

Efficient Buyer Groups for Prediction-of-Use Electricity Tariffs

Valentin Robu
Heriot-Watt University
Edinburgh, Scotland, UK
v.robust@hw.ac.uk

Meritxell Vinyals, Alex Rogers and Nicholas R. Jennings
University of Southampton
Southampton, UK
{mv2y11,acr,nrj}@ecs.soton.ac.uk

Abstract

Current electricity tariffs do not reflect the real cost that customers incur to suppliers, as units are charged at the same rate, regardless of how predictable each customer's consumption is. A recent proposal to address this problem are prediction-of-use tariffs. In such tariffs, a customer is asked in advance to predict her future consumption, and is charged based both on her actual consumption and the deviation from her prediction. Prior work (Vinyals et al. 2014) studied the cost game induced by a single such tariff, and showed customers would have an incentive to minimize their risk, by joining together when buying electricity as a grand coalition. In this work we study the efficient (i.e. cost-minimizing) structure of buying groups for the more realistic setting when *multiple*, competing prediction-of-use tariffs are available. We propose a polynomial time algorithm to compute efficient buyer groups, and validate our approach experimentally, using a large-scale data set of domestic electricity consumers in the UK.

Introduction

Recent years have seen increasing interest in applying artificial intelligence techniques to solve open challenges in smart electricity grids (Ramchurn et al. 2012). One such challenge is demand-side management, and in particular designing tariffs which encourage electricity consumers to have more predictable consumption patterns.

Unfortunately, existing electricity tariffs do not have this property, as they are flat, i.e. each customer pays a fixed amount per kWh consumed, and no information about expected consumption is elicited. A recent proposal to address this problem, which follows related developments in electricity retail industry (Braithwait, Hansen, and O'Sheasy 2007), are *prediction-of-use* (POU) tariffs (Vinyals et al. 2014). In POU tariffs, a customer is asked to predict a baseline for her consumption, and is charged based both on her actual consumption, and her deviation from her baseline prediction (in the sense that units consumed in excess/short of the baseline may be charged different marginal rates).

While the work of Vinyals et al. provides an important first step, it still comes short of a complete solution. In par-

ticular, it does not consider the case of *multiple*, competing prediction-of-use tariffs, which can be offered by different electricity suppliers. To explain, different suppliers may have different acquisition strategies in purchasing electricity. Some suppliers hedge their purchases through forward contracts, and hence are able to offer a flatter tariff, although at a higher cost, to cover the cost of hedging (Team 2011). Other electricity suppliers, in contrast, may choose to buy on the forward market only the amount of electricity that their customers actually predict would be needed, and have to make up any shortfall on the balancing market, where electricity is more expensive. These different buying strategies may result in a range of prediction-of-use tariffs being offered in the market, ranging from flat to more predictive ones.

The availability of multiple tariffs also plays an important role in the success of group buying and collective switching initiatives¹. In such initiatives, a large number of consumers join together to buy electricity as a group, which improves their joint predictability and market power. Our prior work (Vinyals et al. 2014) studied the group buying game induced by a *single* prediction-of-use tariff, and showed it is concave and sub-additive, meaning all consumers buying electricity under such a tariff have an incentive to group and buy together, as a single "grand coalition".

When multiple prediction-of-use tariffs are available, however, the grand coalition does not always form. More specifically, some consumers, who have less uncertainty about their future consumption may group together to join a more predictive tariff, while others, who have more uncertainty, may prefer to form a different group, under a flatter tariff. The problem of determining the most efficient (i.e. cost-minimizing) group buying structure in this setting is complex, and the aim of this paper is to model, characterize and provide the first practical solution to this challenge.

Our work can be seen as having a connection to several strands of related literature. First, the topic of group buying has attracted considerable recent attention from the artificial intelligence and electronic commerce communities (Lu and Boutilier 2012; Meir et al. 2013). While some of this work does look at multiple, competing providers, it only

¹BigSwitch (<https://www.whichbigswitch.co.uk/>) and PeoplePower (<http://www.thepeoplespower.co.uk/>) are examples of successful real-life collective switching initiatives.

considers volume based discounts, which are rather different than the setting with uncertain demand distributions that we consider here. In the energy domain, there has been work studying coalitions or cooperatives between multiple renewable energy resources with uncertain and intermittent supply (Baeyens et al. 2011; Robu et al. 2012). On the demand side, the work of (Kota et al. 2012) proposes the concept of cooperatives for demand-side management, but they do not study the coalition stability properties of such cooperatives using tools from coalitional game theory.

From a conceptual perspective, our work is also related to the newsvendor game problem studied operations research (Muller, Scarsini, and Shaked 2002; Ozen, Norde, and Slikker 2011), but to our knowledge, this literature does not deal with multiple tariffs. Finally, there is a clear connection between our work and the work on optimal coalition structure generation (Elkind, Rahwan, and Jennings 2013), however, to our knowledge, existing work does not deal with coalitions that minimize group buying risk, such as in our work. To summarize, this work can be seen as having several contributions to the state of the art:

- First, we study the selection problem among multiple prediction-of-use tariffs, in the case when each consumer joins a tariff on her own. Assuming that the prediction error of a customer with respect to her baseline is given by a normal probability distribution, we provide a characterization of the tariff choice in terms of the coefficient of variation of each customer. We identify a natural order relationship among the set of non-dominated tariffs, and we show each tariff is cost-minimizing in an interval of the coefficient of variation dimension.
- Next, we study the efficient group formation in case that consumers elect to buy electricity as a group, under one or several tariffs. We identify a restricted case in which a strict order relationship between coefficient of variation and tariff flatness exists, although we show by counterexample such an order may not always hold in the general case. Finally, we formulate a polynomial-time dynamic programming algorithm that computes the optimal structure in such restricted case, and provides high efficiency in practice, in the experiments we performed.
- Finally, we test our example with a large-scale dataset of electricity consumers in the UK. We show, with a set of tariffs ranging from flat to highly predictive, that the tariff choice can differ significantly between the case of individual vs. group buying. More interestingly, we show that in both cases different tariffs may have non-empty market shares, thus, unlike the single tariff case, the grand coalition is not guaranteed to be the most efficient (i.e. cost minimizing) setting.

The remainder of this paper is structured as follows. First, we review prediction-of-use tariffs and introduce our multi-tariff model. Next, we provide a characterization of tariff selection, both individually and using group buying. We then provide a dynamic programming algorithm for determining efficient buyer groups, and validate it using real consumption data. The paper concludes with a discussion.

Prediction-of-Use Tariffs

In this section, we first define the basic properties of prediction-of-use tariffs following the formalization in (Vinyals et al. 2014)². The basic idea of such a tariff is that each customer is asked to provide, in advance, a prediction for her consumption (called *baseline*) during a specific time period. A prediction-of-use tariff is then defined as a tuple $\langle p, \underline{p}, \bar{p} \rangle$, where p is the baseline rate (with $p > 0$) and \underline{p} , \bar{p} are the rates for underconsumed and overconsumed units respectively (with $p > \underline{p} \geq 0$ and $\bar{p} \geq 0$).

Then the payment of a consumer with an ex-ante baseline b and an actual (ex-post) consumption x is determined as:

$$\psi(x, b) = \begin{cases} p \cdot x + \bar{p} \cdot (x - b) & \text{if } b \leq x \\ p \cdot x + \underline{p} \cdot (b - x) & \text{otherwise} \end{cases} \quad (1)$$

This tariff allows us to distinguish between a potentially lower price for the baseline consumption, and a higher penalty for the deviation. Thus, the expected payment of a customer under a prediction-of-use tariff does not depend only on her realised consumption but also on her prediction error. Following Vinyals et al., we assume that the prediction error of a customer with respect to her baseline is given by a normal probability distribution $N(\mu, \sigma)$.

Now, the *optimal baseline* that a customer should report to minimize her expected payment (b^*) is characterized as a quantile of her prediction distribution. Formally,

$$b^* = \mu + \sigma \cdot \Phi^{-1}(r^*)$$

where Φ^{-1} is the quantile function of the standard normal distribution $N(0, 1)$ (thus, it is independent of the specific μ and σ of the consumer) and $r^* = \frac{\bar{p}}{\bar{p} + \underline{p}}$ is the tariff optimal ratio.

Given this optimal baseline, the expected payment of a customer whose prediction error follows a normal distribution $N(\mu, \sigma)$ under a prediction-of-use tariff is:

$$\mathbb{E}^*[\psi(x)] = \underbrace{\mu p}_{\text{Consumption term}} - \underbrace{\sigma(\underline{p} + \bar{p}) \int_0^{r^*} \Phi^{-1}(y) dy}_{\text{Penalty term}} \quad (2)$$

We refer to the factor multiplying σ as the *penalty rate*.

Now, consider that consumers can form buying groups to reduce their expected joint payments. Let N be a set of customers that joined the same group-tariff scheme. Given a subset of customers $S \subseteq N$, let $x_S = \sum_{i \in S} x_i$ be its aggregate prediction error w.r.t. a joint baseline.

A cost game in characteristic form is obtained by assigning, for each possible buying group, its optimal expected payment under the prediction-of-use tariff. Formally,

$$c(S) = \mathbb{E}^*[\psi(x_S)] = \mathbb{E}[\psi(x_S, b^*)] \text{ for all } S \subseteq N \quad (3)$$

Vinyals et al. showed that, assuming normally distributed prediction errors, this game is concave. Concavity implies that the expected joint payment of a set of customers that join the same tariff is always less when grouping together (i.e. the grand coalition is always cost minimizing).

²Interested readers may consult this reference for the full proofs of equations in this section.

The Multi-Tariff Model

Let $\Gamma = \{\tau_1, \dots, \tau_M\}$ be a set of prediction-of-use tariffs offered in the market by competing suppliers. We consider a setting in which there is a set of rational consumers $N = \{1, \dots, n\}$, each of them needs to select an electricity tariff. For each tariff $\tau_k \in \Gamma$ we denote by $p_k, \underline{p}_k, \bar{p}_k$ the baseline, overconsumption and underconsumption rates, respectively, and by τ_k^* the optimal ratio.

The main factor driving the decision of a customer to join a tariff (on her own or through a group-buying scheme) is her *expected payment*. As observed in Eq. 2, the expected payment of a customer in a prediction-of-use tariff depends on her particular prediction error. Thus, for some customers it is expected to be more beneficial to take on higher deviation rates in order to obtain lower baseline prices. Conversely, others will be willing to pay higher baseline prices to avoid deviation penalties as much as possible. Formally, finding the tariff τ^* that minimizes the expected payment of a customer with a prediction error that follows a normal distribution $N(\mu, \sigma)$, involves solving the minimization problem:

$$\tau^* = \arg \min_{\tau_k \in \Gamma} \mathbb{E}^*[\psi_k(x)] \quad (4)$$

where $\mathbb{E}^*[\psi_k(x)]$ is the expected payment of the customer under tariff τ_k defined as in Eq. 2.

To tackle this minimization problem, we will first define the conditions under which a customer will pay less in a tariff than in another. Formally, a customer with a prediction error that follows $N(\mu, \sigma)$ expects to pay strictly less in some tariff τ_k than in tariff τ_l if and only if the following inequality holds:

$$p_k \mu - \sigma(\underline{p}_k + \bar{p}_k) \int_0^{r_k^*} \Phi^{-1}(y) dy < p_l \mu - \sigma(\underline{p}_l + \bar{p}_l) \int_0^{r_l^*} \Phi^{-1}(y) dy$$

Recall that r_k^*, r_l^* and Φ^{-1} are parameters independent of μ and σ . By rearranging terms we obtain:

$$\Delta_{k,l}^P \sigma < \Delta_{l,k}^B \mu, \quad (5)$$

where $\Delta_{l,k}^B = p_l - p_k$ is the difference in baseline price between tariff τ_l and τ_k (where $\Delta_{k,l}^B = -\Delta_{l,k}^B$) and

$$\Delta_{k,l}^P = -(\underline{p}_k + \bar{p}_k) \int_0^{r_k^*} \Phi^{-1}(y) dy + (\underline{p}_l + \bar{p}_l) \int_0^{r_l^*} \Phi^{-1}(y) dy$$

is the difference in penalty rate between τ_k and τ_l .

Definition 1. A tariff τ_l is said to be dominated in expectation iff there exists another tariff $\tau_k \in \Gamma$ such that the expected payment under τ_k is lower than those in τ_l irrespective of the μ and σ of the expected prediction error.

In terms of the above notation, a tariff τ_l is dominated in expectation by τ_k if $\Delta_{k,l}^B < 0$ and $\Delta_{k,l}^P < 0$.

Note that a stronger notion of *ex-post* dominance can also be defined, i.e. if $p_k < p_l$, $\underline{p}_k < \underline{p}_l$, $\bar{p}_k < \bar{p}_l$, then a consumer will pay less under tariff τ_k , regardless of the realised consumption. However, here we are interested in “in expectation” dominance, because the choice of tariff occurs in expectation.

When faced with the ex-ante choice, no risk-neutral consumer will choose a dominated tariff, hence such a tariff gets no market share and w.l.o.g. we can omit it from further analysis. Also w.l.o.g., we can omit duplicate tariffs, i.e. tariffs which have the same baseline and penalty rates. Henceforth, when referring to Γ , we consider eliminated from it all tariffs that are dominated in expectation or duplicates. After these eliminations, for any pair of remaining tariffs $\tau_k, \tau_l \in \Gamma$ it must hold that: if $p_k < p_l$ then $\Delta_{k,l}^P \geq 0$. Now, we can define the relation of *flatness* among tariffs.

Definition 2. Given two tariffs $\tau_k, \tau_l \in \Gamma$ we say that τ_k is *less flat* than τ_l , denoted as $\tau_k \prec^f \tau_l$, if τ_k has a lower baseline price and a higher penalty rate than τ_k (formally, if $\Delta_{k,l}^B < 0$ and $\Delta_{k,l}^P > 0$).

Observe that when applied to a set of tariffs Γ , after we eliminate the dominated and duplicate tariffs, the flatness relationship defines a total ordering among remaining tariffs in Γ (i.e. it is transitive and symmetric).

Now, given two tariffs $\tau_k, \tau_l \in \Gamma$ such that $\tau_k \prec^f \tau_l$ we can rearrange terms in Eq. 5 to obtain:

$$\frac{\sigma}{\mu} < \left[\lambda_{k,l} = \frac{\Delta_{l,k}^B}{\Delta_{k,l}^P} \right] \quad (6)$$

Note that the left side of Eq. 6 is known as **coefficient of variation** (i.e. the ratio between the standard deviation and the mean). The parameter on the right side (which ultimately depends only on the parameters $p, \underline{p}, \bar{p}$ of tariffs τ_k and τ_l) we call the **inter-tariff threshold** between τ_k and τ_l . The following corollary follows directly from previous derivations:

Corollary 1. Given two tariffs $\tau_k \prec^f \tau_l$, there exists a unique threshold $\lambda_{k,l} \in \mathbb{R}$ such that the expected payment of any customer is less in tariff τ_k than in tariff τ_l iff $\frac{\sigma}{\mu} < \lambda_{k,l}$. Otherwise her expected payment is lower in tariff τ_l .

Market Segmentation with Individual Choice

We characterize the partition of the pool of customers in the case each customer chooses to join a prediction-of-use tariff based only on its own expected consumption.

Given Corollary 1, for some non-dominated, duplicate-free set of tariffs Γ we can define a set of $\frac{|\Gamma|(|\Gamma|-1)}{2}$ inter-tariff thresholds, one for each pair of tariffs $\tau_k, \tau_l \in \Gamma$ such that $\tau_k \prec^f \tau_l$. Arranging these thresholds over the coefficient of variation axis, they define a total of $\frac{|\Gamma|(|\Gamma|-1)}{2} + 1$ intervals ($\frac{|\Gamma|(|\Gamma|-1)}{2} - 1$ between thresholds plus two at the axis ends). We label these thresholds in ascending order λ^s where s is the position of the threshold in the ordering.

Lemma 1. Given a set of tariffs Γ , for all $s = 1 \dots \frac{|\Gamma|(|\Gamma|-1)}{2} + 1$, the expected payment of any customers with a coefficient of variation within the range $(\lambda^{s-1}, \lambda^s)$ is minimized under the same tariff.

Proof. Towards a contradiction, consider two consumers $i, j \in N$ with $\frac{\sigma_i}{\mu_i}, \frac{\sigma_j}{\mu_j} \in (\lambda^{s-1}, \lambda^s)$, where i joins tariff τ_k , while j joins τ_l ($k \neq l$). But this implies the existence of

another threshold $\lambda_{k,l} \in (\lambda^{s-1}, \lambda^s)$, which is not possible as the set of all inter-tariffs thresholds is determined. \square

Thus, for each interval, there is one tariff that dominates all the others, for all customers with a $\frac{\sigma}{\mu} \in (\lambda^{s-1}, \lambda^s)$. However, it may not be necessary to keep all these thresholds to characterize the partition because some of the neighbouring intervals will merge, as one tariff dominates in two or more adjoining intervals. Formally, we denote by the set Γ_{ND} the subset of tariffs in Γ that dominate in at least one interval.

It remains to be shown that, if a tariff τ_k in Γ_{ND} dominates in a set of intervals, all these intervals must be adjoining (i.e. consecutive). Again, we prove this by contradiction. Consider a tariff τ_k in Γ_{ND} such that it dominates in two non-consecutive intervals. Then it means that there is at least one interval in the middle in which another tariff τ_l , $\tau_l \neq \tau_k$, dominates. But this would lead to a contradiction with Corollary 1, which states the switch from τ_k to τ_l is defined by a unique point $\lambda_{k,l}$.

Market Segmentation with Group Buying

In this section, we extend our analysis to the case when customers group together to purchase electricity as a coalition (aka group), under any tariff in set Γ . We define the corresponding multiple prediction-of-use tariff cost game $G = \langle N, c \rangle$, where $c : 2^N \rightarrow \mathbb{R}$:

$$c(S, \Gamma) = \min_{\tau_k \in \Gamma} \mathbb{E}[\psi_k(x_S)] \text{ for all } S \subseteq N. \quad (7)$$

In this paper, we take a cooperative game theory approach to this problem, and our goal is to determine which is the optimal buying group structure that minimizes the total cost to all customers (where “cost” is defined as total expected payments under the selected set of tariffs). In general, coalition structure generation involves finding the exhaustive disjoint partition of customers into groups $\vec{S} = \{S_1, \dots, S_m\}$ such that the total cost, $\sum_{k=1}^m c(S_k, \Gamma)$, is minimized. However, as stated by the next lemma, in our case at most one group per tariff will form.

Lemma 2. *The optimal structure of a pool of customers under a set of tariffs Γ can be expressed in terms of $|\Gamma|$ groups, $\{S_1, \dots, S_{|\Gamma|}\}$, being for all $k = 1 \dots |\Gamma|$ S_k a set of customers (possibly empty) that will group under tariff $\tau_k \in \Gamma$.*

Proof. Consider that the optimal structure contains two groups S, S' that join the same tariff τ_k . This leads to a contradiction because, given the concavity (and subadditive) results proved in (Vinyals et al. 2014) for the single tariff game, customers choosing the same tariff reduce their joint expected payment by grouping ($c(S \cup S', \tau_k) < c(S, \tau_k) + c(S', \tau_k)$). \square

Henceforth, we can use the notation $S_k = \vec{S}(k) \subseteq N$ to refer to consumers assigned to τ_k in the structure \vec{S} .

Unfortunately, despite this bound on the maximum number of groups, the group buying structure generation problem is still a complex one because the optimal tariff group of each customer depends on the set of customers that already

joined the group. Hence, in order to tackle this problem, we first single out a tractable grouping case for which the order property based on the coefficient of variation holds.

Optimal Group Structure with Equal σ s

As stated by the next lemma, for the restricted case with (equal σ -s), the grouping follows the same order than in the individual choice case.

Lemma 3. *Consider the case when all customers have the same standard deviation error prediction, σ ($\sigma_1 = \dots = \sigma_n = \sigma$). Then, for any two customers $\forall i, j \in N$ with $\frac{\sigma_i}{\mu_i} < \frac{\sigma_j}{\mu_j}$, if the optimal group structure assigns j to S_k , corresponding to tariff $\tau_k \in \Gamma$, and i to S_l , corresponding to $\tau_l \in \Gamma$ (where $k \neq l$), then it must hold that $\tau_k \prec^f \tau_l$.*

Proof. Consider that the best group buying structure, \vec{S}^* , does not satisfy Lemma 3. This means that there are at least two customers i and j such that $\frac{\sigma_i}{\mu_i} < \frac{\sigma_j}{\mu_j}$ and $i \in S_l^*, j \in S_k^*$ where $\tau_k \prec^f \tau_l$. Towards a contradiction, consider another structure \vec{S}' that contains the same groups as \vec{S}^* , except that in coalitions S'_k and S'_l customers i and j are interchanged. Formally: $S'_m = S_m^*$ for all $m \neq l, k$, $S'_k = S_k^* \setminus \{j\} \cup \{i\}$ and $S'_l = S_l^* \setminus \{i\} \cup \{j\}$. To compare the cost of \vec{S}^* and \vec{S}' , we only need to compare the cost of groups S_l^*, S_k^* with S'_l, S'_k since the rest of groups are the same. We compare these groups based on their two cost components: the consumption term and the penalty term. Since all customers have the same σ and the number of customers in S'_l is the same than in S_l^* then the penalty term is the same among these groups. The same applies between groups S'_k and S_k^* . Regarding consumption terms, the difference between the cost of S_k^* and S'_k is $(\mu_j \cdot p_k - \mu_i \cdot p_k)$ and the difference between the cost of S_l^* and S'_l is $(\mu_i \cdot p_l - \mu_j \cdot p_l)$. Adding both differences we obtain $(p_l - p_k) \cdot (\mu_i - \mu_j)$.

However $p_l > p_k$ (because of the flatness relation) and $\mu_i > \mu_j$ (because of the ratio inequality and $\sigma_i = \sigma_j$), so the cost of \vec{S}' is less than \vec{S}^* , leading to a contradiction. \square

Intuitively, Lemma 3 states that it is optimal to group customers which are poorer predictors (i.e. higher $\frac{\sigma}{\mu}$) in either the same or flatter tariffs than customers that are better predictors (i.e. lower $\frac{\sigma}{\mu}$). However, the partition when grouping can not be directly characterized by the inter-tariff thresholds (which only depend on the tariff themselves) as in the individual case because the coefficient of variation of a customer group depends on other customers that joined that group. Nevertheless, given that such ordering relationship exists, the problem of deciding the optimal coalition structure is not a combinatorial one (as in the unrestricted case). Indeed, in a further section we propose an algorithm that is guaranteed to find such optimal coalition structure in polynomial time.

Group Structure in Unrestricted Cases

Given the tractability results stated for the equal σ 's case, one may wonder if the same coefficient of variation ordering holds for the unrestricted case. Unfortunately, as we next

	p	\bar{p}	$\bar{\bar{p}}$
τ_A	0.17	0.05	0.13
τ_B	0.19	0.01	0.02
τ_C	0.20	0	0




Figure 1: Tariffs used in Example 1 (in £/kWh) and the coefficient of variation intervals with $\frac{\sigma}{\mu}$ for customers $A_1 - A_{10}$ (taken individually) and the group formed by $A_1 - A_9$.

show by means of a counterexample, such ordering restriction does not guarantee optimality in the general case.

Example 1. Consider tariffs τ_A, τ_B, τ_C , their inter-threshold intervals ($\lambda_{AB} = 0.4$ and $\lambda_{BC} = 0.92$) and 10 customers as detailed in Fig. 1. Customers $A_1 - A_9$ are small consumers, whose individual prediction error follows $N(\mu = 1, \sigma = 1)$, with $\frac{\sigma}{\mu} = 1$. Customer A_{10} is a much larger consumer with a prediction error following $N(\mu = 20, \sigma = 15)$ and hence with a lower coefficient of variation ($\frac{\sigma_{10}}{\mu_{10}} = 0.75$). When customers join tariffs on their own, $A_1 - A_9$ will join the flattest tariff τ_C (since $\frac{\sigma_i}{\mu_i} > \lambda_{BC} \forall i = 1 \dots 9$), while A_{10} will join τ_B (since $\lambda_{AB} < \frac{\sigma_{10}}{\mu_{10}} < \lambda_{BC}$). Conversely, in the optimal group buying structure, $A_1 - A_9$ will group together under the most predictive tariff τ_A (since $\frac{\sigma_{\{1:9\}}}{\mu_{\{1:9\}}} = \frac{\sqrt{9}}{9} < \lambda_{AB}$), while A_{10} remains alone under τ_B . This is the optimal (cost minimizing) structure for this example: the grand coalition can not benefit from tariff τ_A (as $\frac{\sigma_{\{1:10\}}}{\mu_{\{1:10\}}} = 0.53 > \lambda_{AB}$, and it results in a higher total payment).

Example 1 is of interest because it shows that the monotonic order of assigning customers to tariffs does not always hold: A_{10} has a lower coefficient of variation than those of $A_1 - A_9$ individually, yet the optimal group buying structure still places her in a flatter (i.e. less predictive) tariff.

While we have shown that such examples can be constructed, we also found they rely on one or a few customers having a very high share of the total consumption, which is rarely the case in practice (at least in the real dataset we have examined). Thus, albeit in the general case there is no theoretical guarantee that the dynamic programming approach that we propose in the next section returns the optimal partition, the quality of the solution is nearly-optimal for real-world distributions of consumption. An empirical examination of such a case is provided in the experimental section.

Dynamic Programming Search for Optimal Group Structure

We present an algorithm that given a set of tariffs Γ computes the best group structure $\{S_1^*, \dots, S_{|\Gamma|}^*\}$ such that $\forall i, j \in N$ with $\frac{\sigma_i}{\mu_i} < \frac{\sigma_j}{\mu_j}$ $i \in S_k, j \in S_l, k \neq l$ then

Algorithm 1 Finding P^* in the $\frac{\sigma}{\mu}$ interval with monotonic group tariff selection

Require: $\{1, \dots, n\}$ (customers ordered by ascending coefficient of variation) and Γ (tariffs ordered by increasing flatness)

```

1: for  $i = 1$  to  $n$  do
2:    $C[i] \leftarrow c(\{1 : i\}, \Gamma)$ ;
3:    $T[i] \leftarrow k \mid \Gamma(k)$  is the optimal tariff for group  $\{1 : i\}$ 
4:    $P[i] \leftarrow 0$ ;
5:   if  $i == n$  then
6:      $\bar{\tau} = |\Gamma| - 1$ ; /*At least one tariff must remain*/
7:   else
8:      $\bar{\tau} = |\Gamma| - 2$ ; /*At least two tariffs must remain*/
9:   end if
10:  for  $j = 1$  to  $i - 1 \mid T[j] \leq \bar{\tau}$  do
11:     $c' \leftarrow c(\{j + 1 : i\}, \Gamma) + C[j]$ ;
12:     $\tau' \leftarrow k \mid \Gamma(k)$  is the optimal tariff for group  $\{j + 1 : i\}$ 
13:    if  $c' < C[i]$  and  $\tau > T[j]$  /*Update*/ then
14:       $(C[i], P[i], T[i]) = (c', j, \tau')$ ;
15:    end if
16:  end for
17: end for
18:  $P^*, T^* \leftarrow \{\}; i \leftarrow n$ ;
19: while  $i > 0$  do
20:    $P^* \leftarrow P^* \cup \{P[i] + 1 : i\}$ 
21:    $T^* \leftarrow T^* \cup T[i]$ 
22:    $i \leftarrow P[i]$ ;
23: end while
24: return  $(C[n], P^*, T^*)$ 

```

$\tau_k \prec^f \tau_l$. The core of our approach takes its starting point in a Dynamic Programming (DP) algorithm formulated in (Rothkopf, Pekec, and Harstad 1995) to solve combinatorial auctions with neighbouring assets. When applied to our problem, this algorithm computes the best group buying structure when groups of consecutive consumers in the coefficient of variation interval are considered. Although this DP algorithm has a deterministic polynomial complexity³, it does not exploit the strictly monotonic tariff selection that characterizes our problem. To this end, we extend the algorithm to restrict search to structures in which groups are assigned to strictly increasing flatter tariffs. Thus, we develop an enhanced version of the DP algorithm that takes into account the specific structure of the problem. Specifically, our algorithm constrains search such that multiple groups cannot form around the same tariff (c.f. Lemma 2) and agents are assigned to groups in monotonic order of their flatness.

The pseudocode of the method is given in Algorithm 1. This algorithm requires, as an input, the set of customers ordered by ascending coefficient of variation and the set of available tariffs ordered by increasing flatness. Then, for each customer $i = 1 \dots n$, the algorithm computes the subproblem corresponding to the best group-buying structure among i and other more predictable consumers (lines 1-17). The cost of such a subproblem is stored in $C[i]$, the first customer not included in the group of i is stored in $P[i]$, while the best tariff for such group is stored in $T[i]$. At the be-

³Quadratic computational requirements and linear computational requirements to the number of consumers (Rothkopf, Pekec, and Harstad 1995).




Figure 2: Market segmentation when customers join tariffs on their own (left); and through group-buying schemes (right).

gining of each iteration, the corresponding group-buying structure is initialized to one in which all customers join a single group (lines 2-4). Then, the algorithm computes the minimum number of tariffs that must not be selected to allow the formation of a new group (i.e. one in the case of the last customer and two for the rest). Then, for each $j < i$ such that the subproblem j left insufficient tariffs ($T[j] < \bar{\tau}$), the cost of such structure is compared (and updated if it is the case) with the cost of the group $\{(j+1) : i\}$ plus the cost of the best group structure among the first j customers, $C[j]$ (lines 10-16). Moreover, the structure is only updated if the tariff assigned to the last formed group, $T[j]$, is less flatter than those assigned to the new group, τ (line 13). By so doing, partitions that do not satisfy the flatness order among tariffs are discarded.

At the end of this process, the cost of the best group-buying structure explored during the execution of the algorithm is found in $C[n]$. Finally, the algorithm executes an iterative procedure (lines 18-23) to recover the groups that form such a structure (P^*) and the set of tariffs to which they are assigned (T^*). Although Algorithm 1 has the same worst-case quadratic computational complexity as the basic DP algorithm, in the practical instances (such as those considered in the experiments), the number of operations is significantly reduced. In particular, for the two tariffs case ($|\Gamma| = 2$), Algorithm 1 reduces to linear complexity (i.e. $\bar{\tau}$ is 0 for all iterations with exception of the last).

Experimental Evaluation

Our experimental analysis makes use of a large dataset of around 3000 households (i.e. customers) in the UK. For each customer, the dataset included her electricity consumption for every half hour during a three-month period. We take the sample mean over the consumption realizations of each customer as a point estimate for her μ and the standard deviation as an unbiased estimator for her σ .

The evaluation considers three tariffs (F, P and P+) de-

tailed as follows. Tariff F (*Flat*), corresponds to a flat tariff in which customers pay a fixed price (£0.205) per kW consumed. Tariff P (*Predictive*) reduces the baseline price of tariff F at the cost of charging a penalty of £0.01/£0.03 for each kW underconsumed/overconsumed respectively. Finally, tariff P+ (*Highly Predictive*) offers the lowest baseline price but severely penalizes any imbalance (with penalties of £0.17/£0.26 per Kw underconsumed/overconsumed). These figures ensure the baseline prices of these tariffs are roughly similar to the prevalent prices in the UK tariffs (see (Team 2011)). However, the penalty rate of tariff P+ is set particularly high to favour the formation of multiple buying groups, since they are of more interest given the focus on group structure generation of this paper. Next, we analyze the resulting market segmentation when consumers choose tariffs on their own and through group buying schemes.

Market segmentation with individual customer choice.

Fig. 2 (left) shows the coefficient of variation of each customer (ordered from the most to the least predictable), in the colour of its preferred tariff when joining individually. The figure also shows the share of customers that each tariff obtains as a pie chart. Observe that the first 2366 customers with the lowest coefficient of variation prefer to join tariff P, whereas the remaining 622 customers prefer tariff F (i.e. tariff P+ gets no market share). Therefore, although a high percentage (i.e. 79.18%) of customers could benefit from a lower baseline price through tariff P, there remains a significant percentage (i.e. 20.82%) that are too unpredictable to benefit from this tariff, and choose the flat tariff F. Most importantly, the risk of penalty imbalance of tariff P+ remains too high for any customer, when taking an individual decision. The total expected payment of the customers per day under this market segmentation is £8097.

Market segmentation with group buying. We use the DP algorithm presented in the previous section to compute the most efficient group buying structure, under the constraint that customers with increasing coefficient of variation are

assigned to flatter tariffs. While, as discussed before, this does not provide a theoretical guarantee of optimality in the general case, we found this constraint reasonable for the real dataset employed here. Moreover, while the DP algorithm has a worst-case quadratic performance in the number of customers, because of the additional structure of the problem it exploits, we found that in practice the number of operations is much less (around 10^5 or 2%) of the number of operations prescribed by the quadratic bound (which is $|N|^2/2 = 4.5 \cdot 10^6$). Fig. 2 (right) shows the coefficient of variation of the group joined by each customer (ordered from the most to the least predictable), again in the colour of the selected tariff. As before, the figure includes a pie chart of the share of the market obtained by each tariff. Observe that, by allowing grouping, the first 2758 customers, with the lowest coefficient of variation, prefer to group together under tariff P^+ whereas the remaining 230 customers prefer to group under tariff P . The total expected payment of the customers per day under this market segmentation is £7952. Hence, there is a payment reduction for grouping this set of customers under such computed group-buying structure of £145 per day. It is worth noticing that the structure formed in the grouping case is completely different to the one from the individual choice case. Now, tariff F gets no market share, and the reduction of risk imbalance from grouping lets a high percentage of the customers (i.e. 92.30%) benefit from the most predictive P^+ tariff.

Conclusions and Future Work

This work provides a thorough analysis of market segmentation with multiple prediction-of-use tariffs. Besides giving a formal characterization of the problem, both for a single individual and group perspective, we provide an enhanced DP algorithm for determining efficient group buying structures, and validate our approach on a large dataset of electricity consumers from the UK.

There are several directions we plan to explore in future work. From a practical perspective, it would be interesting to explore the connection between prediction-of-use and the more widely known time-of-use (TOU) tariffs (Braithwait, Hansen, and O'Sheasy 2007).⁴ POU tariffs are a new, different concept than TOU tariffs, but the two could be naturally combined. For example, some retailers may offer, based on their hedging strategy in the forward/balancing market, tariffs in which customers are only asked to predict their future consumption only for peak-time periods. Such tariffs could use a prediction-of-use structure for the peak times, but a flat tariff structure during the rest of the day.

From a broader AI perspective, while coalitional formation and structured coalitional games are a growing area of research in the AI community (e.g. (Chalkiadakis, Elkind, and Wooldridge 2011; Chalkiadakis, Markakis, and Jennings 2012)), there are few works that discuss what types of algorithms and techniques are relevant and applicable in concrete domains. This work can be seen as opening the way

for a wider range of techniques from coalition formation and coalitional game theory to be applied in the tariff design domain, and we intend to pursue some of these in future work.

References

- Baeyens, E.; Bitar, E.; Khargonekar, P. P.; and Poolla, K. 2011. Wind energy aggregation: A coalitional game approach. In *CDC-ECE*, 3000–3007.
- Braithwait, S.; Hansen, D.; and O'Sheasy, M. 2007. *Retail Electricity Pricing And Rate Design In Evolving Markets*. Edison Electric Institute.
- Chalkiadakis, G.; Elkind, E.; and Wooldridge, M. 2011. *Computational Aspects of Cooperative Game Theory*. Morgan & Claypool Publishers.
- Chalkiadakis, G.; Markakis, E.; and Jennings, N. R. 2012. Coalitional stability in structured environments. In *AAMAS*, 779–786.
- Elkind, E.; Rahwan, T.; and Jennings, N. 2013. Computational coalition formation. *Multiagent Systems (2nd Edition)*, MIT Press 329–380.
- Kota, R.; Chalkiadakis, G.; Robu, V.; Rogers, A.; and Jennings, N. R. 2012. Cooperatives for demand side management. In *ECAI*, 969–974.
- Lu, T., and Boutilier, C. 2012. Matching models for preference-sensitive group purchasing. In *ACM EC*, 723–740.
- Meir, R.; Lu, T.; Tennenholtz, M.; and Boutilier, C. 2013. On the value of using group discounts under price competition. In *AAAI*.
- Muller, A.; Scarsini, M.; and Shaked, M. 2002. The newsvendor game has a nonempty core. *GEB* 38:118–126.
- Ozen, U.; Norde, H.; and Slikker, M. 2011. On the convexity of newsvendor games. *International Journal of Production Economics* 133:35–42.
- Ramchurn, S. D.; Vytelingum, P.; Rogers, A.; and Jennings, N. R. 2012. Putting the 'smarts' into the smart grid: a grand challenge for artificial intelligence. *Communications of the ACM* 55(4):86–97.
- Robu, V.; Kota, R.; Chalkiadakis, G.; Rogers, A.; and Jennings, N. R. 2012. Cooperative virtual power plant formation using scoring rules. In *AAAI*, 370–376.
- Rothkopf, M. H.; Pekec, A.; and Harstad, R. M. 1995. Computationally manageable combinatorial auctions. *Management Science* 44(8):1131–1147.
- Team, G. M. 2011. Electricity and gas supply market report. Technical Report 176/11, The Office of Gas and Electricity Markets (Ofgem).
- Vinyals, M.; Robu, V.; Rogers, A.; and Jennings, N. 2014. Prediction-of-use games: a cooperative game theory approach to sustainable energy tariffs. In *AAMAS*. To appear. Available from <http://eprints.soton.ac.uk/360819/>.

⁴In TOU tariffs, customers are charged different marginal rates per kWh, depending on the time of day in which the consumption occurs.