Demand Side Energy Management via Multiagent Coordination in Consumer Cooperatives

Andreas Veit[⋆] **Ying Xu**[†] **Ronghuo Zheng**[†] **Nilanjan Chakraborty**[‡] **Katia Sycara**[‡]

[⋆]Karlsruhe Institute of Technology
Kaiserstrasse 12
76131 Karlsruhe, Germany.

[†]Tepper School of Business
Carnegie Mellon University
Pittsburgh, PA 15213

[‡]Robotics Institute, School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

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 renewable supply. In this paper, we propose a partially-centralized organization of consumers (or agents), namely, a consumer cooperative for purchasing electricity from the market. We propose a novel multiagent coordination algorithm, to shape the energy consumption of the cooperative. In the cooperative, a central coordinator buys the electricity for the whole group and consumers make their own consumption decisions, based on their private consumption constraints and preferences. To coordinate individual consumers under incomplete information, we propose an iterative algorithm, in which a virtual price signal is sent by the coordinator to induce consumers to shift their demands when required. This algorithm provably converges to the central optimal solution and minimizes the electric energy cost of the cooperative. Additionally, we perform simulations based on real world consumption data to characterize (a) the convergence properties of our algorithm and (b) understand the effect of different parameters that characterize the electricity consumption profile on the potential cost reduction through coordination by our algorithm. The results show that as the participants' flexibility of shifting their demands increases, cost reduction increases. We also observe that the cost reduction is not very sensitive to the variation in consumption patterns of the consumers (e.g., whether the consumers use more electricity during the evening or during the day). Finally, our simulations indicate that the convergence time of the algorithm scales linearly with the agent population size.

1 Introduction

According to the US Department of Energy, the creation of a sustainable and energy-efficient society is one of the greatest challenges of this century, as traditional non-renewable sources of energy are depleting and adverse effects of carbon emissions are being felt (DOE, 2003). Two key issues in creating a sustainable and energy-efficient society are reducing peak energy demands and increasing the penetration of renewable energy sources. In order to achieve a reliable operation of the electricity distribution system, supply and the load have to be balanced within a tight tolerance in real time. One way, which is most commonly used, to achieve the demand supply balance is to supply all the required 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, so that flexible power plants required to operate for meeting peak demands are used as little as possible (Palensky & Dietrich, 2011). Managing the demand side becomes more critical when the uncertainty in the energy supply increases as is the case with increasing penetration of renewable energy in the electricity market (Medina, Muller, & Roytelman, 2010). In order to adjust consumer demands, various demand response programs have been introduced. Demand Response (DR) can be 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)). The first type of programs are incentive based programs (IBP), where customers receive payments for their participation in the programs. A typical example of an IBP are Direct Load Control (DLC) programs, in which utilities have the ability to remotely control the power consumption of consumers' appliances by switching them on/off. In small scale pilot studies, DLC has been successful in reducing peak energy consumption. However, the biggest drawback of DLC is that consumers may not be comfortable with utility companies having direct control over their appliances (Rahimi & Ipakchi, 2010; Medina et al., 2010).

The second type of programs are price based programs (PBP), which are based on variable pricing rates, so that energy rates 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 by shifting energy from peak to off-peak times. The basic example of PBP are TOU programs, in which the rate during peak times is higher than the rate during other off-peak times. Recent technological advances in smart meters and smart appliances have enabled direct and real time participation (RTP) of an individual consumer in the energy market through the use of software agents. This allows price based systems with hourly prices depending on the actual cost of generation. A key feature of RTP programs is that each customer communicates with the utility companies individually. However, there are two key problems in realizing this potential. First, despite the presence of small pilot programs, the end users are usually not of sufficient size for the utilities to be considered for demand response services. Second, if end users participate in the market directly, without control by the utility companies, the stability of the system may be compromised, due to uncontrolled distributed interactions (Ramchurn, Vytelingum, Rogers, & Jennings, 2012).

In (Mohsenian-Rad, Wong, Jatskevich, Schober, & Leon-Garcia, 2010) it is argued that a good demand side management program should focus on controlling the aggregate load (which is also important for economic load dispatching (Wood & Wollenberg, 1996)) of a group of consumers instead of individual consumers. Therefore, in this paper the problem of coordinating a group of consumers called *consumer cooperatives* is introduced and studied. A consumer cooperative 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 consumption, and direct participation in the energy markets. 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 for cost effective electricity allocation. It has to ensure that the demand goals and constraints of the group members (consumers) are fulfilled, while also enabling to flatten out peak demands for the group. The consumers espouse the goals of the group, but they are not willing to completely disclose their demand goals and constraints to either other firms or the coordinator. Moreover, the members autonomously decide how to shift their loads to help the group flatten peak demands. Real world consumer groups coordinated in the above manner can be formed naturally in many application scenarios, especially when they are geographically co-located, e.g., industrial parks/technology parks, commercial estates or large residential complexes.

Consumer cooperatives offer advantages to both energy utility companies and consumers. From the utility'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. The partially centralized model offers several advantages to the stakeholders: 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 is of importance, since the mediated participation of the consumer group in the market 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 (which is not the case in current long term fixed contract structures (Kirschen, 2003)). Second, by buying collectively, the group can benefit from volume discounts. The situation here is analogous to group insurance programs in companies. Third, in negotiated electricity contracts, the price usually consists of two components, one coming from 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 in reducing the premium against volatility, provided the demands of the group members are coordinated, so that their total demand is more stable and lower during demand peaks. Thus, in this paper, we study the problem of coordinating the electricity demand of agents who are purchasing electricity as a consumer cooperative.

The objective of this coordination is to minimize the cost function of the total cost of procuring electricity. The technical challenge in designing a coordinated demand management for a consumer cooperative 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 will depend on the aggregate consumption profile of all agents. However, the agents may not want to share their consumption patterns or consumption constraints with other agents. Therefore, in this paper, an algorithm is designed that enables the central agent to coordinate the consumers to achieve the optimal centralized consumption load, while the individual agents retain their private knowledge about their consumption constraints.

We first consider the simple but practically relevant problem with a planning horizon of only two time slots with different electricity prices, where agents have demand constraints that have to be satisfied, but they do not have any cost for shifting demands between different time slots. This will be called the *basic setting*. We design a simple iterative algorithm, where in each iteration the coordinator computes *virtual price signals* and sends them to the consumers, who then compute their optimal consumption profiles based on this price signal and send it back to the coordinator. This will be called the *basic coordination algorithm*. We show that in the problem setting with only two time slots and no cost for shifting demand, this iterative algorithm converges to the optimal schedule. We then consider settings with a planning horizon of more than two time slots with different electricity prices and with individual costs for the agents to shift demand. This will be called the *general setting*. Since the basic algorithm does not guarantee the optimal solution for general setting, we design an iterative algorithm that includes an additional phase. This general algorithm first runs the basic algorithm until it converges and then the additional phase if necessary. In that additional phase, in each iteration, the agents compute their marginal valuation for their electricity demand in addition to their optimal consumption profiles and send them back to the coordinator. The coordinator uses the additional information to adapt the virtual price signals. We show that this general iterative algorithm converges to the optimal schedule in the general setting. These provably optimal demand scheduling algorithms for consumer cooperatives are the primary contribution of this paper. We also conduct simulation studies to characterize (a) the convergence properties of our algorithm and (b) understand the effect of different parameters that characterize the electricity consumption profile on the potential cost reduction through coordination by the algorithm. The results show that as the participants' flexibility of shifting their demands increases, cost reduction increases. We also observe that the cost reduction is not very sensitive to the variation in consumption patterns of the consumers (e.g., whether the consumers use more electricity during the evening or during the day). Finally, our simulations indicate that the convergence time of the algorithm scales linearly with the agent population size. A preliminary version of this work appeared in (Veit, Xu, Zheng, Chakraborty, & 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 for which we prove in Section 4.2 that it converges to the optimal solution in basic settings with only two time slots and no cost of shifting demand. In Section 5 we introduce the general iterative algorithm for which we prove in Section 5.3 that it converges to the optimal solution in general settings. Subsequently, in Section 6 we evaluate the coordination algorithms using simulations based on ral world consumption data. In particular, in Section 6.1 we describe the data sets used for the simulations and explain the parameterization of the simulation scenarios. In Section 6.2 we discuss the results of the simulations, regarding the potential cost reduction through coordination as well as the convergence properties of the algorithm. Finally in Section 7 we summarize the main contributions of this paper and give an outlook on future work.

2 Related Work

The literature on demand management and demand response is extensive. As mentioned in the introduction, the demand response programs vary from classical direct load control (DLC) to price based programs with real time prices (RTP). 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.

The initial pricing schemes utilizing variable price signals in order to influence consumer demand used prices for electricity that varied according to the time of the day or day of the year. These traditional time of use (TOU) pricing schemes penalize certain periods of time with a higher electricity price, so that customers respond to these signals by adjusting their consumption to reduce their own cost. Thereby, the electricity price is set to be high at times having typically high consumption so that demand during peak load times can be reduced. However, TOU does not necessarily reduce the overall energy demand, but only the consumption patterns are influenced (Rahimi & Ipakchi, 2010; Kirschen, 2003; Medina et al., 2010; Albadi & El-Saadany, 2007; Palensky & Dietrich, 2011). Most of the literature looking at 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 can be applied to all long-term contract markets and day-ahead markets if the planning horizon is sufficiently short (see (Vytelingum, Ramchurn, Rogers, & Jennings, 2010) and (Ramchurn, Vytelingum, Rogers, & Jennings, 2011)). Known electricity prices are used so that a central manager can deliver the information to consumers to incentivize consumers to shift demand to times with low-prices.

Current literature on demand shaping mostly operates in the paradigm that it is desirable to have a system where the consumer can set its preferences and constraints about the timing of the operation of the appliances (or loads), and have an automated (or autonomous) system that ensures the electricity cost is minimized while the user preferences are met. It is (implicitly or explicitly) assumed that the utility companies can send a price signal to software agents at those systems (or smart meters) that respond to this price and schedule the appliances for the future.

The different approaches for demand scheduling proposed in the literature differ in some central characteristics. First they differ in the *objective of the demand scheduling* and secondly they differ in the *level at which the problem is solved*. 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. The different levels at which the problem has been studied include the level of the single user, the market maker, as well as the grid operator.

When the demand scheduling problem is studied at the grid operator level, the grid stability (not the energy cost) is the main objective. In particular, objectives include the minimization of power flow fluctuations (Tanaka, Uchida, Ogimi, Goya, Yona, Senjy, Funabashi, & Kim, 2011), integration of green energy (Wu, Mohsenian-Rad, & Huang, 2012) and the minimization of power losses and voltage deviations (Clement-Nyns, Haesen, & Driesen, 2010).

Most work on demand scheduling focuses at the level of the end users. An important characteristic of the end user is that it cannot influence the electricity prices, i.e., the end user is a price taker. In demand scheduling for end users the usual objective is to minimize the cost of electricity consumption of a residential or commercial end user by optimally scheduling the demand. However, the studies differ in some key characteristics. Differences include the particular problem they solve, the assumption on the knowledge about the prices and demand and whether an optimal solution to the optimization problem is guaranteed. In (Chu & Jong, 2008) the authors assume known electricity prices and demand for their cost optimization. However they restrict the considered loads to air conditioning systems and in contrast to this paper the authors focus on load shedding and not load shifting. The authors in (Pedrasa, Spooner, & MacGill, 2010) also assume known prices and demand, but for their optimization they use Particle Swarm Optimization. In (Philpott & Pettersen, 2006) and (Samadi, Mohsenian-Rad, Wong, & Schober, 2013) the prices for electricity are known , but no central knowledge of the demand is assumed, only an estimate. In (Philpott & Pettersen, 2006) the authors study how to optimally bid in the day ahead electricity market, if the actual demand is uncertain and the difference to the bid has to be bought from the real time market. In (Samadi et al., 2013) the authors formulate an optimization problem for the real time residential load management that only requires some statistical estimates of the future load demand.

The next studies focus on a setting where the demand is assumed to be known but there is price uncertainty. In (Conejo, Morales, & Baringo, 2010) the authors present a procedure to adjust the hourly load level of a household in response to real-time electricity prices, where only the price for the next hour is known. In (Kim & Poor, 2011) the scheduler only has statistical knowledge about the future prices and the scheduling problem is thus modeled as a Markov decision process. In (Mohsenian-Rad & Leon-Garcia, 2010) the energy consumption of a household is also scheduled in response to time-varying prices. In addition the authors claim that the use of an inclining block rate as pricing scheme can prevent load synchronization. However they only consider one consumer that has perfect knowledge about the load of all appliances.

In practice, it may be difficult for utility companies or grid operators to deal with individual end users in such demand response programs. The demand shift of an individual end user 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. Moreover there have been concerns voiced that the stability of the system may be compromised with such uncontrolled distributed interactions (Kirschen, 2003). In our work the demand scheduling problem is solved at the level of a consumer cooperative where the objective is to minimize the electricity procurement cost for the cooperative. The demand constraints and preferences are private knowledge to the consumers and not known to the coordinator. The cooperative is also a price taker, because the cooperative buys the electricity from a utility company. From the utility's perspective, the consumer cooperative is large enough to be useful in demand response programs.

There are also studies in the extant literature that focus on the demand scheduling at the level of the market maker. In contrast to the previous approaches this means that the coordinator can set the prices for electricity. The objective in these studies is mostly to minimize the total cost of power generation. In (Dietrich, Latorre, Olmos, & Ramos, 2012) an electric system with high wind penetration is modeled in order to compare different demand response programs. They compare demand shifting vs. peak shaving as well as centralized vs. decentralized approaches. They conclude that the centralized approach reaches higher overall cost savings, but has the disadvantage that central knowledge of consumers' constraints and preferences is necessary. The authors in (Parvania & Fotuhi-Firuzabad, 2010) present a stochastic model to schedule reserves provided by demand response in the wholesale electricity markets. In order to create consumers that are of sufficient size for demand response programs they introduce demand response providers that aggregate end consumers. Some other studies use game theoretic approaches to model the consumer behavior. In (Atzeni, Ordonez, Scutari, Palomar, & Fonollosa, 2013) the authors formulate the day-ahead grid optimization problem, whereby each user on the demand-side minimizes its individual cost. In (Vytelingum, Voice, Ramchurn, Rogers, & Jennings, 2011) the authors look at the challenges of the adoption of micro-storage devices for the energy system. The authors characterize the competition equilibrium of the amount of storage that will be adopted by the population. In (Wu, Mohsenian-Rad, Huang, & Wang, 2011) the authors propose a demand side management method to tackle the temporal variations in wind power generation. Using game theory, they analyze the interactions among users to efficiently utilize the available renewable and conventional energy. Some studies consider further objectives. In (Nguyen, Song, & Han, 2012) the authors focus on reducing the peak-to-average ratio (PAR) of consumption. In order to reduce the PAR, users request their energy demands to an energy provider, who dynamically updates the energy prices based on the loads of the users. Another objective is used by the authors in (Baharlouei, Hashemi, Narimani, & Mohsenian-Rad, 2013). They develop an electricity billing mechanism that focuses on both reducing cost and ensuring fairness. A key feature of many of those works is that the agents communicate directly with the utilities and hence, the focus is on controlling the overall load by interacting directly with each consumer. There is no interaction among the consumers. However (Ramchurn et al., 2011) mentions that such variable price signals can only work for a small number of houses whereas in a large scale with many customers there is a possibility it may not reduce demand peak and even cause instabilities like *herding phenomenon*. Thereby agents synchronize their load, because they move their consumptions towards the low price times and thus cause a spike in demand, leading to increasing energy cost (Ramchurn et al., 2012). Although some heuristics have been proposed to address this issue, there is no algorithm with provable guarantees to solve this problem. In order to stabilize the system in (Voice, Vytelingum, Ramchurn, Rogers, & Jennings, 2011) agents are charged an additional fee based on how much they change their storage profile from one period to the next. In (Ramchurn et al., 2011) the authors introduce an adaptive mechanism controlling the rate and frequency at which the agents are allowed to adapt their loads and to readjust their heating profile. In (Vytelingum et al., 2010) a compensation signal is sent to the agents providing an estimate of how much they should aim to change their behavior.

Another approach is proposed by (Mohsenian-Rad et al., 2010) focusing on a utility or a generator controlling the load of a group of consumers. The authors accomplish this by allowing the individual consumers to interact with one another. The problem is formulated in a game-theoretic framework and the consumers coordinate in an iterative manner and exchange their demand profiles (but not their consumption constraints with each other). In our paper, consumers interact as a coordinated group with the utility in order to prevent load synchronization, but the consumer architecture is different. The consumers form a cooperative with a central coordinator and the agents do not share their consumption profiles with other agents.

3 Problem Formulation

In this paper we consider a setting in which a central coordinator purchases electricity from a supplier to support the demands of a consumer group. Figure 1a shows an example of such a consumer group. The group consists of several commercial consumers that use one central coordinator to purchase the necessary electricity from a supplier or from the electricity market. The flow of money is illustrated by the green dotted arrows. The yellow dashed arrows show the flow of electricity. For the coordination there is an exchange of information between the coordinator and the individual agents. This flow of information is illustrated by the blue arrows. We assume that the price of electricity is known over the whole planning horizon. This is true when the agent group has a long-term electricity contract (say yearly) and the agents planning horizon is shorter (say 1 day). The contract is not a flat-rate contract, since in this case there would be no economic incentive for the agents to shift their demands.

We consider the consumer group to consist of N members with the planning period divided into M discrete time slots. The amount of discrete time slots depends on the market price structure, which can be different in practice, based on the utility companies. Figure 1b shows the logical schema of such a consumer group. The Figure illustrates how the coordinator aggregates the demands of the N agents of the planning horizon of M discrete time slots. Note that $M = 2$ for time-of-use pricing with different prices during day and night, whereas $M = 24$ in an hourly pricing scheme. 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, $i \in \{1, 2, \ldots, N\}$. 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 consumption levels. For each time slot, every unit of electricity consumed below a specified threshold is charged at a lower price, while any additional unit exceeding that threshold is charged at a higher price. This is an example of a non-flat electricity pricing rate, which has also been called a "two-level inclining block rate" by (Mohsenian-Rad & Leon-Garcia, 2010). 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 consumption level is the payment increment (decrement) for adding (reducing) one unit of electricity. Figure (2a) shows an example of a two-level increasing threshold pricing model adopted from BC Hydro.[1] The marginal price of a two-level threshold structure can formally be written as follows: $\mathrm{p}_j^m(\rho_j) = \begin{cases} p_j^H & \rho_j > h_j \\ p_j^L & \rho_j \le h_j \end{cases}$ with $p_j^H > p_j^L$, where h_j is the threshold for consumption in time slot j. We further assume that the high price in any time slot is greater than the low price in any other time slot, i.e., $p_j^H > p_k^L,\ \forall j, k$. Let for the further analysis $x^+ = \max\{0, x\}$ and $x^- = \min\{0, x\}$. The total energy cost for time slot j is thus the integral of the marginal prices. Figure (2b) shows the total electricity cost for the aggregated demand based on the

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

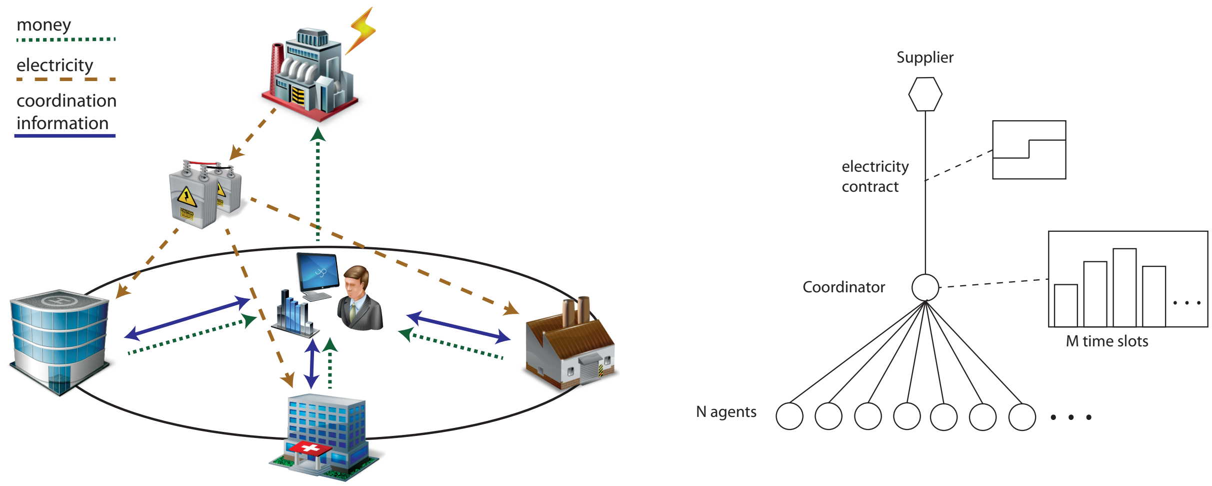

(a) Example setup of real world cooperative.

(b) Logical schema of a cooperative.

Figure 1: System overview of the architecture as considered in this paper including the cooperative and supply.

two-level level threshold pricing model. The total electricity cost can be computed as:

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

The demand profile of each agent $\mathbf{r}_i$ must satisfy its individual constraints. We 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. The overall demand can come from two types of loads, shiftable loads and non-shiftable loads. We will consider loads where the consumption constraints are given by a constraint set $\mathcal{X}_i$ which is private knowledge of the agent i. An agent does not share this constraint set $\mathcal{X}_i$; neither with other firms nor with the coordinator. Unless otherwise specified this constraint set is assumed to be a convex polytope. In some application scenarios, when an agent determines its energy consumption profile, it has to consider additional costs associated with the consumption schedule. For example, in any given factory the energy is most commonly used for production. Changing the energy consumption 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 *central 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. 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.

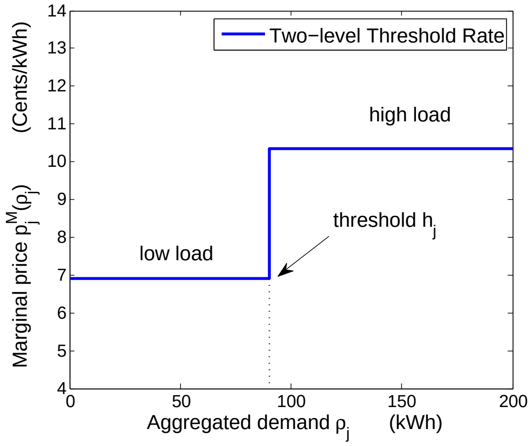

(a) Marginal electricity prices.

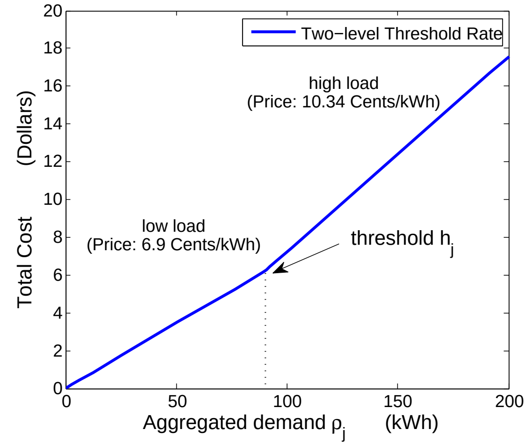

(b) Total electricity cost.

Figure 2: Sample two-level increasing threshold pricing model as used by BC Hydro

4 Solution Approach

Although the problem in (2) is a convex optimization problem, since the constraints and preferences of the agents are private knowledge, the optimal consumption profiles cannot be computed directly by the central coordinator. However since the constraints in Problem (2) are agent-specific, they are naturally separable. 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 consumption of all agents ρ_j. 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 consumption 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 initial demand profile. Let Figure 3a depict such an aggregated initial 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 are below, $\rho_1 < h_1, \rho_3 < h_3$. It follows that 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 coordinating the demand would be for the coordinator to inform the 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 constraints and consumption preferences of the other agents in the group, while their costs strongly depend on the consumption 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 peak of the demand in that time slot and thus

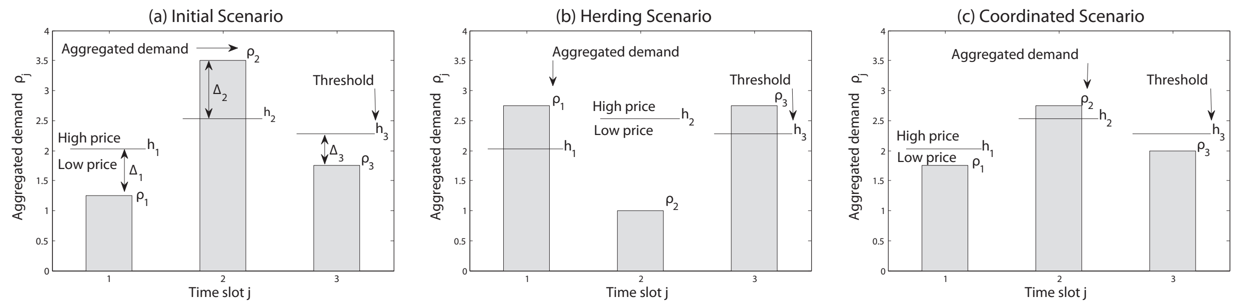

Figure 3: A comparison of demand profiles in Initial, Herding and Coordinated scenarios.

increasing the total cost. The effect of this herding phenomenon is shown in Figure 3b, 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 price signal which the coordinator sends to the agents.

In this section we propose a novel virtual price signal that the coordinator uses to guide the agents' demand profiles. *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 goal of designing the virtual price signal is to enable 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 demand profile $\mathbf{R}$, which is therefore added into the price function. For ease of readability the aspect of 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 (3a), this amount, denoted by Δ_j, $j = 1, 2, 3$, is the difference between the total aggregated demand and the threshold in each time slot. To avoid herding the amount Δ_j 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 how much demand they can 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

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

Note that the above problem, like the central problem, is a convex optimization problem and is 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 3c shows this case, where the total demand in the second time slot remains above the threshold, because the whole Δ_2 could not 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.

The rest of the section is structured as follows: In the first subsection, we introduce an iterative algorithm for coordinating the agents for the simplified setting of two time slots and no shifting cost. In the subsequent discussion, this will be called the ”basic setting” and the algorithm will be called the ”basic coordination algorithm”. Then we prove that the basic coordination algorithm converges for any setting. However, it converges to the *optimal solution* for the basic setting. Therefore we introduce an iterative algorithm for general settings. In particular, we state the conditions under which it is necessary to extend the basic algorithm with an additional phase. We also present a detailed description of the additional phase and finally prove that the general algorithm converges to the optimal solution.

4.1 Basic Coordination Algorithm

We will now present the basic coordination algorithm and the details of 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. The 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 so that the central optimal solution can be reached. 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. Second, the individual agents use the virtual price signal to solve their individual cost optimization problem and report their new demand profile to the coordinator.

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: All agents compute an initial energy consumption profile $\mathbf{r}_i$ by solving Problem (3) based on the market prices and send it to the coordinator.

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

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

3. After receiving the virtual price signal, all agents individually calculate their new demand profiles $\mathbf{r}_i'$ according to their 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. Otherwise, it sets $\mathbf{R} = \mathbf{R}'$ and goes to step (1).

4.1.2 Coordination with virtual price signal

In this section we explain how the central coordinator coordinates the firms in demand shifting via virtual price signals. The key idea is to use the marginal electricity cost in each time slot to create the virtual price signals for the agents. Shifting demand from time slots with high marginal cost to those with low

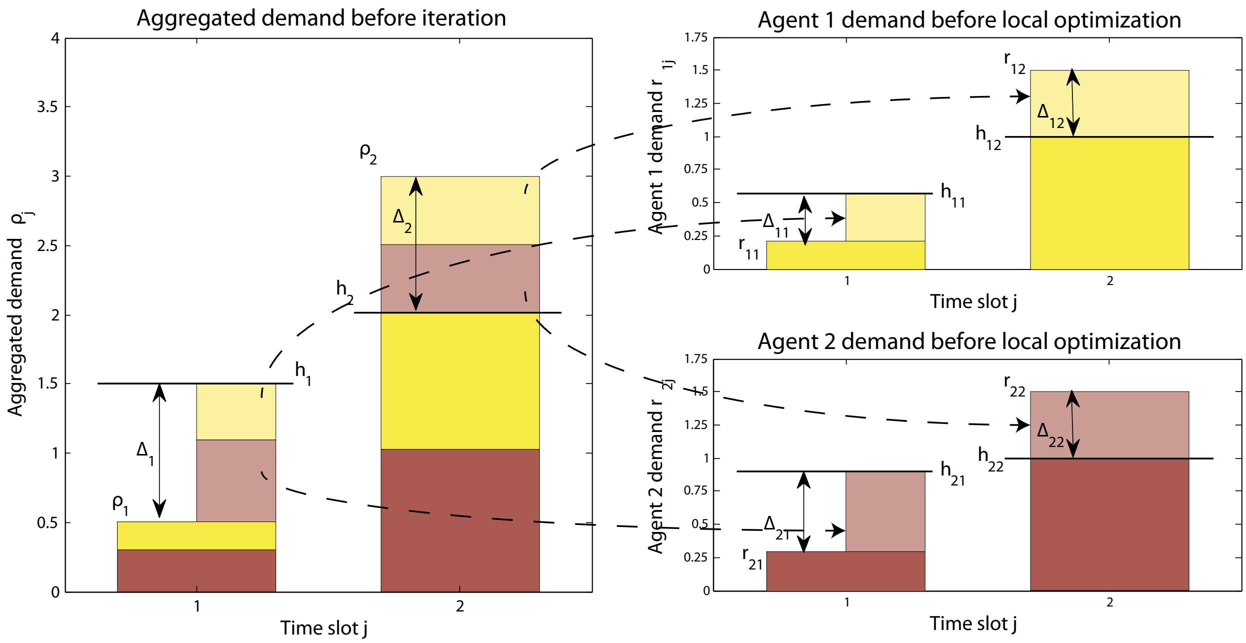

Figure 4: Division of Δ among agents for the price signal. Here Δ_1 is the amount of demand that can be shifted into time slot 1 while Δ_2 indicates the amount of demand that needs to be shifted out of time slot 2.

marginal cost can be beneficial provided an appropriate amount is shifted. An appropriate shift implies that the resulting aggregated demand profile does not lead to higher marginal price in the former cheap time slots. As an example, in Figure (3a) a demand shift from time slot 2 to time slot 1 could be beneficial, as long as the shifted amount does not exceed Δ_1. In order to limit the agents demand shifts they have to be able to foresee the price changes caused by their demand shifting.

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 exceeding demand is charged at a higher price. The virtual price signal is therefore parameterized by the low marginal price, p_j^L, the high marginal price, p_j^H, and the consumption threshold, $\mathrm{h}_{ij}(\mathbf{R})$, that specifies the demand levels at which the prices apply. The virtual price, s_{ij}^v, for agent i in time slot j is computed based on these parameters as follows

$$\mathrm{s}_{ij}^v(r_{ij}|\mathbf{R}) = \begin{cases} \frac{p_j^L\,\mathrm{h}_{ij}(\mathbf{R}) + p_j^H(r_{ij} - \mathrm{h}_{ij}(\mathbf{R}))}{r_{ij}} & r_{ij} > \mathrm{h}_{ij}(\mathbf{R}) \\ p_j^L & r_{ij} \leq \mathrm{h}_{ij}(\mathbf{R}) \end{cases} \quad (4)$$

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:

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

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.

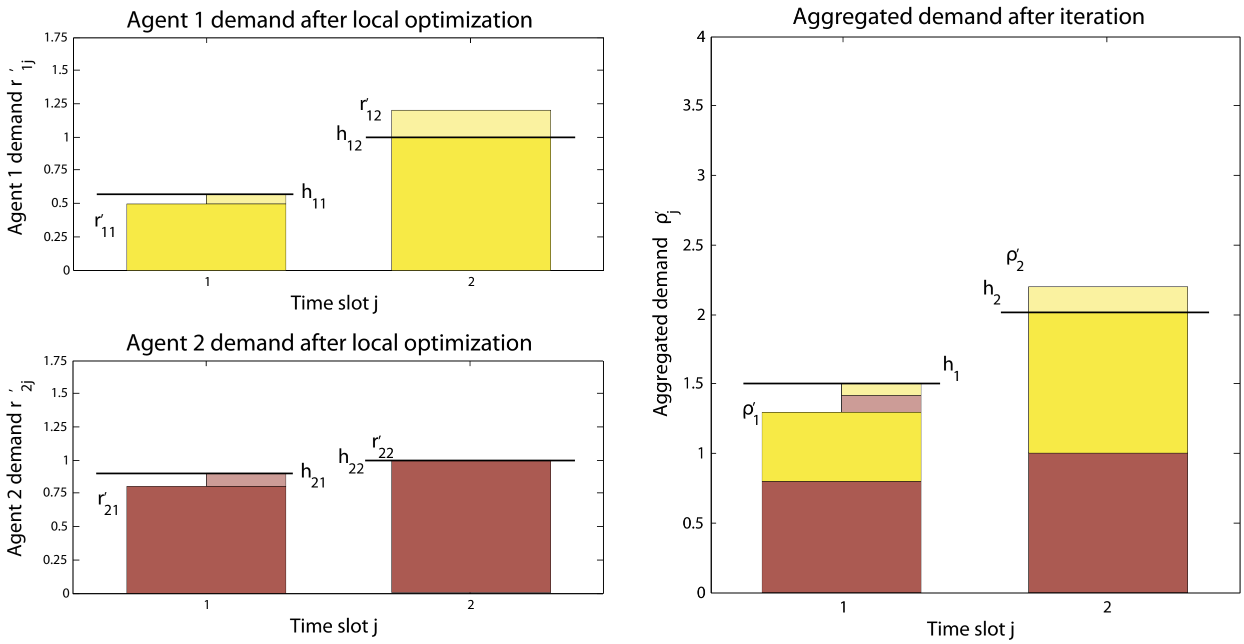

Figure 5: Demand after agent's local optimization. Agent 1 cannot reduce its entire demand from time slot 2 (i.e., $r_{12} > h_{12}$ because of its own consumption constraints (left top plot). Thus the aggregated load in time slot 1 is still below the threshold (i.e., $\rho_1' < h_1$). Thus the central coordinator again sends a modified price signal to the agents and the process continues until agents stop shifting their demands.

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 load change of the agents to be less than Δ_j the coordinator 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 total demand in that time slot, i.e., $\Delta_{ij} = \frac{\Delta_j r_{ij}}{\sum_i r_{ij}}$. Figure 4 shows how the total load change is divided among the agents in order to create the individual virtual price signals. The left side shows the initial aggregated loads of two agents. The yellow load belongs to agent 1 and the red load 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 needs 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 load 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 firms will independently optimize their demand profiles in order to minimize their cost according to Problem (3). Together with the virtual price signal 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 virtual price signal divides the amounts to be shifted among the agents so that 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 5 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 according to their virtual price signal. In comparison to Figure 4 on the right side it can be seen that agent 2 shifted the whole amount as was allocated, but agent 1 only shifted some part of it, due to its constraints. The right side 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.

4.2 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 $M = 2$ and $\mathrm{g}_i(\cdot) = 0$. In Lemma 1, it is shown 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 it is shown that, when $M = 2$ and $\mathrm{g}_i(\cdot) = 0$, the converged solution is an optimal solution. Subsequently it is shown that the algorithm can get stuck in a suboptimal solution in general settings if $M > 2$ (Lemma 2) or if $\mathrm{g}_i(\cdot) \neq 0$ (Lemma 3).

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

Proof. Let's first introduce some notation that will be used throughout this proof. Let $\mathbf{R}$ be the demand profile at the end of iteration round t and $\mathbf{R}'$ be the demand profile at the end of round $t+1$. Similarly, let $\mathrm{C}(\mathbf{R})$ be the electricity cost at the end of iteration round t and $\mathrm{C}(\mathbf{R}')$ be the electricity cost at the end of round $t+1$. The virtual prices for round $t+1$ are computed by the central manager using $\mathbf{R}$. Let $\mathrm{C}_i^v(\mathbf{r}_i|\mathbf{R})$ be the cost of electricity 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 (given by the objective function in Problem (2)) for the consumer group based on market prices is equal to the sum of the individual cost of the agents (give by the objective function in Problem (3)) based on the virtual price signals $\sum_{i=1}^{N} \mathrm{C}_i^v(\mathbf{r}_i|\mathbf{R}) = \mathrm{C}(\mathbf{R})$:

$$
\begin{aligned}
&\sum_{i=1}^{N}\sum_{j=1}^{M} \mathrm{s}_{ij}^v(r_{ij}|\mathbf{R})\, r_{ij} + \sum_{i=1}^{N} \mathrm{g}_i(\mathbf{r}_i) \\
&= \sum_{i=1}^{N}\sum_{j=1}^{M} \left[p_j^H \left(r_{ij} - \left(r_{ij} + \frac{(h_j - \rho_j)\, r_{ij}}{\sum_i r_{ij}} \right) \right)^+ \right. \\
&+ p_j^L \left(r_{ij} - \left(r_{ij} + \frac{(h_j - \rho_j)\, r_{ij}}{\sum_i r_{ij}} \right) \right)^- \\
&\left. + p_j^L \left(r_{ij} + \frac{(h_j - \rho_j)\, r_{ij}}{\sum_i r_{ij}} \right) \right] + \sum_{i=1}^{N} \mathrm{g}_i(\mathbf{r}_i) \\
&= \sum_{j=1}^{M} \left[p_j^H (\rho_j - h_j)^+ + p_j^L (\rho_j - h_j)^- + p_j^L h_j \right] + \sum_{i=1}^{N} \mathrm{g}_i(\mathbf{r}_i) \\
&= \sum_{i=1}^{N}\sum_{j=1}^{M} \mathrm{p}_j(\rho_j)\, r_{ij} + \sum_{i=1}^{N} \mathrm{g}_i(\mathbf{r}_i)
\end{aligned}
\tag{8}
$$

If the algorithm has not stopped, at least one agent has changed its demand profile, i.e., $\exists i$ with $\mathbf{r_i'} \neq \mathbf{r_i}$. According to Problem (3) agents only change their demand profile, if that reduces their cost. 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 $\mathbf{r_i}$:

$$
\mathrm{C}_i^v\left(\mathbf{r}_i'|\mathbf{R}\right) < \mathrm{C}_i^v(\mathbf{r}_i|\mathbf{R}) \tag{9}
$$

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 given the new aggregated demand is also lower or equal to the sum of the agents' individual cost given 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 total 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 the ease of readability. Please consult Appendix A for the complete proof. For any time slot, j, we have

$$
\begin{aligned}
&\sum_{i=1}^{N} \mathrm{p}_j\left(\rho_j'\right) r_{ij}' - \sum_{i=1}^{N} \mathrm{s}_{ij}^v\left(r_{ij}'|\mathbf{R}\right) r_{ij}' \\
&= \begin{cases} \sum_{i:\, r_{ij}' \le h_{ij}}^{N} \left(p_j^H - p_j^L\right)\left(r_{ij}' - \mathrm{h}_{ij}(\mathbf{R})\right) & \rho_j' > h_j \\ \sum_{i:\, r_{ij}' > h_{ij}}^{N} \left(p_j^L - p_j^H\right)\left(r_{ij}' - \mathrm{h}_{ij}(\mathbf{R})\right) & \rho_j' \le h_j \end{cases} \\
&= \begin{cases} \left[\sum_{i=1}^{N} \left(p_j^H - p_j^L\right)\left(r_{ij}' - \mathrm{h}_{ij}(\mathbf{R})\right)^- \right] \le 0 & \rho_j' > h_j \\ \left[\sum_{i=1}^{N} \left(p_j^L - p_j^H\right)\left(r_{ij}' - \mathrm{h}_{ij}(\mathbf{R})\right)^+ \right] \le 0 & \rho_j' \le h_j \end{cases}
\end{aligned}
\tag{10}
$$

Since $\sum_{i=1}^{N} \mathrm{p}_j\left(\rho'_j\right) r'_{ij} \leq \sum_{i=1}^{N} \mathrm{s}^v_{ij}\left(r'_{ij}|\mathbf{R}\right) r'_{ij}\ \forall j$, it also holds for the sum of all time slots: $\mathrm{C}\left(\mathbf{R}'\right) \leq \sum_{i=1}^{N} \mathrm{C}^v_i\left(\mathbf{r}'_i|\mathbf{R}\right)$. From Equations 8, 9 and 10 we can conclude that:

$$\mathrm{C}\left(\mathbf{R}'\right) \leq \sum_{i=1}^{N} \mathrm{C}^v_i\left(\mathbf{r}'_i|\mathbf{R}\right) < \sum_{i=1}^{N} \mathrm{C}^v_i\left(\mathbf{r}_i|\mathbf{R}\right) = \mathrm{C}\left(\mathbf{R}\right) \tag{11}$$

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

Theorem 1. *The basic iterative algorithm for solving problem* (2) *always converges for any value of M.*

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. □

Theorem 2. *Assuming the basic setting with $M = 2$ and $\mathrm{g}_i(\cdot) = 0$, the converged solution of the basic algorithm $\mathbf{R}$ is optimal.*

Proof. We now prove by contradiction that when the algorithm has converged to the solution $\mathbf{R}$, then 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 $\mathrm{p}^m_j(\rho_j) > \mathrm{p}^m_k(\rho_k)$, because otherwise the new cost would not be lower. In the proof below, three cases are considered.

First, if $\rho_j \neq h_j$, $\rho_k \neq h_k$ we get that all agents have active constraints such that a shift from j to k would be infeasible for the agents, because otherwise a shift from j to k would be beneficial for at least one agent and the algorithm would not have stopped. Since there is no feasible shift from j to k it follows $\rho'_j \geq \rho_j$ and $\rho'_k \leq \rho_k$. With $\mathrm{C}(\mathbf{R}') < \mathrm{C}(\mathbf{R})$ it follows that $\mathrm{p}^m_j(\rho_j) < \mathrm{p}^m_k(\rho_k)$, which leads to a contradiction since our assumption implies that $p^m_j(\rho_j) > p^m_k(\rho_k)$.

If $\rho_j = h_j$, $\rho_k = h_k$ then $\mathrm{p}^m_j(\rho_j) = p^L_j$ (as ρ_j decreases) and $\mathrm{p}^m_k(\rho_k) = p^H_k$ (as ρ_k increases). From the market price structure we have $p^H_k > p^L_j$. It follows $\mathrm{p}^m_j(\rho_j) < \mathrm{p}^m_k(\rho_k)$, which leads to a contradiction.

If $\rho_j \neq h_j$, $\rho_k = h_k$ and if $\rho_j < h_j$ then $\mathrm{p}^m_j(\rho_j) = p^L_j$ and $\mathrm{p}^m_k(\rho_k) = p^H_k$ thus $\mathrm{p}^m_j(\rho_j) < \mathrm{p}^m_k(\rho_k)$. If $\rho_j > h_j$ then $\mathrm{p}^m_j(\rho_j) = p^H_j$ and $\mathrm{p}^m_k(\rho_k) = p^H_k$. We have $p^H_k > p^H_j$, 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 $\mathrm{p}^m_j(\rho_j) < \mathrm{p}^m_k(\rho_k)$, 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 2. *The basic algorithm can converge to a suboptimal solution in settings with $M > 2$.*

Proof. We now prove by presenting a counterexample that the basic algorithm can converge to a suboptimal solution in general settings with $M > 2$. 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:

$$(p^L_1, p^H_1) = (3, 6),\ h_1 = 10$$
$$(p^L_2, p^H_2) = (2, 5),\ h_2 = 10$$
$$(p^L_3, p^H_3) = (1, 4),\ h_3 = 10$$

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

$$r_{11} \in [0,3],\ r_{12} \in [0,10],\ r_{13} \in [0,10],\ r_{11} + r_{12} + r_{13} = \tau_1 = 17$$
$$r_{21} \in [0,10],\ r_{22} \in [0,10],\ r_{23} \in [9,15],\ r_{21} + r_{22} + r_{23} = \tau_2 = 17$$

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 $\mathbf{r}_1^{(1)} = (0, 7, 10)$, $\mathbf{r}_2^{(1)} = (0, 2, 15)$. The cost based on the initial demand profiles is $\mathrm{C}\left(\mathbf{R}^{(1)}\right) = 88$.

At convergence the final profiles of the two agents are $\mathbf{r}_1^{(T)} = (3, 7+\frac{7}{9}, 6+\frac{2}{9})$, $\mathbf{r}_2^{(T)} = (5+\frac{7}{9}, 2+\frac{2}{9}, 9)$ (see Figure 6a for a graphical representation). The cost based on these demand profiles is $\mathrm{C}\left(\mathbf{R}^{(T)}\right) = 77+\frac{2}{9}$.

However a different demand profile of the agents exists with $\mathbf{r}'_1 = (3, 10, 4)$, $\mathbf{r}'_2 = (7, 0, 10)$ (see Figure 6b for a graphical representation). This profile is feasible and leads to lower total cost $\mathrm{C}\left(\mathbf{R}'\right) = 76$. It follows that the algorithm has stopped in a suboptimal solution. □

Lemma 3. *The basic algorithm can converge to a suboptimal solution in settings with* $\mathrm{g}_i(\cdot) \neq 0$.

Proof. In Appendix B a counterexample is presented, to prove that the basic algorithm can converge to a suboptimal solution in general settings with $g_i(\cdot) \neq 0$. □

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 which show 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 will present a general coordination algorithm, for the energy allocation in general settings with more than two time slots and non-zero individual costs, $\mathrm{g}(\cdot) \neq 0$, for shifting demand. First, using an example, we explain the reason for the basic algorithm to get stuck at a sub-optimal solution in the general settings. Then we introduce an additional phase to the basic algorithm and state the general algorithm. Subsequently in Theorem 3 we prove that the converged solution of the general algorithm is optimal in general settings.

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 is at a threshold. The coordinator only knows that their marginal valuations are somewhere in between the low price below the threshold p_j^L and the high price above the threshold p_j^H. However the agents might have different valuations so that a shift from an agent with a lower valuation in that time slot to an agent with a higher valuation might be beneficial. The following examples illustrate the statements made above.

In the counterexample for proving Lemma 2 the aggregated demand in the final demand profile $\mathbf{R}^{(T)}$ is at the threshold in time slot 2. The demand profile in the converged profile is shown in Figure 6a. The valuation for one additional unit of energy in time slot 2 for the agents is given by the cost savings generated by the possible reduction in the other time slots. The demand constraints of the agents for each time slot are indicated in Figure 6 by the dashed lines. In any time slot, the lower horizontal dashed line indicates

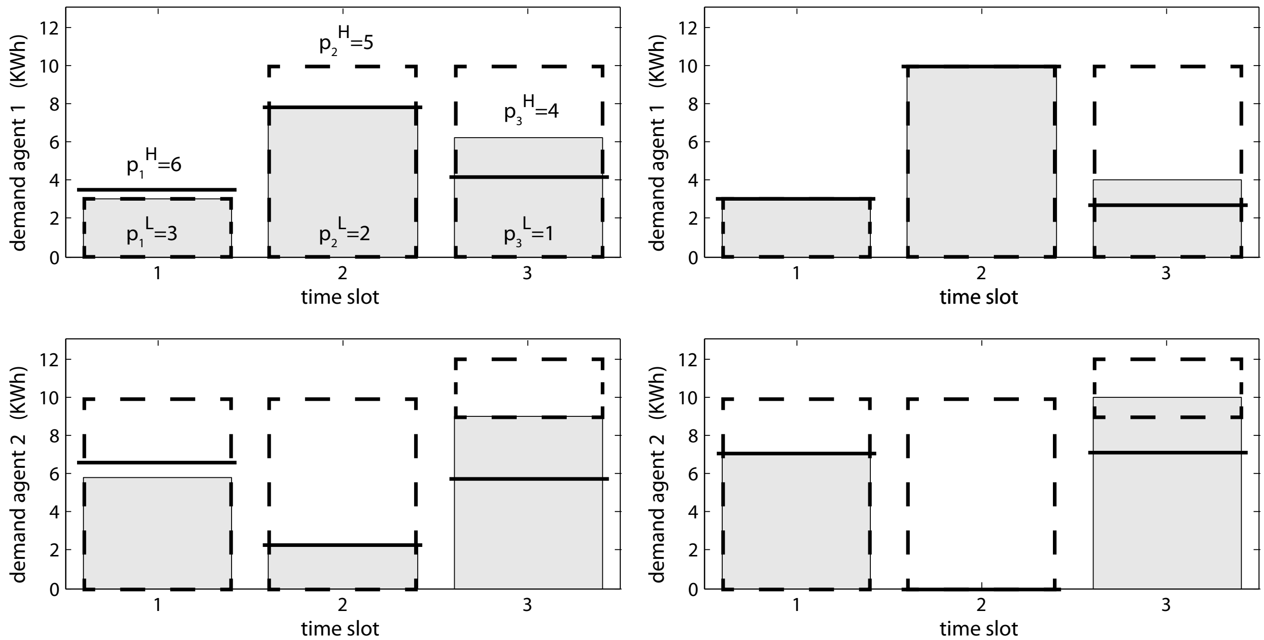

(a) Converged Solution $\mathbf{R}^{(T)}$

(b) Optimal Solution $\mathbf{R}'$

Figure 6: Converged and Optimal solution for the example in Lemma 2. 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 lower bound and the upper horizontal dashed line indicates the upper bound of the electricity demand of the agent. In order to increase demand in time slot 2 agent 1 could decrease demand in either time slot 1 or 3. Since the marginal price for electricity is higher in time slot 3, agent 1's valuation is 4. Because of its constraints agent 2 can only decrease demand in time slot 1 so that its valuation is 3. Both valuations are between the low price $p_2^L = 2$ and the high price $p_2^H = 5$, but agent 1's valuation is higher than agent 2's. A shift from agent 2 to agent 1 would therefore reduce the total cost. However the coordinator does not know which agent has a higher valuation, because it does not know the agents' constraints and thus cannot adjust the thresholds to induce the shift that would decrease the cost. The algorithm thus stops in a suboptimal solution. Similar reasoning can be used with the counterexample given for proving Lemma 3 to understand why the basic algorithm may converge to a suboptimal solution when the individual costs for shifting demand is non-zero.

In order to resolve the problem arising when in the converged solution the aggregated demand is at a threshold in some time slots, let's now introduce an additional phase to the algorithm. The basic idea of this additional phase is that the coordinator queries the agents for their individual valuation for an additional unit of energy in those time slots. The coordinator then uses this information to adjust the virtual price signals.

When the original 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 iteration. In this iteration the coordinator sends an additional query to the agents besides the virtual price signal. The coordinator requests the agents' marginal valuation for an increase and for a decrease of energy in all time slots that are at the threshold.

Having received the virtual price signal and the additional request from the coordinator the agents compute their optimal demand profile $\mathbf{r}_i'$, based on the virtual price signal. In addition the agents compute the vectors $\mathbf{v}_i^+$ and $\mathbf{v}_i^-$. The vector $\mathbf{v}_i^+$ consists of a $v_{ij}^+(\mathbf{R})$ for every time slot at the threshold, $\rho_j = h_j$, indicating the change in cost that would occur to agent i if the coordinator would increase the threshold in that time slot by ϵ. The vector $\mathbf{v}_i^-$ consists of a $v_{ij}^-(\mathbf{R})$ for every time slot at the threshold indicating the change in cost that would occur to agent i if the coordinator would decrease the threshold in that time slot by ϵ. After this computation the agents send $\mathbf{r}_i$, $\mathbf{v}_i^+$ and $\mathbf{v}_i^-$ back to the coordinator.

For the following analysis let $\mathbf{R}^{ij+}$ be the demand profile from which in the resulting individual cost minimization problem of agent i the threshold in time slot j is increased by ϵ, $h_{ij}(\mathbf{R}^{ij+}) = h_{ij}(\mathbf{R}) + \epsilon$. Similarly let $\mathbf{R}^{ij-}$ be the demand profile from which in the resulting individual cost minimization problem of agent i the threshold in j is decreased by ϵ, $h_{ij}(\mathbf{R}^{ij-}) = h_{ij}(\mathbf{R}) - \epsilon$. With that $v_{ij}^+(\mathbf{R})$ and $v_{ij}^-(\mathbf{R})$ are computed as follows:

$$\begin{aligned} v_{ij}^+ &= \mathrm{C}_i^v\left(\mathbf{r}_i^{j+}|\mathbf{R}^{ij+}\right) - \mathrm{C}_i(\mathbf{r}_i') \\ v_{ij}^- &= \mathrm{C}_i^v\left(\mathbf{r}_i^{j-}|\mathbf{R}^{ij-}\right) - \mathrm{C}_i(\mathbf{r}_i') \end{aligned} \tag{12}$$

Having received the demand profiles and marginal valuations of the agents the coordinator computes the virtual price signals for the next iteration. If the aggregated demand is at the threshold, $\rho_j = h_j$, the coordinator finds the agent with the lowest cost for an ϵ raise of the threshold: $l = \arg\min_i \left\{v_{ij}^+\right\}$ and the agent with the lowest cost for an ϵ reduction of the threshold: $k = \arg\min_i \left\{v_{ij}^-\right\}$. If the combined change in cost is negative $v_{kj}^- + v_{lj}^+ < 0$, meaning a beneficial shift exists from agent l to agent k, the coordinator updates the thresholds of agents l, k based on

$$\begin{aligned} \Delta_{lj} &= \epsilon \\ \Delta_{kj} &= -\epsilon \end{aligned}$$

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

$$v_{lj}^+ \geq -v_{kj}^-,\ \forall l, k \in \{1, ..., N\} \text{ and } \forall j \text{ s.t. } \rho_j = h_j \tag{13}$$

because 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 $\mathrm{CalculateVirtualPriceSignals}$ (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).

This general algorithm first runs the basic algorithm as long as the solution does not get stuck and then changes to the more complicated algorithm with the additional phase. In particular, the check for the additional phase is the following: 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 overall algorithm reaches the optimal solution when $\mathbf{R} = \mathbf{R}'$ and either $\rho_j = h_j, \forall j$ or in the additional phase $v_{lj}^{+} \geq -v_{kj}^{-},\ \forall l, k \in \{1, ..., N\}, \forall j$ s.t. $\rho_j = h_j$.

Algorithm 1: Overall algorithm.

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

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

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 Lemma 4 we show the cases for which the central problem is equal in cost to the sum of the individual virtual costs. Then in Theorem 3 we prove that the solution of the general algorithm lies within an ε-neighborhood of the optimal solution.

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. $\min\left\{v_{lj}^+ - v_{kj}^-\right\}$;
  Adapt demand to be shifted for agent l: $\Delta_{lj} \leftarrow \varepsilon$;
  Adapt demand to be shifted for agent k: $\Delta_{kj} \leftarrow -\varepsilon$;
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

Lemma 4. *Given a virtual price signal $\mathrm{s}_{ij}^v\left(\cdot|\mathbf{R}\right)$ and a consumption profile $\{x_{1j}, x_{2j}, ...x_{Nj}\}$ in a time slot j, the real market cost, $\sum_{i=1}^N \mathrm{p}_j\left(\Sigma_{i=1}^N x_{ij}\right) x_{ij}$, equals the virtual cost $\sum_{i=1}^N \mathrm{s}_{ij}^v\left(x_{ij}|\mathbf{R}\right) x_{ij}$, if either $\forall i$, $x_{ij} \geq h_{ij}\left(\mathbf{R}\right)$ or $\forall i$, $x_{ij} \leq h_{ij}\left(\mathbf{R}\right)$.*

Proof. We will prove this by simple algebraic calculations. Assume that either $\forall i$, $x_{ij} \geq h_{ij}\left(\mathbf{R}\right)$ or $\forall i$, $x_{ij} \leq h_{ij}\left(\mathbf{R}\right)$. From Equation (10) we get that the difference between the market cost and the virtual cost, $\sum_{i=1}^N \mathrm{p}_j\left(\Sigma_{i=1}^N x_{ij}\right) x_{ij} - \sum_{i=1}^N \mathrm{s}_{ij}^v\left(x_{ij}|\mathbf{R}\right) x_{ij}$, is given by either

$$\sum_{i:\ x_{ij} \leq h_{ij}}^N \left(p_j^H - p_j^L\right)\left(x_{ij} - \mathrm{h}_{ij}\left(\mathbf{R}\right)\right), \text{ if } \sum_{i=1}^N x_{ij} > h_j$$
$$\sum_{i:\ x_{ij} > h_{ij}}^N \left(p_j^L - p_j^H\right)\left(x_{ij} - \mathrm{h}_{ij}\left(\mathbf{R}\right)\right), \text{ if } \sum_{i=1}^N x_{ij} \leq h_j$$

If $\sum_{i=1}^{N} x_{ij} > h_j$ then by our assumption $\forall i$, $x_{ij} \geq h_{ij}(\mathbf{R})$ so that the two cost are equal. If $\sum_{i=1}^{N} x_{ij} \leq h_j$ then by our assumption $\forall i$, $x_{ij} \leq h_{ij}(\mathbf{R})$ so that the two cost are equal. It follows that the market cost equals the virtual cost. □

Lemma 5. *The converged solution* $\mathbf{R}$ *of the basic and general algorithm is the optimal solution* $\mathbf{R}^*$*, if* $\forall j$, $\mathbf{R}$ *satisfies* $\rho_j \neq h_j$.

Proof. We now prove by contradiction that when the algorithm has converged to the solution $\mathbf{R}$ and $\forall j$, $\mathbf{R}$ satisfies $\rho_j \neq h_j$, then there is no other solution $\mathbf{R}'$ with lower cost. Suppose there exists a solution $\mathbf{R}'$ having a lower 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 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 loads 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 4, $\sum_{i=1}^{N} \mathrm{C}_i^v\left(\mathbf{r}_i^{\#}(\alpha)|\mathbf{R}\right) = \mathrm{C}(\mathbf{R}^{\#}(\alpha))$. Moreover we have $\mathrm{C}(\mathbf{R}^{\#}(\alpha)) < \mathrm{C}(\mathbf{R})$ and from Equation (8) we also have $\sum_{i=1}^{N} \mathrm{C}_i^v(\mathbf{r}_i|\mathbf{R}) = \mathrm{C}(\mathbf{R})$. It follows that $\exists i$, 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 Problem 3. □

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, $\mathbf{r}_i = \arg\min_{\mathbf{x}_i \in \mathcal{X}_i} \mathrm{C}_i^v(\mathbf{x}_i|\mathbf{R})$, 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 Lemma 5 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.

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 consumption on the threshold, $r_{lj}^{\#} = \mathrm{h}_{lj}(\mathbf{R}^{\#\#})$ and $r_{kj}^{\#} = \mathrm{h}_{kj}(\mathbf{R}^{\#\#})$. Thus, with Lemma 4 we get that the sum of the virtual cost is equal to the central cost.

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

which implies either

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

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

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

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 (12) 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

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

which implies either

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

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

$$\sum_{l\in L} \mathrm{C}_l^v \left(\mathbf{r}_l^{\#}(\alpha)|\mathbf{R}^{\#\#}(\alpha)\right) + \sum_{k\in K} \mathrm{C}_k^v \left(\mathbf{r}_k^{\#}(\alpha)|\mathbf{R}^{\#\#}(\alpha)\right) < \sum_{l\in L} \mathrm{C}_l^v (\mathbf{r}_l|\mathbf{R}) + \sum_{k\in K} \mathrm{C}_k^v (\mathbf{r}_k|\mathbf{R})$$

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

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.

6 Simulation

In the previous section, we presented provably-good coordination algorithms that minimize the electricity procurement cost for energy demand scheduling in a consumer cooperative. For practical applications we need to have some insight on the effect of the different variables (e.g., consumption data, electricity prices) on the potential gains that would be achieved through coordination. Furthermore, we do not have any bounds on the number of iterations (or time) it takes for the algorithm to converge. Thus, for practicability, another key aspect is to understand the time it takes for the algorithm to converge in realistic settings. Since the potential gain of coordination and the convergence time are dependent on the actual values of the electricity consumption of agents as well as the electricity prices, it is difficult to have an analytic characterization of these two aspects. Therefore, a multiagent simulation is used for this evaluation.

Each agent in a consumer cooperative is characterized by its total electricity demand and demand constraints for each time slot. The electricity contract is defined by the electricity prices (which is a two level inclining block rate) at different time slots. For the purpose of these simulations we will assume that the constraints of each agent are given by upper and lower bounds of consumption during each time slot. Thus, the total number of variables required for each simulation setting is of the order of the product of the number of agents and the number of timeslots. Furthermore, the typical consumption profile of agents in the cooperative may also vary (e.g., some agents may consume more during the day whereas other agents may consume more during the evening). Since it is difficult to control all the variables individually in a simulation setting, we use four parameters to generate all the simulation data. We use real consumption data from consumers and these four parameters to generate data for the multiagent simulation. Two of the parameters along with real consumption data from consumers are used to generate electricity demand data of the agents. We use a parameter, $p\_fracAgent$ that controls the distribution of the typical consumption profile. The parameter $p\_flexShift$ is used to model the flexibility of the consumers to change their electricity requirement in a time slot. The other two parameters are used for generating the thresholds for each time slot for the electricity prices. The parameter $p\_distThresh$ models the difference between the typical demand in a time slot and the threshold (h_j) whereas $p\_flatThresh$ models the fluctuation of the threshold across the time slots. The parameters will be discussed in more detail in the Subsection 6.1.2.

6.1 Simulation data generation

As stated earlier, we use real consumption data to generate the various scenarios in our simulations. We will now discuss the different data sets used here and then describe the procedure of generating the simulation data.

6.1.1 Datasets

CER electricity consumption data: The agents' demand profiles are generated using real electricity consumption data of 46 small and medium enterprises. This information 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 was divided into time intervals of 30 minutes. Thus, one data point represents the average electricity consumption in kilowatts of one participant during a 30 minute time interval. The participants of the study also answered detailed questionnaires. The questionnaires provide information about the properties of the enterprises and private households. For example the small and medium enterprises answered questions about their number of employees and typical hours of operation. In the context of this analysis we use the data from small and medium enterprises, to simulate an exemplar consumer cooperative like industrial park or commercial area. We will refer to this data set as the `CER_data_set`[2].

EEX electricity prices market data: The cooperative's electricity prices for the simulation scenarios are generated using real day-ahead market electricity prices gathered from the European Energy Exchange (EEX), which is the leading energy exchange in Europe. In particular the hourly day-ahead prices have been collected from the EPEX Spot market, that operates the day-ahead auctions as well as an intra-day market with a yearly volume in 2012 of 339 TWh. [3] As input, the average hourly day-ahead spot market prices from all 20 Tuesdays from January 1st 2013 to May 14th 2013 have been used. Figure 7 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`[4].

6.1.2 Simulation parameters

Each agent is defined by its total electricity requirement over the whole planning horizon, $\tau_i = \sum_{j=1}^{M} r_{ij}$, and the agent's constraints that determine the distribution of electricity consumption over the planning horizon. In this simulation we assume that the constraints are given by upper and lower bounds on the electricity consumption in each time slot, i.e., $r_{ij} \in \left[\underline{r_{ij}}, \overline{r_{ij}}\right]$. Note that the total consumption has to satisfy the agent's individual consumption constraints, $\tau_i \in \left[\sum_{j=1}^{M} \underline{r_{ij}}, \sum_{j=1}^{M} \overline{r_{ij}}\right]$. Since we do not have data on industry specific individual cost functions we assume $\mathrm{g}_i = 0$ for our simulations.

[2]The data set is available at http://www.ucd.ie/issda/data/commissionforenergyregulation/.
[3]http://cdn.eex.com/document/123681/20130114_EEX_Jahresrueckblick.pdf
[4]The data is available at http://www.eex.com/en/Market%20Data.

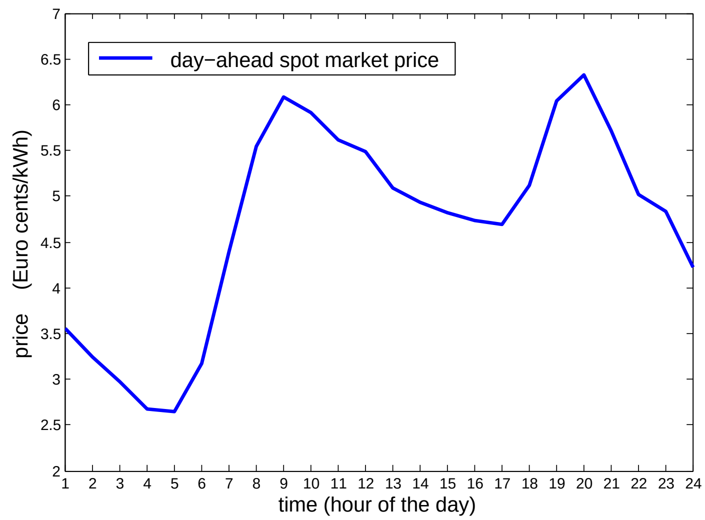

Figure 7: Day-ahead spot market prices from EPEX Spot for an average Tuesday in 2013

Modeling the diversity of cooperatives: There are a variety of possible individual demand characteristics of agents in a cooperative. However the consumption patterns of some customers may be more similar than others. To reflect this diversity of the agents within the simulations, we used the `CER_data_set` to identify two important classes of consumers with shared characteristics. Here, Class 1 represents consumers that consume most of their electricity during the day and have a low load at night, whereas Class 2 represents consumers that have a stable consumption during the day, but have a higher consumption at night. With these classes different scenarios with varying compositions for the consumer cooperatives can be simulated. The fraction of agent in Class 1 is defined by $p\_fracAgent \in (0, 1)$. Thus, the fraction of agents from Class 2 is $1 - p\_fracAgent$. The parameter $p\_fracAgent$ allows us to vary the composition of the cooperative across the different simulations.

To find the classes of consumers with similar characteristics the participants of the CER study were clustered using a semi-automatic approach. In a first step the participants were clustered using k-means clustering based on the information from the questionnaires and the electricity consumption data. From the questionnaires the area of business, number of employees, approximate hours of operation and whether the premises also operate during the weekends were taken into account for the clustering. From the electricity consumption data the consumption profiles of Tuesdays were used, since the planning horizon in the scope of this paper is one day and Tuesday is the most average work day with no effects like businesses which do not operate on Mondays. Note that the intention of the clustering was not to find a model that fits all participants into classes, but only to identify important classes of consumers to populate the simulations. Thus, in a second step the two classes were separated manually. The separation was based on the clustering results showing that a group of participants exists with a significant consumption profile during the day and another group exists with participants mainly consuming during the night (see Figure 8).

Modeling the flexibility of consumers: Given the class of an agent, we generate the demand profile of the agent based on the real consumption data. Although the consumption data in the `CER_data_set` is

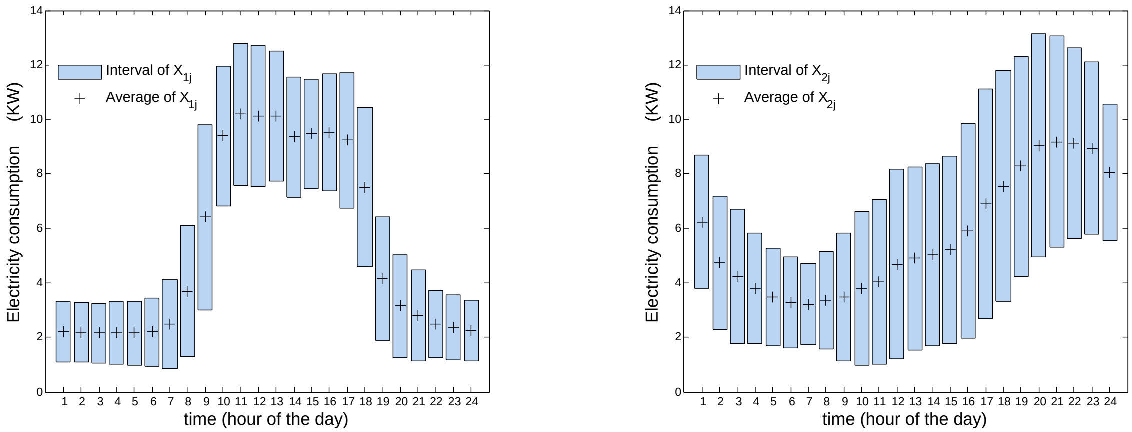

(a) Class 1.

(b) Class 2.

Figure 8: Distributions for the demand profiles, X_{Cj}, for the different consumer classes

divided into time intervals of 30 minutes, for this simulation time intervals of 60 minutes are used, because the electricity market prices are also based on intervals of 60 minutes. Let thus $j = 1$ represent the time slot "Tuesdays 12:00 AM to 1:00 AM" and $j = 24$ represent the time slot "Tuesdays 11:00 PM to 12:00 AM". Let $\overline{\overline{y}_{Cj}}$ be the mean of the average consumptions 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 the participants' average consumptions in time slot j. The demand profile for agent at time slot j is then drawn randomly from the uniform distribution $X_{Cj} \sim U\left[\overline{\overline{y}_{Cj}} - s\left(\overline{y}_{Cj}\right), \overline{\overline{y}_{Cj}} + s\left(\overline{y}_{Cj}\right)\right]$, where X_{Cj} denotes the random variable for demand at time slot j. The distributions of X_{Cj} for Class 1 are shown in Figure 8a and for Class 2 in Figure 8b. From these figures it is apparent that agents from Class 1 consume most of their electricity during the day and that agents from Class 2 have a high consumption at night. The *nominal demand* of an agent i for time slot j is generated by one outcome x_{ij} of X_{Cj}. Note that x_{ij} is not the actual demand of agent i in time slot j. It is a variable that is used for generating the total demand of an agent and other data as we will describe below. The total consumption of an agent for the whole day, can be computed by $\tau_i = \sum_{j=1}^{24} x_{ij}$.

The ability of the agents to shift their demand is expressed through the agents' upper and lower bound constraints on their electricity consumption in each time slot. To vary the flexibility of consumption between scenarios we use the parameter $p\_flexShift \in (0, 1)$, which gives the percentage by which agents can vary their demands around their x_{ij} in each time slot. This means the higher $p\_flexShift$ the further apart are the upper and lower bounds in each time slot. Based on the flexibility of consumption $p\_flexShift$ the upper and lower bounds on onsumption, can be computed as

$$\begin{aligned} \underline{r_{ij}} &= x_{ij}\left(1 - p\_flexShift\right) \\ \overline{r_{ij}} &= x_{ij}\left(1 + p\_flexShift\right) \end{aligned} \tag{14}$$

Modeling the electricity price: The prices of the electricity contract of the cooperative are defined by the values for marginal electricity prices p_j^L, for the low load, and p_j^H, for the high load, in each time slot and

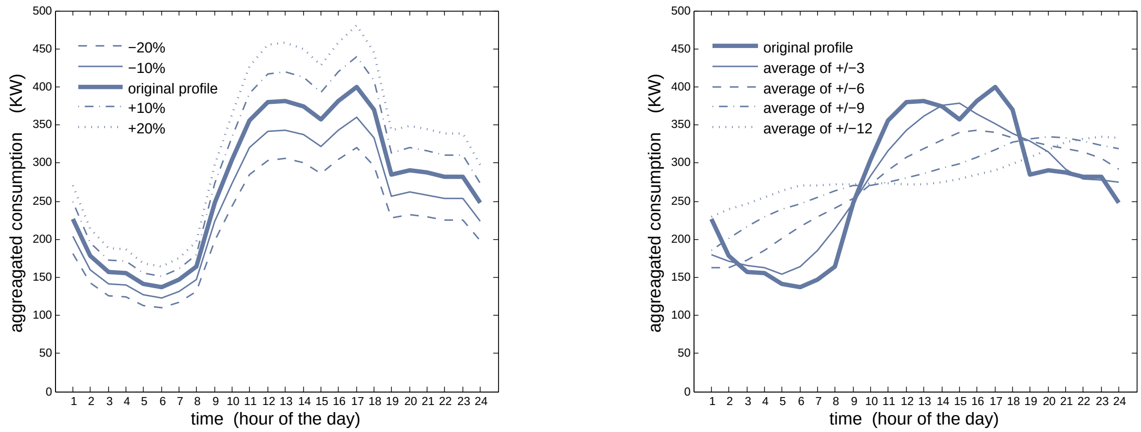

(a) As $p\_distThresh$ increases, the thresholds increase.

(b) As $p\_flatThresh$ increases, the thresholds become more flat.

Figure 9: Effect of input parameters $p\_distThresh$ and $p\_flatThresh$ on the consumption thresholds

the thresholds h_j specifying at which consumption levels the marginal prices apply. We use the real market 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. The marginal electricity prices are computed by

$$\begin{aligned} p_j^L &= mp_j \\ p_j^H &= mp_j + \max_{x \in \{1,\ldots,M\}} \{mp_x\} - \min_{x \in \{1,\ldots,M\}} \{mp_x\} \end{aligned} \tag{15}$$

Equation (15) ensures that all high marginal electricity prices are higher than the low prices in every time slot (in accordance with our assumption). These prices are fixed across all the simulations.

The thresholds of the electricity contract on the other hand are not fixed but vary across the different simulation scenarios. Therefore the thresholds are parameterized. The height of the thresholds can be varied in order to understand how the algorithm is affected by the thresholds being below or above the nominal demand. The height of the thresholds is described by $p\_distThresh \in (-0.2, 0.2)$ that specifies by what percentage the thresholds are above or below the aggregated loads. The thresholds are thus computed based on $(1 + p\_distThresh) \sum_i x_{ij}$. Figure 9a illustrates the variation of the height of the thresholds with $p\_distThresh$. The flatness of the thresholds across time slots is described by $p\_flatThresh \in (0, 24)$ that specifies the range of time slots over which the thresholds are flattened by a moving average. This means that the threshold in time slot j is calculated as the average of the loads from the time slots $j - p\_flatThresh$ to $j + p\_flatThresh$. Figure 9b shows that with an increasing $p\_flatThresh$ the thresholds get more flat. For $p\_flatThresh = 0$ the thresholds follow exactly the load profile whereas for $p\_flatThresh = 24$ the thresholds are flat, i.e. they are the same in every time slot. Thus, the thresholds are computed by

$$h_j = \frac{1}{1 + 2p\_flatThresh} \sum_{j-p\_flatThresh,\ j+p\_flatThresh} \left((1 + p\_distThresh) \sum_i x_{ij} \right) \tag{16}$$

Simulation scenarios: To create different consumer cooperative scenarios, the input parameters were varied as follows: $p\_fracAgent \in \{0, 0.25, 0.5, 0.75, 1\}$, $p\_flexShift \in \{0.1, 0.2, 0.3\}$, $p\_distThresh \in \{-0.2, -0.1, 0, 0.1, 0.2\}$ and $p\_flatThresh \in \{0, 12, 24\}$. This resulted in a total of 225 different scenarios. To reduce the variance in the results the simulations were repeated 4 times with different random seeds. Our observations are based on the mean over the 4 repetitions. Additionally, in order to analyze the convergence rate, the parameter ε was varied as $\varepsilon \in \{0.5, 1, 2\}$ and the population size was varied as $\#agents \in \{20, 40, 60, 80, 100\}$. The simulations were stopped at convergence of the algorithm, which for this simulations was defined to be reached, when the cost reduction in one iteration got less than 0.00001%, $\mathrm{C}(\mathbf{R}) / \mathrm{C}(\mathbf{R}') < 1.0000001$.

6.2 Simulation results

We will first present the effects of the different simulation parameters on the potential gain through coordination and subsequently discuss the results on convergence time. The main purpose of this analysis of the potential gains of coordination is that a coordinator of a real consumer cooperative can decide whether the implementation of the coordination algorithm is beneficial for its cooperative. In addition, an indication of the potential gain can also help to decide on the composition of the cooperative.

For the evaluation we sampled a total of 225 scenarios. A subset of the results for $p\_fracAgent = 0, 0.5, 1$ are shown in Figure 10. 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 Figure 10a the results from cooperatives consisting of consumers mainly consuming electricity at night are shown. In Figure 10b the results from the cooperatives that consist of both consumer classes in equal proportions are shown. Finally in Figure 10c the results from the cooperatives consisting of only consumers with their main consumption during the day are shown. In all the figures there are scenarios with high cost reduction (white) and scenarios with less cost reduction (dark red). Over all 225 scenarios a mean cost reduction of 2.62% was observed, with the results varying from 0% to 7.06% cost reduction. The figures also allow a first guess on the effect of the different parameters. First it can be seen that as the flexibility of shifting demand increases, i.e. $p\_flexShift \nearrow$, the cost reduction increases. Second it can be seen that as the absolute distance from the thresholds to the load profile decreases, i.e. $|p\_distThresh| \searrow$, cost reduction increases. From the composition of the cooperative, i.e. $p\_fracAgent$, as well as from the flatness of the thresholds, i.e. $p\_flatThresh$, no linear effect can be observed. However there seems to be an effect of $p\_fracAgent$ and $p\_flatThresh$ together. This absence of linear effect but presence of interaction effect is plausible, because when the thresholds follow the demand curve, the main difference of demand profiles between the agent populations get neutralized. In order to analyze these effects in detail a sensitivity analysis is performed.

6.2.1 Sensitivity analysis

The purpose of this sensitivity analysis is to understand the effect of the different input parameters on the potential gain of coordination with the aim of developing trends that allow understanding the behavior of the system in scenarios that are not simulated. For the sensitivity analysis a multiple linear regression was performed to check the main effects of the input parameters and the cost reduction through coordination. The cost reduction as fraction of the initial cost was the dependent variable (criterion), the input parameters $p\_fracAgent$, $p\_flatThresh$, $p\_flexShift$, the absolute value of the parameter $p\_distThresh$, i.e. $|p\_distThresh|$, and the interaction terms $p\_distThresh * p\_flexShift$ and $p\_fracAgent * p\_flatThresh$ were included as independent variables (predictors) in the calculation. The input parameter $p\_distThresh$ was included as absolute values, because looking at the absolute distance is similar to separately looking at

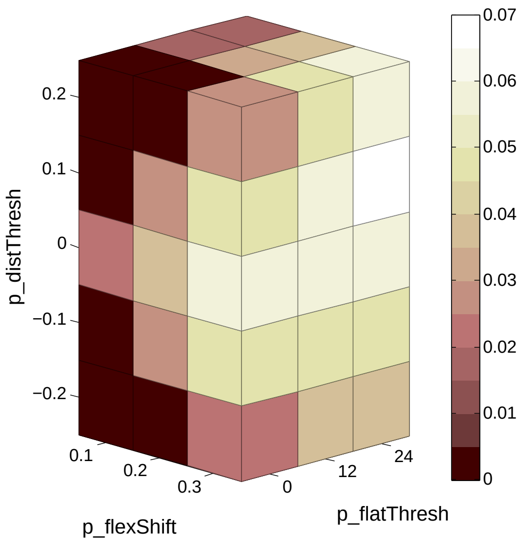

(a) Only consumers who mainly consume at night, $p\_fracAgent = 0$.

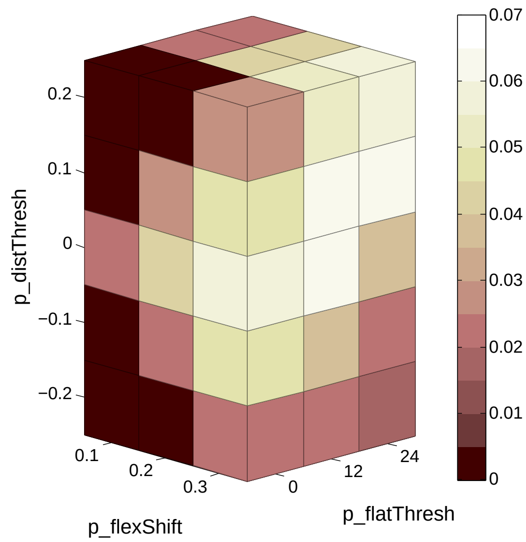

(b) Consumers from both classes have equal fractions, $p\_fracAgent = 0.5$

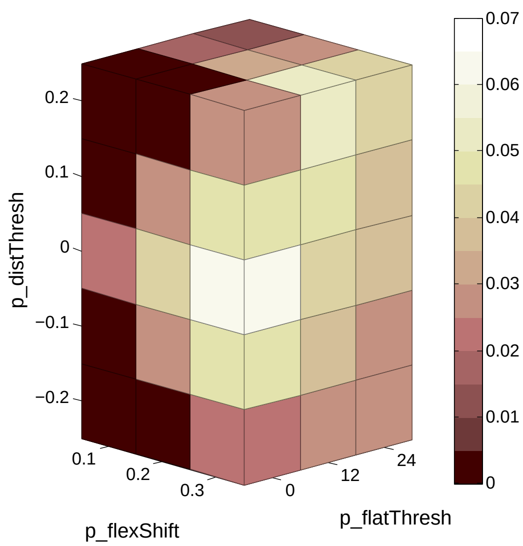

(c) Only consumers who mainly consume during day, $p\_fracAgent = 1$.

Figure 10: The cost reduction through the coordination algorithm as a function of the four input parameters $p\_fracAgent$, $p\_flatThresh$, $p\_distThresh$ and $p\_flexShift$

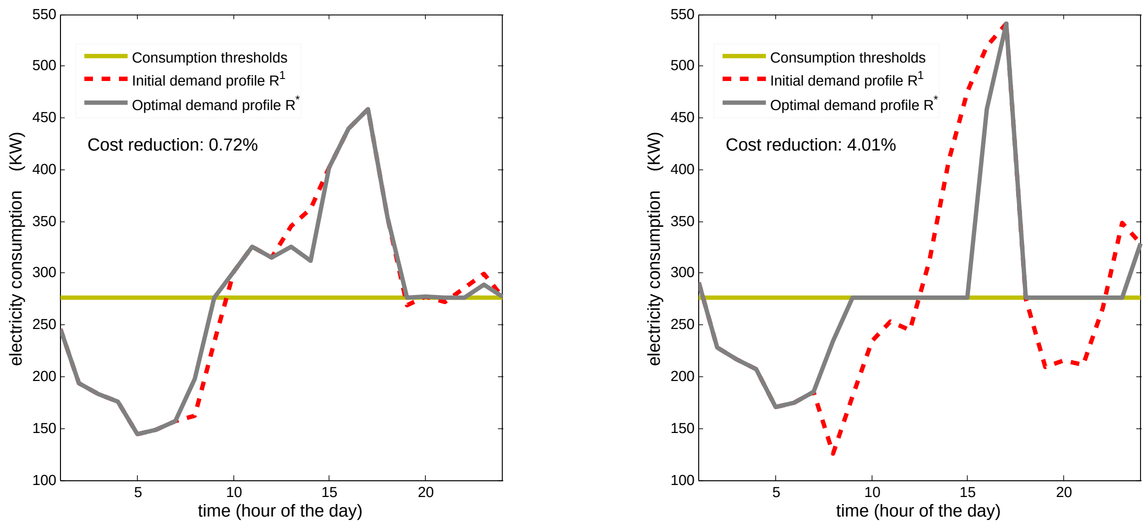

(a) Low flexibility, $p\_flexShift = 0.1$.

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

Figure 11: Example scenarios illustrating the effect of the agents' flexibility of shifting their demands ($p\_flexShift$) on the potential gain through the algorithm.

the effect of the thresholds lying below the load and the effect of the thresholds lying above the load. Other interaction effects were not included, because their effect on the model's quality was very small. The model explained 76.36% of the variance (Adjusted R^2 = .7571). The remaining variance cannot be explained by the model, because the agents' consumption constraints are generated randomly. Table 1 shows the information about the included predictors. Except for $p\_fracAgent$ with $p > 0.05$ all evaluation dimensions are highly significant predictors with $p < 0.01$.

Table 1: Multiple linear regression model $R^2 = 0.7636$ $corrR^2 = 0.7571$

	Coefficient	p-value	Standard Error
$p\_fracAgent$	0.0019	0.0910	0.0011
$p\_flatThresh$	0.0077	1.66E-05	0.0017
$p\_flexShift$	0.0114	1.35E-25	0.0010
$\|p\_distThresh\|$	-0.0072	2.08E-16	0.0008
$p\_distThresh * p\_flexShift$	0.0018	6.67E-17	0.0002
$p\_fracAgent * p\_flatThresh$	-0.0016	0.0035	0.0005

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 1 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 cost reduction.

Intuitively, with higher flexibility, more demand from the time slots with high prices can be shifted to those time slots with lower prices and consequently a higher cost reduction can be achieved. For example if the flexibility is $p\_flexShift = 0$, no demand can be shifted and thus it is impossible to reduce the cost. While higher flexibility allows for more possibilities to coordinate the demand, there is another reason for the cost reduction. Recall that the gain of coordination is the difference between the cost resulting from the uncoordinated demand profile of the cooperative $\mathbf{R}^1$ and the coordinated demand profile $\mathbf{R}^*$. For computing $\mathbf{R}^1$ the agents optimize their demand according to the hourly market prices. Consequently, the more the agents' flexibility of shifting demand increases the more the agents adapt their demand schedules to the hourly prices. However, since the agents are not coordinated every agent shifts as much demand as possible to cheap time slots. These demand shifts lead to load synchronization, where the total consumption in the cheap slots exceeds the thresholds leading to higher consumption costs (*herding behavior*). Thus, high flexibility leads not only to more possibilities to coordinate the demand, but also to an uncoordinated profile with highly synchronized demand. Figure 11 illustrates this effect in settings with low ($p\_flexShift = 0.1$) and higher ($p\_flexShift = 0.3$) flexibility. The figure shows clearly that the consumption peak is higher in the scenario with more flexibility. The area between the red dashed curve of the initial demand profile $\mathbf{R}^1$ and the optimal demand profile $\mathbf{R}^*$ in gray multiplied by the respective marginal prices shows the cost reduction through coordination. With the increasing freedom in Figure 11b it can be seen that the area between the two curves increases. Note that in the optimal demand profile the height of the peak is still the same. The reason for this is that in the time slots below the thresholds the demand is at the upper bounds of the agents. Thus, there is a need for consumption above the thresholds. According to the hourly prices as depicted in Figure 7 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 consumption 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 1 shows that the effect of $|p\_distThresh|$ is highly significant with $p < 0.01$. The negative value of the coefficient (-0.0072) shows that the further away the thresholds are to the load profile of the cooperative the less cost reduction can be achieved through coordination.

The group of consumers reduces its cost when they shift demand from time slots with the aggregated demand above the threshold to time slots with the aggregated demand below the threshold. The maximum demand that can possibly be shifted from the high consumption levels to the low consumption levels is given by the minimum of the demand above the thresholds and the unused demand below the thresholds. Intuitively, with the thresholds being above the load profile, the demand that can be shifted gets smaller, because the demand above the thresholds is smaller than the unused demand below the thresholds. For example in the extreme case where the thresholds are above the aggregated demand in every time slot, no beneficial shift exists. Similarly, with the thresholds being below the load profile fewer possibilities to shift exist, because the unused demand below the thresholds is smaller than the demand above the thresholds. Again, if the thresholds are below the aggregated demand in every time slot, no beneficial shift is possible. Consequently, the further away the thresholds are from the load profile the lower is the cost reduction that can be achieved. Figure 12 illustrates this effect in settings with low ($p\_distThresh = -0.2$), normal ($p\_distThresh = 0$) and high ($p\_distThresh = 0.2$) 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. This result is also effected by the flatness of the thresholds, i.e. $p\_flatThresh$, because the flatter the thresholds get, the larger the two described areas become. This would mean that with flatter thresholds the cost reduction increases as well. The multiple linear regression model summarized in

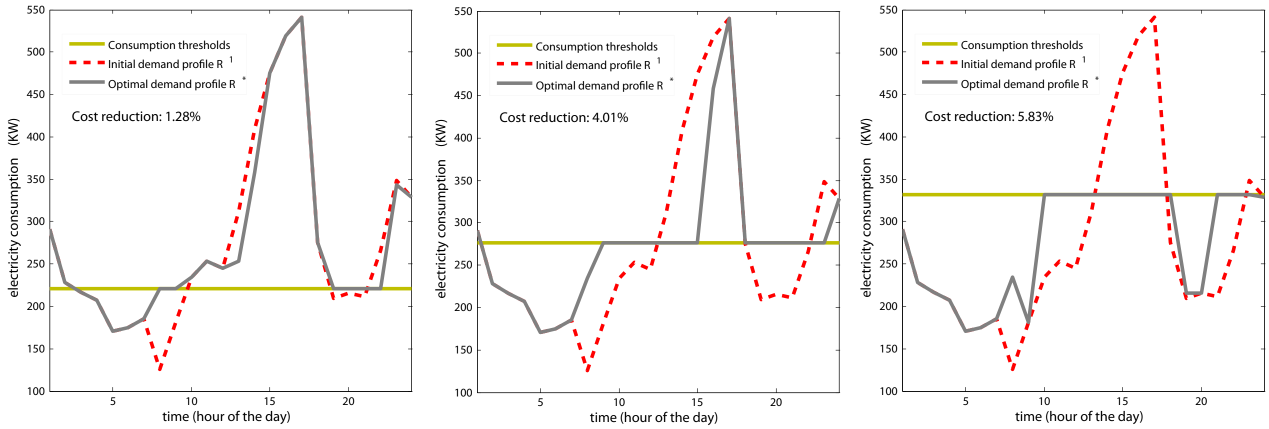

(a) Thresholds below the load, $p\_distThresh = -0.2$. (b) Thresholds at load, $p\_distThresh = 0$. (c) Thresholds above the load, $p\_distThresh = 0.2$.

Figure 12: Example scenarios illustrating the effect of the distance of the thresholds to the load profiles ($p\_distThresh$) on the algorithm's performance.

Table 1 supports this result, because it shows that the effect of $p\_flatThresh$ is has a positive coefficient (0.0077) and is highly significant with $p < 0.01$.

However, the cost reduction in the scenario with high thresholds ($p\_distThresh = 0.2$) in Figure 12c is greater than in the case of thresholds with normal height ($p\_distThresh = 0$) in Figure 12b. This indicates that in addition to the decreased cost reduction caused by the absolute distance from the demand profile there seems to be an effect that increases the cost reduction towards higher thresholds. In order to understand the positive effect towards higher thresholds, the flexibility of shifting demand has to be taken into account. The regression model summarized in Table 1 shows that the effect of the interaction term $p\_distThresh * p\_flexShift$ is highly significant with $p < 0.01$. When the flexibility is low, it might happen that not the whole intended demand can be shifted. That is because the agents hit their lower consumption bounds in the time slots above the thresholds or their upper consumption bounds in time slots below the thresholds. Since the absolute distance between the upper and lower bounds on consumption is smaller in time slots with low demand, the consumption 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 that are below the thresholds. Thus, more possibilities for shifting demand are created. Consequently, if the flexibility of shifting demand is a limiting factor, the potential cost reduction increases when the thresholds are slightly increased towards the peaks.

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 summarized in Table 1 shows that the parameter $p\_fracAgent$ is not a significant predictor for the cost reduction with $p > 0.05$. This result supports the observation that we can see cost reductions from the coordinated behavior in each of the two classes (the day consumers and the night consumers) as well as mixed groups. Figure 13 illustrates the effect of coordination in settings with only night consumers ($p\_fracAgent = 0$), a mixed group ($p\_fracAgent = 0.5$), and only day consumers ($p\_fracAgent = 1$). However, in these examples a relatively large difference in cost reduction can be observed between cooperatives with different compositions. With an increasing fraction of day consumers the cost reduction decreases much stronger than indicated by the parameter $p\_fracAgent$ in the regression

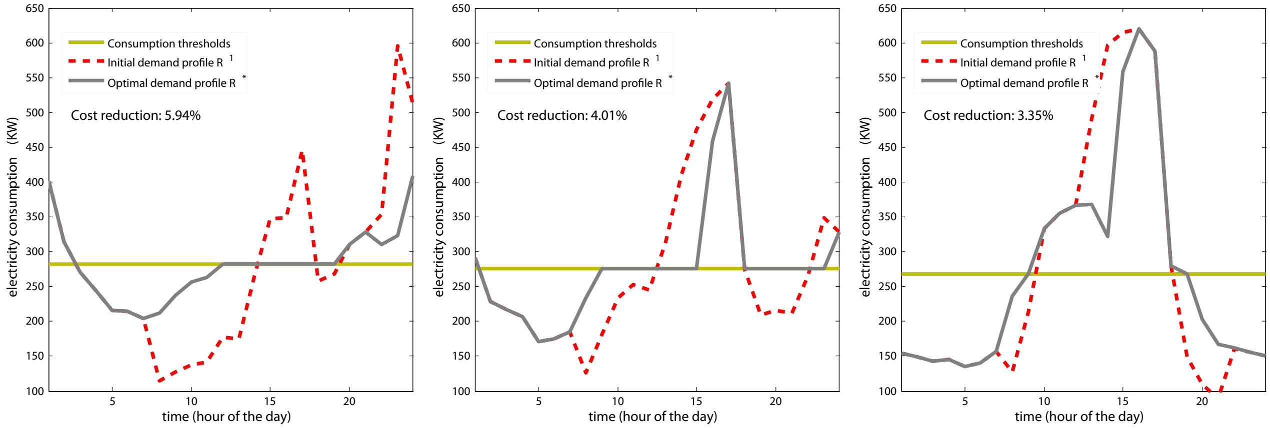

(a) Only consumers who mainly consume at night, $p\_fracAgent = 0$.

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

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

Figure 13: Example scenarios illustrating the effect of composition of the cooperative ($p\_fracAgent$) on the algorithm's performance.

model.

The reason for the effect mentioned above is that there is 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 (-0.0016) shows that as the flatness of the thresholds increases and the fraction of day consumers also increases, cost reduction decreases. In the examples shown in Figure 13 it can be seen that the thresholds are lying outside of the agents' bounds on consumption in many time slots, because the thresholds are the same in each time slot. In this situation not the thresholds, but the consumption bounds limit the potential cost reduction through coordination. Consequently, as $p\_flatThresh$ is getting larger and thus the thresholds flatter, the thresholds get out of reach of the agents' consumption bounds. Since the load profiles of the day consumers have a more significant difference between the peak and low load, this effect is stronger for groups with a larger fraction of day consumers.

6.2.2 Convergence properties of the algorithm

The main purpose of this analysis of the convergence properties of the algorithm is that a coordinator of a real consumer cooperative can assess how much time would be necessary to perform the coordination algorithm. Therefore the coordinator needs an indication of how many iterations of the algorithm are necessary in order to converge. In the following analysis the first iteration is defined as the submission of the uncoordinated initial demand profiles by the agents. Consequently the second iteration in this analysis is the iteration where the coordinator for the first time sends the virtual price signals to the agents. Please recall that in the beginning only the basic algorithm is performed and then after the basic algorithm converged the additional phase is started if necessary. Please also recall that all simulations were stopped at convergence of the algorithm, which for this simulations was defined to be reached, when the cost reduction in one iteration got less than 0.00001%, $\mathrm{C}(\mathbf{R}) / \mathrm{C}(\mathbf{R}') < 1.0000001$.

Result 4: Our simulations indicate that possibly in most cases the basic algorithm can provide good results. Although there are counter examples showing that the basic algorithm can converge to a suboptimal solution in some cases, in our simulations we found that in 80.89% of all sampled cases the basic algorithm achieved at least 99% of the potential cost reduction. Furthermore in every sampled scenario at least 94% of the potential cost reduction was achieved. This indicates that possibly in most cases the basic algorithm can provide good results.

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

In the general algorithm the variable ε defines the step size of the algorithm. Intuitively, larger steps lead to a faster convergence, but at a cost of reduced accuracy. Table 2 shows the convergence results for varying $\varepsilon \in \{0.5, 1, 2\}$. In particular, the table shows the convergence accuracy as the average share of the potential cost reduction that was achieved by the algorithm and the convergence time as the average number of iterations that were necessary to converge. For all simulations, population sizes of 40 agents were used. The results show that the number of iterations to convergence decreases as ε increases and the accuracy decreases. However, the reduction in accuracy is very small and the algorithm achieves good results for all ε.

Table 2: Convergence results for varying epsilon and constant number of agents ($= 40$)

ε	Share of potential cost reduction achieved	Average Number of iterations until convergence
0.5	0.9978	36.52
1	0.9967	26.07
2	0.9945	17.91

Result 6: Our simulations indicate that the convergence time scales linearly with the agent population size.

In our simulations the number of agents took values of 20, 40, 60, 80 and 100. Table 3 shows the convergence results. Similar to the previous result, the table shows the convergence accuracy as the average share of the potential cost reduction that was achieved by the algorithm and the convergence time as the average number of iterations that were necessary to converge. For all simulations the step size was set to $\varepsilon = 1$. The results indicate that the number of iterations to convergence increases linearly with the number of agents. The accuracy does not change much as we increase the agent population.

Table 3: Convergence results for varying number of agents and constant ε ($= 1$)

Number of agents	Share of potential cost reduction achieved	Average Number of iterations until convergence
20	0.9970	22.09
40	0.9967	26.07
60	0.9963	33.91
80	0.9962	34.44
100	0.9964	40.03

7 Conclusion and Future Work

In this paper, we presented iterative coordination algorithms to minimize the energy cost of a consumer cooperative given that information about the agents' individual demand constraints remains private. We first considered the simple, but practically relevant problem with a planning horizon of only two time slots with different electricity prices where agents have demand constraints that have to be satisfied, but they do not have any cost for shifting demands. We designed a simple iterative algorithm, where in each iteration, the coordinator computes *virtual price signals* and sends them to the consumers, who then compute their optimal consumption profiles based on this price signal and send it back to the coordinator. We designed a virtual price signal that coordinates the consumers' demand shifts such that the total cost is reduced in each iteration. We then showed that in the problem setting with only two time slots and no cost for shifting demand, this iterative algorithm converges to the optimal schedule. We then considered general settings with a planning horizon of more than two time slots with different electricity prices and with individual costs for the agents to shift demand. Considering that the basic algorithm does not guarantee the optimal solution for a general setting we designed a general iterative algorithm that includes an additional phase.This general algorithm first performs the basic algorithm as long as the solution does not get stuck and then changes to the more complicated algorithm with an additional phase. In that additional phase, in each iteration, the agents compute their marginal valuation for their electricity demand, in addition to their optimal consumption profiles and send these back to the coordinator. The coordinator uses the additional information to adapt the virtual price signals. We then showed that this general iterative algorithm converges to the optimal schedule in the general setting. These provably optimal demand scheduling algorithms for consumer cooperatives are the primary contribution of this paper. Additionally, we conducted extensive evaluations on the algorithm using multiagent simulation based on real world consumption data. Through simulations, we characterized the convergence properties of our algorithm. We also characterized the effects of different consumption characteristics of the agent population on the potential cost reduction through coordination. The results show that as the participants' flexibility of shifting their demands increases, cost reduction increases. We also observe cost reductions from the coordinated behavior in each of the consumer classes as well as mixed groups. Finally, our simulations indicate that the convergence time scales linearly with the agent population size.

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 a problem where the individual problems for the agents is no longer convex, and thus no agent can solve its individual problem optimally. The authors in (Luo, Chakraborty, & 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 consumption, 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.

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. *Smart Grid, IEEE Transactions on*, *4*(2), 866–876.

Baharlouei, Z., Hashemi, M., Narimani, H., & Mohsenian-Rad, H. (2013). Achieving optimality and fairness in autonomous demand response: Benchmarks and billing mechanisms. *Smart Grid, IEEE Transactions on*, *4*(2), 968–975.

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. *Power Systems, IEEE Transactions on*, *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. *Power Systems, IEEE Transactions on*, *25*(1), 371–380.

Conejo, A., Morales, J., & Baringo, L. (2010). Real-time demand response model. *Smart Grid, IEEE Transactions on*, *1*(3), 236–242.

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

DOE, U. (2003). Grid 2030: A national vision for electricity's second 100 years..

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

Kim, T., & Poor, H. (2011). Scheduling power consumption with price uncertainty. *Smart Grid, IEEE Transactions on*, *2*(3), 519–527.

Kirschen, D. (2003). Demand-side view of electricity markets. *Power Systems, IEEE Transactions on*, *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. *Smart Grid, IEEE Transactions on*, *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. *Smart Grid, IEEE Transactions on*, *1*(3), 320–331.

Mohsenian-Rad, A., & Leon-Garcia, A. (2010). Optimal residential load control with price prediction in real-time electricity pricing environments. *Smart Grid, IEEE Transactions on*, *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 *Computer Communications Workshops (INFOCOM WKSHPS), 2012 IEEE Conference on*, pp. 91–96. IEEE.

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

Parvania, M., & Fotuhi-Firuzabad, M. (2010). Demand response scheduling by stochastic scuc. *Smart Grid, IEEE Transactions on*, *1*(1), 89–98.

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

Philpott, A., & Pettersen, E. (2006). Optimizing demand-side bids in day-ahead electricity markets. *Power Systems, IEEE Transactions on*, *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. *Smart Grid, IEEE Transactions on*, *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. *Smart Grid, IEEE Transactions on*, *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*.

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

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.

Wu, C., Mohsenian-Rad, H., Huang, J., & Wang, A. Y. (2011). Demand side management for wind power integration in microgrid using dynamic potential game theory. In *GLOBECOM Workshops (GC Wkshps), 2011 IEEE*, pp. 1199–1204. IEEE.

Appendix

A Proof: Virtual cost is an upper bound on the total cost

Proof. Here we give the full algebraic 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 $\sum_{i:\ r'_{ij}>h_{ij}}^{N}+\sum_{i:\ r'_{ij}\le h_{ij}}^{N}$.

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

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.

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

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

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

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, $\mathrm{C}\left(\mathbf{R}'\right) \le \sum_{i=1}^{N}\mathrm{C}_i\left(\mathbf{r}'_i\right)$:

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

$\square$

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 for the two time slots 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' consumption are:

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

In addition to the electricity cost the two agents have to consider additional cost associated with the demand schedule. This additional cost for is given as:

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

Let $\mathbf{R}^{(t)}$ denote the demand profile of the agents in the tth iteration and let $t = 1$ be the initial demand profile. 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 the initial demand profiles is $\mathrm{C}\left(\mathbf{R}^{(1)}\right) = 109$. Subsequently, the coordinator computes the thresholds for the virtual price signals by comparing the initial demand profiles to the thresholds of the price function. The thresholds for the virtual price signals are

$$h_{11}(\mathbf{R}^{(1)}) = 1.8,\ h_{12}(\mathbf{R}^{(1)}) = 5.5$$
$$h_{21}(\mathbf{R}^{(1)}) = 7.2,\ h_{22}(\mathbf{R}^{(1)}) = 5.5$$

Based on the virtual price signals the agents compute their next demand profiles $\mathbf{r}_1^{(2)} = (1.5, 5.5)$, $\mathbf{r}_2^{(2)} = (4.5, 5.5)$. The cost based on this next demand profiles is $\mathrm{C}\left(\mathbf{R}^{(2)}\right) = 107.5$. Subsequently the coordinator computes the thresholds for the new virtual price signals

$$h_{11}(\mathbf{R}^{(2)}) = 2.25,\ h_{12}(\mathbf{R}^{(2)}) = 5.5$$
$$h_{21}(\mathbf{R}^{(2)}) = 6.75,\ h_{22}(\mathbf{R}^{(2)}) = 5.5$$

Since no agent can improve its cost given this virtual price signal the demand profiles keep unchanged so that $\mathbf{R}^{(T)} = \mathbf{R}^{(2)}$. Thus, the algorithm ends. However a different demand profile of the agents, $\mathbf{r}_1' = (1, 6)$, $\mathbf{r}_2' = (5, 5)$, is feasible and leads to lower total cost $\mathrm{C}(\mathbf{R}') = 107$. It follows that the algorithm has stopped in a suboptimal solution.