^L\left(\rho'_j - h_j\right)^- + p_j^L h_j \\
&= \begin{cases} p_j^H\left(\rho'_j - h_j\right) + p_j^L h_j & \rho'_j > h_j \\ p_j^L\left(\rho'_j - h_j\right) + p_j^L h_j & \rho'_j \le h_j \end{cases} \\
&= \begin{cases} \sum_{i=1}^{N}\left(p_j^H\left(r'_{ij} - h_{ij}\right)\right) + \sum_{i=1}^{N}\left(p_j^L h_{ij}\right) & \rho'_j > h_j \\ \sum_{i=1}^{N}\left(p_j^L\left(r'_{ij} - h_{ij}\right)\right) + \sum_{i=1}^{N}\left(p_j^L h_{ij}\right) & \rho'_j \le h_j \end{cases}
\end{aligned}
$$

The sum of the agents' individual cost is given in Equation 7 as

$$\begin{aligned}&\sum_{i=1}^{N} \mathrm{s}_{ij}\left(r'_{ij}\right) r'_{ij},\ \forall j\\ &= \sum_{i=1}^{N}\left[p_j^H\left(r'_{ij}-h_{ij}\right)^+ + p_j^L\left(r'_{ij}-h_{ij}\right)^- + p_j^L h_{ij}\right]\end{aligned}$$

For the difference of the two costs we get:

$$\begin{aligned}&\sum_{i=1}^{N} \mathrm{p}_j\left(\rho'_j\right) r'_{ij} - \sum_{i=1}^{N} \mathrm{s}_{ij}\left(r'_{ij}\right) r'_{ij},\ \forall j\\ &= \begin{cases}\sum_{i=1}^{N}\left[p_j^H\left(r'_{ij}-h_{ij}\right) + p_j^L h_{ij} - p_j^H\left(r'_{ij}-h_{ij}\right)^+ - p_j^L\left(r'_{ij}-h_{ij}\right)^- - p_j^L h_{ij}\right] & \rho'_j > h_j\\ \sum_{i=1}^{N}\left[p_j^L\left(r'_{ij}-h_{ij}\right) + p_j^L h_{ij} - p_j^H\left(r'_{ij}-h_{ij}\right)^+ - p_j^L\left(r'_{ij}-h_{ij}\right)^- - p_j^L h_{ij}\right] & \rho'_j \le h_j\end{cases}\end{aligned}$$

Since $\sum_{i=1}^{N}\left(p_j^L h_{ij}\right)$ and $-\sum_{i=1}^{N}\left(p_j^L h_{ij}\right)$ equal each other out they and can be removed.

$$= \begin{cases}\sum_{i=1}^{N}\left[p_j^H\left(r'_{ij}-h_{ij}\right) - p_j^H\left(r'_{ij}-h_{ij}\right)^+ - p_j^L\left(r'_{ij}-h_{ij}\right)^-\right] & \rho'_j > h_j\\ \sum_{i=1}^{N}\left[p_j^L\left(r'_{ij}-h_{ij}\right) - p_j^H\left(r'_{ij}-h_{ij}\right)^+ - p_j^L\left(r'_{ij}-h_{ij}\right)^-\right] & \rho'_j \le h_j\end{cases}$$

Lets now write $\sum_{i=1}^{N} p_j^H\left(r'_{ij}-h_{ij}\right)^+$ as $\sum_{i:\ r'_{ij}>h_{ij}}^{N} p_j^H\left(r'_{ij}-h_{ij}\right)$ and $\sum_{i=1}^{N} p_j^L\left(r'_{ij}-h_{ij}\right)^-$ analog as $\sum_{i:\ r'_{ij}\le h_{ij}}^{N} p_j^H\left(r'_{ij}-h_{ij}\right)$. Lets also write $\sum_{i=1}^{N}$ as $\sum_{i:\ r'_{ij}>h_{ij}}^{N} + \sum_{i:\ r'_{ij}\le h_{ij}}^{N}$.

$$= \begin{cases}\sum\limits_{i:\ r'_{ij}>h_{ij}}^{N}\left[p_j^H\left(r'_{ij}-h_{ij}\right) - p_j^H\left(r'_{ij}-h_{ij}\right)\right] + \sum\limits_{i:\ r'_{ij}\le h_{ij}}^{N}\left[p_j^H\left(r'_{ij}-h_{ij}\right) - p_j^L\left(r'_{ij}-h_{ij}\right)\right] & \rho'_j > h_j\\ \sum\limits_{i:\ r'_{ij}>h_{ij}}^{N}\left[p_j^L\left(r'_{ij}-h_{ij}\right) - p_j^H\left(r'_{ij}-h_{ij}\right)\right] + \sum\limits_{i:\ r'_{ij}\le h_{ij}}^{N}\left[p_j^L\left(r'_{ij}-h_{ij}\right) - p_j^L\left(r'_{ij}-h_{ij}\right)\right] & \rho'_j \le h_j\end{cases}$$

For $\rho'_j > h_j$ the terms $\sum_{i:\ r'_{ij}>h_{ij}}^{N} p_j^H\left(r'_{ij}-h_{ij}\right)$ and $-\sum_{i:\ r'_{ij}>h_{ij}}^{N} p_j^H\left(r'_{ij}-h_{ij}\right)$ equal each other out and can be removed. For $\rho'_j \le h_j$ remove $\sum_{i:\ r'_{ij}\le h_{ij}}^{N} p_j^L\left(r'_{ij}-h_{ij}\right)$ respectively.

$$= \begin{cases}\sum_{i:\ r'_{ij}\le h_{ij}}^{N}\left[p_j^H\left(r'_{ij}-h_{ij}\right) - p_j^L\left(r'_{ij}-h_{ij}\right)\right] & \rho'_j > h_j\\ \sum_{i:\ r'_{ij}>h_{ij}}^{N}\left[p_j^L\left(r'_{ij}-h_{ij}\right) - p_j^H\left(r'_{ij}-h_{ij}\right)\right] & \rho'_j \le h_j\end{cases}$$

Factor $\left(p_j^H - p_j^L\right)$ and $\left(p_j^L - p_j^H\right)$ out.

$$\begin{aligned}&= \begin{cases}\sum_{i:\ r'_{ij}\le h_{ij}}^{N}\left(p_j^H - p_j^L\right)\left(r'_{ij}-h_{ij}\right) & \rho'_j > h_j\\ \sum_{i:\ r'_{ij}>h_{ij}}^{N}\left(p_j^L - p_j^H\right)\left(r'_{ij}-h_{ij}\right) & \rho'_j \le h_j\end{cases}\\ &= \begin{cases}\sum_{i=1}^{N}\left(p_j^H - p_j^L\right)\left(r'_{ij}-h_{ij}\right)^- & \rho'_j > h_j\\ \sum_{i=1}^{N}\left(p_j^L - p_j^H\right)\left(r'_{ij}-h_{ij}\right)^+ & \rho'_j \le h_j\end{cases}\\ &= \begin{cases}\left[\sum_{i=1}^{N}\left(p_j^H - p_j^L\right)\left(r'_{ij}-h_{ij}\right)^-\right] \le 0 & \rho'_j > h_j\\ \left[\sum_{i=1}^{N}\left(p_j^L - p_j^H\right)\left(r'_{ij}-h_{ij}\right)^+\right] \le 0 & \rho'_j \le h_j\end{cases}\end{aligned}$$

Since this inequality is true for every time slot, it also holds for the sum over all time slots. It follows that the sum of the agents' individual cost according to their virtual price signals is an upper bound on the total central cost at market prices, $\mathrm{C}\left(\mathbf{R}'\right) \leq \sum_{i=1}^{N} \mathrm{C}_i\left(\mathbf{r}'_i\right)$:

$$\sum_{i=1}^{N}\sum_{j=1}^{M} \mathrm{p}_j\left(\rho'_j\right) r'_{ij} + \mathrm{g}_i(\mathbf{r}'_i) \leq \sum_{i=1}^{N}\sum_{j=1}^{M} \mathrm{s}_{ij}\left(r'_{ij}\right) r'_{ij} + \mathrm{g}_i(\mathbf{r}'_i)$$

□

Appendix B. Proof by Counterexample: Basic Algorithm Can Converge to a Suboptimal Solution in Case of $g_i(\cdot) \neq 0$

We present a counterexample to prove that the basic algorithm can converge to a suboptimal solution in general settings with $g_i(\cdot) \neq 0$. Consider a population of 2 agents, $N = 2$, and a planning horizon of 2 time slots, $M = 2$. The agents' constraints are such that in the converged solution, the aggregated demand is equal to the threshold in one time slot. Let the price function be given as:

$$(p_1^L, p_1^H) = (3, 8),\ h_1 = 9\ ;\ (p_2^L, p_2^H) = (3, 8),\ h_2 = 11$$

The individual constraints on the agents' demand are:

$$r_{11} \in [1, 3],\ r_{12} \in [4, 6],\ r_{11} + r_{12} = \tau_1 = 7$$
$$r_{21} \in [4, 6],\ r_{22} \in [4, 6],\ r_{21} + r_{22} = \tau_2 = 10$$

The agents' individual cost associated with the demand schedule is given as:

$$\mathrm{g}_1\left(\mathbf{r}_1\right) = \mathbf{r}_1 \begin{bmatrix} 5 \\ 1 \end{bmatrix},\ \mathrm{g}_2\left(\mathbf{r}_2\right) = \mathbf{r}_2 \begin{bmatrix} 6 \\ 3 \end{bmatrix}$$

At the beginning, the agents compute their initial demand profiles based on the market prices $\mathbf{r}_1^{(1)} = (1, 6)$, $\mathbf{r}_2^{(1)} = (4, 6)$. The cost based on these demand profiles is $\mathrm{C}\left(\mathbf{R}^{(1)}\right) = 109$.

At convergence, the final profiles of the two agents are $\mathbf{r}_1^{(T)} = (1.5, 5.5)$, $\mathbf{r}_2^{(T)} = (4.5, 5.5)$. The cost based on these demand profiles is $\mathrm{C}\left(\mathbf{R}^{(T)}\right) = 107.5$.

However a different demand profile of the agents exists with $\mathbf{r}'_1 = (1, 6)$, $\mathbf{r}'_2 = (5, 5)$. This profile is feasible and leads to lower total cost, i.e., $\mathrm{C}\left(\mathbf{R}'\right) = 107$. It follows that the algorithm has stopped in a suboptimal solution.

References

Albadi, M., & El-Saadany, E. (2007). Demand response in electricity markets: An overview. In *Power Engineering Society General Meeting, 2007. IEEE*, pp. 1–5. IEEE.

Atzeni, I., Ordonez, L., Scutari, G., Palomar, D., & Fonollosa, J. (2013). Demand-side management via distributed energy generation and storage optimization. *IEEE Transactions on Smart Grid*, *4*(2), 866–876.

Ausubel, L. M., & Milgrom, P. (2006). The lovely but lonely vickrey auction. *Combinatorial auctions*, 17–40.

Barbose, G., Goldman, C., & Neenan, B. (2004). A survey of utility experience with real time pricing. Tech. rep., Ernest Orlando Lawrence Berkeley National Laboratory, Berkeley, CA, USA.

Bertsekas, D. P., & Tsitsiklis, J. N. (1989). *Parallel and distributed computation: numerical methods*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.

Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press.

Chu, C., & Jong, T. (2008). A novel direct air-conditioning load control method. *IEEE Transactions on Power Systems*, *23*(3), 1356–1363.

Clement-Nyns, K., Haesen, E., & Driesen, J. (2010). The impact of charging plug-in hybrid electric vehicles on a residential distribution grid. *IEEE Transactions on Power Systems*, *25*(1), 371–380.

Dietrich, K., Latorre, J., Olmos, L., & Ramos, A. (2012). Demand response in an isolated system with high wind integration. *IEEE Transactions on Power Systems*, *27*(1), 20–29.

Green, J., & Laffont, J.-J. (1977). Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica: Journal of the Econometric Society*, 427–438.

Hurwicz, L. (1975). On the existence of allocation systems whose manipulative nash equilibria are pareto-optimal. In *3rd World Congress of the Econometric Society*.

Jellings, C. W., & Chamberlin, J. H. (1993). *Demand Side Management: Concepts and Methods*. PennWell Books, Tulsa, OK, USA.

Kirschen, D. (2003). Demand-side view of electricity markets. *IEEE Transactions on Power Systems*, *18*(2), 520–527.

Luo, L., Chakraborty, N., & Sycara, K. (2013). Distributed algorithm design for multi-robot generalized task assignment. In *Proceedings of IEEE International Conference on Intelligent Robots and Systems*.

Medina, J., Muller, N., & Roytelman, I. (2010). Demand response and distribution grid operations: Opportunities and challenges. *IEEE Transactions on Smart Grid*, *1*(2), 193–198.

Mohsenian-Rad, A., Wong, V., Jatskevich, J., Schober, R., & Leon-Garcia, A. (2010). Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *IEEE Transactions on Smart Grid*, *1*(3), 320–331.

Mohsenian-Rad, A., & Leon-Garcia, A. (2010). Optimal residential load control with price prediction in real-time electricity pricing environments. *IEEE Transactions on Smart Grid*, *1*(2), 120–133.

Nguyen, H. K., Song, J. B., & Han, Z. (2012). Demand side management to reduce peak-to-average ratio using game theory in smart grid. In *IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pp. 91–96. IEEE.

Palensky, P., & Dietrich, D. (2011). Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE Transactions on Industrial Informatics*, *7*(3), 381–388.

Pedrasa, M., Spooner, T., & MacGill, I. (2010). Coordinated scheduling of residential distributed energy resources to optimize smart home energy services. *IEEE Transactions on Smart Grid*, *1*(2), 134–143.

Philpott, A., & Pettersen, E. (2006). Optimizing demand-side bids in day-ahead electricity markets. *IEEE Transactions on Power Systems*, *21*(2), 488–498.

Rahimi, F., & Ipakchi, A. (2010). Demand response as a market resource under the smart grid paradigm. *IEEE Transactions on Smart Grid*, *1*(1), 82–88.

Ramchurn, S., Vytelingum, P., Rogers, A., & Jennings, N. (2012). Putting the 'smarts' into the smart grid: A grand challenge for artificial intelligence. *Communications of the ACM*, *55*(4), 86–97.

Ramchurn, S. D., Vytelingum, P., Rogers, A., & Jennings, N. (2011). Agent-based control for decentralised demand side management in the smart grid. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pp. 5–12.

Samadi, P., Mohsenian-Rad, H., Wong, V., & Schober, R. (2013). Tackling the load uncertainty challenges for energy consumption scheduling in smart grid. *IEEE Transactions on Smart Grid*, *4*(2), 1007–1016.

Tanaka, K., Uchida, K., Ogimi, K., Goya, T., Yona, A., Senjy, T., Funabashi, T., & Kim, C. (2011). Optimal operation by controllable loads based on smart grid topology considering insolation forecasted error. *IEEE Transactions on Smart Grid*, *2*(3), 438–444.

Veit, A., Xu, Y., Zheng, R., Chakraborty, N., & Sycara, K. (2013). Multiagent coordination for energy consumption scheduling in consumer cooperatives. In *Proceedings of 27th AAAI Conference on Artificial Intelligence*, pp. 1362–1368.

Voice, T., Vytelingum, P., Ramchurn, S., Rogers, A., & Jennings, N. (2011). Decentralised control of micro-storage in the smart grid. In *Proceedings of 25th AAAI Conference on Artificial Intelligence*, pp. 1421–1427.

Vytelingum, P., Ramchurn, S., Rogers, A., & Jennings, N. (2010). Agent-based homeostatic control for green energy in the smart grid. In *First International Workshop on Agent Technologies for Energy Systems (ATES 2010)*.

Vytelingum, P., Voice, T., Ramchurn, S., Rogers, A., & Jennings, N. (2011). Theoretical and practical foundations of large-scale agent-based micro-storage in the smart grid. *Journal of Artificial Intelligence Research*, *42*(1), 765–813.

Wood, A., & Wollenberg, B. (1996). *Power Generation Operation and Control*. John Wiley & Sons.

Wu, C., Mohsenian-Rad, H., & Huang, J. (2012). Wind power integration via aggregator-consumer coordination: A game theoretic approach. In *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, pp. 1–6. IEEE.