

Prediction-of-use games: a cooperative game theory approach to sustainable energy tariffs

Meritxell Vinyals

Valentin Robu

Alex Rogers

Nicholas R. Jennings

School of Electronics and
Computer Science
University of Southampton
Southampton, UK
{mv2y11, vr2, acr, nrj}@ecs.soton.ac.uk

ABSTRACT

Current electricity tariffs do not reflect the real costs that a customer incurs to a supplier, as units are charged at the same rate, regardless of the consumption pattern. In this paper, we propose a prediction-of-use tariff that better reflects these costs, which asks customers to predict a baseline consumption, and charges them based both on their actual consumption, and the deviation from their prediction. We show how under this tariff no customer would have an incentive to consume in excess of their actual needs, and derive closed form expressions for their optimal prediction and expected payments.

Second, using principles from cooperative game theory, we study how customers can collectively reduce their potential deviation by aggregating under a group-buying scheme. We prove that the associated cost game is concave, which means grouping reduces the total expected bill and that this payment can be fairly allocated among customers by their Shapley values. Third, considering a model where customers can join the group online, we propose marginal payment allocation schemes that incentivise them to commit early, thus preventing start-up inertia. Finally, we validate our model using real data from a set of 3000 consumers from the UK.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems

Keywords

Cooperative games, group buying, collective energy tariffs, payment redistribution

1. INTRODUCTION

Recent years have seen significant efforts to switch to a more sustainable *smart energy grid*, which incentivises the development of renewable energy resources, as well as lower levels of consumption. However, these efforts have sometimes met with resistance from consumers because they can lead to a rise in electricity prices. Moreover, with such price rises millions of households are being pushed towards fuel poverty. Given this, we argue, a priority of current research should be to develop tools that empower consumers in interacting with energy providers to obtain the best deals.

Existing electricity tariffs are not well-suited to dealing with these challenges. In most cases, a customer is charged at a flat rate, based

only on the number of units actually consumed and regardless of their consumption pattern. However, this matches poorly with the structure of the costs that energy suppliers face [17]. In more detail, in most countries, suppliers purchase electricity supplied to their customers through forward contracts, in which they commit to acquire a baseline demand, to be supplied at a later period. These forward purchases can be done through long term bilateral contracts (for electricity bought several months in advance), or through the spot markets (for electricity bought in a shorter time horizon, e.g. a day in advance). Any deviation between the baseload and the actual consumption is resolved through the *balancing market*, where prices, especially at peak times, can be much higher than those obtained through forward contracts. For example, if a supplier has a shortfall of electricity acquired, it may have to pay considerably more in the balancing market for the extra units. Equivalently, units bought in excess also have a cost, because typically the prices per unit sold back in the balancing market, especially in periods of low demand, are also much lower.

From this, we can see that each consumer has a hidden *predictability cost* for the supplier: consumers with a stable consumption base and predictable loads involve lower costs than consumers with an unpredictable future demand. Suppliers recover these additional predictability costs by increasing the unit price of flat electricity tariffs. Therefore, current tariffs involve a considerable degree of hidden *cross-subsidisation*: predictable customers end up paying more to compensate the risk of others with more unpredictable consumption patterns.

In this paper we address these challenges by proposing a tariff that better matches the structure of current electricity markets. Essentially, our proposal involves a new tariff structure, the *prediction-of-use* (POU) tariff, which asks customers to predict their baseline consumption and charges them based on both their actual consumption and the deviation from their prediction. As we show in Section 3, this tariff reflects explicitly the predictability cost by separating between a potentially lower price for the baseline consumption from a higher penalty for the deviation. In this way, our pricing structure incentivises more reliable consumption patterns and increase rate fairness, as customer would no longer subsidise other customers for their predictability cost. Note that tariffs with this structure are already a reality in a number of markets. For example Braithwait et al. [16] describe a real-time tariff offered by several energy suppliers in the US, where customers commit to a self-selected baseline load, and are charged a standard rate for their consumption baseline, a penalty rate for usage in excess of their baseline, and receive return credit for underconsumed units. However, so far, most of these offerings are addressed to large-scale industrial consumers, as individual domestic consumers are typically

Appears in: *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), Lomuscio, Scerri, Bazzan, Huhns (eds.), May, 5–9, 2014, Paris, France.*

Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

too small to have a significant influence on costs.

A potential downside of this tariff is that it transfers part of the prediction risk faced by the supplier to the consumers. Now, one way that customers can collectively act to reduce the risk associated to their consumption imbalance is by joining the electricity tariff as a single virtual consumer (i.e. aggregating their demand). In the Internet era, such group buying has been shown to be both popular and effective (e.g. due to the success of e-sites such as *Groupon* or *LivingSocial*). In the context of energy, real-world initiatives such as the *BigSwitch*¹ or *thePeoplesPower*² achieved significant discounts by bringing a large number of consumers together and negotiating a better deal on their behalf with suppliers.

Against this background, we study the underlying properties for the type of tariff that we propose, from a coalitional game theory perspective. By so doing, we ensure that, in expectation, any individual consumer or subset of consumers will pay less by being part of the group initiative. Second, we identify cost allocation schemes that can fairly allocate the expected bill among customers taking part in an electricity group buying initiative.

From an application perspective, another important issue with group buying is *inertia*: consumers are reluctant to join a collective switching initiative which only a few people have joined. This may be due to strategic considerations or to the fact that users are unconvinced the initiative will be successful until a “critical mass” is reached [5]. To deal with this, we also propose new payment allocation schemes that incentivize customers to commit early.

In summary, our work makes the following contributions:

- We propose a new kind of tariff that introduces a *prediction-of-use* component to price electricity and study the theoretical properties of such a tariff. We show that no agent has an incentive to consume electricity in excess of its actual requirements, regardless of the baseline it reported. Moreover, using a probabilistic model of the uncertainty over future consumption, we derive closed form expressions for the optimal baseline an agent should report, as well as its expected payment using this baseline under the posted tariff.
- We study the game that our tariff induces among a set of consumers. We show that the game is concave,³ hence the bill is always reduced by grouping and it can be distributed fairly using the well-known Shapley value solution concept.
- Considering a realistic model where consumers may join a group online, over a period of time, we study the problem of designing payment distributed schemes that incentivise them not to delay their arrival (i.e. joining at the earliest time a consumer is available guarantees her the lowest possible payment). We develop several such allocation schemes, based both on strict arrival order and on arrival windows.
- We study and validate experimentally all the mechanisms proposed using a dataset of 3000 domestic consumers in the UK. We quantify the incentives for joining groups and allocation payments, for consumers with different prediction accuracies.

This paper is organised as follows. Section 2 gives some preliminaries and related work. Section 3 introduces our POU tariffs and their theoretical properties. The associated cooperative cost game is studied in Section 4, while Section 5 proposes our online incentive compatible allocation schemes. Section 6 evaluates the proposed tariffs using real data and Section 7 concludes.

¹<https://www.whichbigswitch.co.uk/>

²<http://www.thepeoplespower.co.uk/>

³Concavity in cost games is equivalent to the concept of convexity in utility maximizing games.

2. RELATED WORK & PRELIMINARIES

Recently, group buying has been an active area of research in the multi-agent and AI communities [7, 8, 9]. However, this strand of work studies how to incentivise buyers in online markets through volume-based discounts, without pursuing greater reliability or behavioural change.

In the energy domain, several works have proposed the formation of cooperatives or coalitions among renewable energy producers, with uncertain supply [2, 12]. Baeyens et al. [2] use a coalitional game theory approach to propose a fair mechanism for dividing revenue in wind energy aggregation. However, the game they study is not always convex, hence Shapley value payment allocation are not core-stable. By contrast, we assume that consumer prediction errors follow independent normal distributions (i.e. a common assumption in statistics theory for prediction errors [4]), and show in this case the prediction-of-use game we consider is always concave.

Several works have considered incentivising the formation of coalitions among electricity consumers, thus have very similar aims to this work. For instance, Rose et al. [13] propose a mechanism that insures truthful revelation of demands and prevents overconsumption (a property we also consider in this work). Kota et al. [6] and Akasiadis & Chalkiadakis [1] propose the formation of coalitions between consumers in order to improve reliability and shift peak-time electricity loads. However, these works do not study the concavity or core stability of coalitions in energy purchasing.

Finally, the work on online group formation, where consumers join over time, is inspired by recent advances in online mechanism design (see Parkes [11] for an overview). While the basic aim is the same (that of assuring participants will have no incentive to delay their engagement), work in that area typically focuses on allocation of resources, while here we study incentives for group formation.

2.1 Cooperative game theory

In this section we provide an overview of basic concepts of cooperative game theory [3] used later in the paper. Let $N = \{1, \dots, n\}$ be a set of agents. A subset $S \subseteq N$ is termed a coalition. Then a *cost game* is a tuple $\langle N, c \rangle$ where $c : 2^N \rightarrow \mathbb{R}$ is a *characteristic function* that assigns to every coalition a real value representing the cost that the coalition incurs by itself ($c(\emptyset) = 0$). All the properties and solution concepts are presented in terms of cost games (i.e. minimizing cost) although connections with the corresponding concepts in utility games (i.e. maximizing) are outlined.

Given a cost game, there are many ways to divide the cost of the game (i.e. assuming transferable cost) among its members. A vector $\varphi = \{\varphi_1, \dots, \varphi_n\}$ that assigns some cost to each agent $i \in N$ is called an *allocation*. We denote $\sum_{i \in S} \varphi_i$ as $\varphi(S)$. An allocation is an *imputation* for N , if it is *efficient* ($\varphi(N) = c(N)$) and *individually rational* ($\varphi(\{i\}) \leq c(\{i\}) \forall i \in N$). The cost allocation issue has been extensively studied in cooperative game theory under considerations of fairness and stability. The most prominent solution aiming at stability is the *anticore* (corresponding to the core in utility games).

DEFINITION 1. *The anticore of a cost game $\langle N, c \rangle$ is composed of all imputations of N such that $\varphi(S) \leq c(S) \forall S \subset N$.*

Note that φ is in the anticore of a cost game iff no coalition $S \subset N$ can improve upon φ (i.e. has no cross-subsidies).

Analogous to the anticore, the most prominent solution aiming at fairness is the *Shapley value*. The *Shapley value* is based on the intuition that the payment that each agent receives should be proportional to its contribution. To formally define the *Shapley value*, we need first to define the notion of marginal contribution.

Let Π_N denote the set of all orderings of agents N . Given an ordering $\pi \in \Pi_N$, let $S_{\leq i}^{\pi}$ be the set of predecessors of i in π .

DEFINITION 2. The marginal contribution of an agent i w.r.t. an ordering π is defined as $\Delta^\pi(i) = c(S_{<i}^\pi \cup \{i\}) - c(S_{<i}^\pi)$.

Then, the *Shapley* value of an agent i , is simply the average of all possible marginal contributions, that is w.r.t all orderings of N . Formally, the *Shapley* allocation scheme assigns:

$$\varphi(i) = \frac{1}{n!} \sum_{\pi \in \Pi_N} \Delta^\pi(i) \quad \forall i \in N \quad (1)$$

One interesting property that a cost game might satisfy is concavity (corresponding to convexity in utility games).

DEFINITION 3. A cost game $\langle N, c \rangle$ is concave if $c(T \cup \{i\}) - c(T) \leq c(S \cup \{i\}) - c(S)$ for all $i \in N$ and all $S \subset T \subseteq N \setminus \{i\}$.

Hence, in concave games the marginal contribution of an agent to any coalition is greater than its marginal contribution to a larger coalition. Thus, concavity implies that it is always beneficial (i.e. reduces cost) to group agents together, meaning the grand coalition is efficient. Moreover, concavity means that any payment scheme that is a linear combination of marginal contributions (including the Shapley value) are in the core.

3. PREDICTION-OF-USE TARIFFS

In this section we define the class of prediction-of-use tariffs. First, in Section 3.1 we formalize the tariff scheme, discuss its economically viability and exemplify its application to several consumption patterns. Next, in Section 3.2 we study some of its theoretical properties.

3.1 The tariff scheme

In a POU tariff, each customer is asked in advance to predict an expected *baseline* for her consumption, denoted henceforth by b . Intuitively described, the tariff works by asking each customer to pay a rate p (with $p > 0$) for her consumption, while deviations are charged proportionally as follows. Units actually consumed, but not predicted in the baseline are charged an additional marginal rate of \bar{p} (with $\bar{p} \geq 0$). For the units whose consumption is predicted, but are not actually consumed, the tariff also charges a penalty rate \underline{p} (with $0 \leq \underline{p} < p$), typically much smaller than p .

Formally, given a prediction-of-use tariff described by a tuple $\langle p, \bar{p}, \underline{p} \rangle$ the payment for a consumer with a predicted baseline b and an actual (realised) consumption x is determined as:

$$\psi(x, b) = \begin{cases} p \cdot x + \bar{p} \cdot (x - b) & \text{if } b \leq x \\ p \cdot x + \underline{p} \cdot (b - x) & \text{otherwise} \end{cases} \quad (2)$$

A first observation about this tariff (c.f. Eq. 2) is that for $\bar{p} = 0$ and $\underline{p} = 0$, the model corresponds to the traditional flat tariff, in which customers pay a fixed price (p) per kW consumed, regardless of their prediction accuracy.

As discussed in [16], tariffs with this structure are economically viable from a supplier's perspective. In more detail, by setting the prices accordingly, the supplier can receive essentially the same revenue from each customer under such a tariff that it would have received had the customer remained on the standard one. Moreover, changes in consumptions patterns induced by these tariffs can potentially provide benefits not only to customers but also to the suppliers, thus incentivising the latter to offer them. Next, we illustrate how a supplier can make a POU tariff economically viable by means of three examples. The parameters $\langle p, \bar{p}, \underline{p} \rangle$ (listed in Table 1) are set to roughly match long-term averages from the UK balancing market⁴.

⁴Historical balancing prices for every Settlement Period in a particular day are available at <http://www.elexon.co.uk/reference/credit-pricing/imbalance-pricing>

	p	\underline{p}	\bar{p}
P^+ redictive	0.05	0.03	0.05
Predictive	0.06	0.02	0.02
Flat	0.08	0	0

Table 1: Examples of POU tariffs (in £/kWh).

	Demand (in kW)	Baseline (in kW)	Payment (in £)		
			P^+ redictive	Predictive	Flat
Bob	300	500	24	22	24
Annie	300	100	35	34	24
John	300	300	15	18	24

Table 2: Example of monthly consumptions, predicted baselines, and payments under tariffs of Table 1.

EXAMPLE 1 (P^+ REDICTIVE). This tariff gives the best possible price for the predicted baseline but severely penalizes any imbalance with respect to it. To make this tariff economically viable, we can imagine the supplier contracting the baseline predicted by the consumer in the forward market and charging the consumer's imbalance at the prices in the balancing market. Hence p is set as the price for kWh obtained the forward market (£0.05); \bar{p} is set as the difference between the expected price to buy in the imbalance market and the baseline price (£0.1 - £0.05 = £0.05); and \underline{p} is set as the difference between the baseline price and the expected price to sell in the imbalance market (£0.05 - £0.02 = £0.03).

EXAMPLE 2 (PREDICTIVE). This tariff reduces the penalty for imbalances at the cost of increasing the baseline price. This tariff is economically viable for a supplier when contracting not only the baseline at the forward market but also some extra quantity to account for potential imbalances. Although the baseline price increases slightly (from £0.05 to £0.06) it also reduces the imbalance prices, from £0.03 to £0.02 for underconsumption and from £0.05 to £0.02 for overconsumption.

EXAMPLE 3 (FLAT). A flat tariff with $\bar{p} = £0.08$ per kWh, roughly similar to wholesale costs in the prevalent tariffs in the UK.

Next, we illustrate the computation of actual payments w.r.t. these tariffs. Consider three customers, Bob, Annie and John, that contracted a POU tariff with a predicted monthly baseline of 500kW, 200kW and 300kW respectively. Although at the time of demand, all of them realise a consumption of 300kW, their payments may vary because of the different predicted baselines. Table 2 shows the respective payments of each consumer, in our three example tariffs. Bob is penalized for his 200 units of overconsumption, Annie for her 200 units of underconsumption, whereas John pays his consumption at the baseline price with no penalties. As we observe in this example, POU tariffs reward more predictable consumers: John always pays the cheapest rate in all tariffs. We also observe that the benefit of John is greater in the P^+ redictive than not in the Predictive or the Flat tariff (in which indeed there is no reward for being a good predictor).

3.2 Theoretical properties

This section studies the monotonicity of POU tariffs, and derives closed form expressions for the optimal baseline an agent should report, as well as for its expected payment.

3.2.1 Monotonicity w.r.t. consumption

A first important property to be guaranteed in this setting is monotonicity w.r.t. to the realized consumption. Specifically, we need to guarantee that, irrespective of the predicted baseline b , a higher actual consumption will always result in a higher payment. This property is important because otherwise an agent could have the incentive to artificially inflate its consumption just to meet its prediction (e.g. a domestic consumer could let electric heating run, even if there is no need for it).

PROPERTY 1. A tariff is called strictly monotonic w.r.t. realisations if $\forall b, \forall x_1, x_2$ with $x_1 < x_2$, then $\psi(x_1, b) < \psi(x_2, b)$.

In words, regardless of b , an agent will not pay more for consuming x_1 than x_2 whenever $x_1 < x_2$.

THEOREM 1. The tariff in Equation 2 is strictly monotonic.

PROOF. The proof needs to consider four cases: 1. $x_1, x_2 < b$, 2. $x_1, x_2 \geq b$, 3. $(x_1 < b) \wedge (x_2 \geq b)$ and 4. $(x_1 \leq b) \wedge (x_2 > b)$. For cases 1 and 2, the inequality follows immediately, by just applying Eq. 2. For cases 3 and 4, assume by contradiction that $\psi(x_1, b) > \psi(x_2, b)$. For both cases, applying Eq. 2 yields: $px_1 + \underline{p}(b - x_1) > px_2 + \bar{p}(x_2 - b)$, resulting in $p(x_1 - x_2) + \underline{p}(b - x_1) - \bar{p}(x_2 - b) > 0$. By replacing in the first term x_2 by $x_1 + (b - x_1) + (x_2 - b)$ and regrouping, we obtain: $(\underline{p} - p)(b - x_1) - (\underline{p} + \bar{p})(x_2 - b) > 0$. We know from construction of the tariff that $0 \leq \underline{p} < p, 0 \leq \bar{p}$ and hence $(\underline{p} - p)$ and $-(\underline{p} + \bar{p})$ are strictly negative. Moreover, for case 3, $(b - x_1) \geq 0$ and $(x_2 - b) > 0$ and for case 4, $(b - x_1) > 0$ and $(x_2 - b) \geq 0$. Thus, all terms on the left side must be less or equal to 0 and at least one of them must be strictly negative (the second one in case 3 and the first one in case 4), leading to a contradiction. \square

3.2.2 Expected payment

We assume that the predicted consumption of a customer (related to time scale at which the tariff requires to predict) is modelled as a random variable x . Now, given the tariff structure presented in Eq. 2, the expected payment of a customer for a predicted baseline b is:

$$\begin{aligned} \mathbb{E}[\psi(x, b)] &= \underline{p} \int_0^b (b - x) f(x) dx + \bar{p} \int_b^\infty (x - b) f(x) dx + p \int_0^\infty x f(x) dx \\ &= \underline{p} \int_0^\infty (b - x) f(x) dx + (\bar{p} + \underline{p}) \int_b^\infty (x - b) f(x) dx + p \int_0^\infty x f(x) dx \\ &= \underline{p} b \int_0^\infty f(x) dx + (\bar{p} + \underline{p}) \int_b^\infty (x - b) f(x) dx + (p - \underline{p}) \int_0^\infty x f(x) dx \\ &= \underline{p} b + (\bar{p} + \underline{p}) \int_b^\infty (x - b) f(x) dx + (p - \underline{p}) \mathbb{E}[x] \end{aligned} \quad (3)$$

Here f denotes the probability density function of the expected consumption, F is the cdf of this function, while $\mathbb{E}[x]$ is its expectation. Note the third equality holds by $\int_0^b (b - x) f(x) dx = \int_0^\infty (b - x) f(x) dx - \int_b^\infty (b - x) f(x) dx$. The fourth equality holds by $\int_0^\infty f(x) dx = 1$.

3.2.3 Optimal baseline

Given the expected payment presented in Eq. 3, the next question is which quantity an agent should report as a baseline. In other words, we are interested in determining the optimal baseline b^* that minimises the expected payment of a customer, namely:

$$b^* = \arg \min_b \mathbb{E}[\psi(x, b)] \quad (4)$$

Thus, in order to find b^* , we take the derivative of the expected payment function defined in Eq. 3 and set it to zero to find:

$$\frac{d\mathbb{E}[\psi(x, b)]}{db} = \underline{p} - (\bar{p} + \underline{p})(1 - F(b)) = 0 \quad (5)$$

$$\Rightarrow F(b) = 1 - \frac{\underline{p}}{\bar{p} + \underline{p}} = \frac{\bar{p}}{\bar{p} + \underline{p}} \quad (6)$$

Here we call the ratio $\frac{\bar{p}}{\bar{p} + \underline{p}}$ the *optimal ratio* and we refer to it as r^* . Note that since the cdf F is a bijective (thus reversible) function, b^* is uniquely determined as:

$$b^* = F^{-1}(r^*) \quad (7)$$

Similarly testing the second derivative proves that b^* is a global maximum (proof is omitted for lack of space).

Next, the expected payment of the agent given in Eq. 3 when restricted to the optimal baseline $b = b^*$ can be simplified as:

$$\begin{aligned} \mathbb{E}[\psi^*(x)] &= \underline{p} b^* + (\bar{p} + \underline{p}) \int_{b^*}^\infty (x - b^*) f(x) dx + (p - \underline{p}) \mathbb{E}[x] \\ &= (\bar{p} + \underline{p}) \int_{b^*}^\infty x f(x) dx + (p - \underline{p}) \mathbb{E}[x] \\ &= (\bar{p} + p) \int_0^\infty x f(x) dx + (-\underline{p} - \bar{p}) \int_0^{b^*} x f(x) dx \\ &= (\bar{p} + p) \mathbb{E}[x] + (-\underline{p} - \bar{p})(b^* \cdot F(b^*) - \int_0^{b^*} F(x) dx) \\ &= (\bar{p} + p) \mathbb{E}[x] + (-\underline{p} - \bar{p}) \int_0^{F(b^*)} F^{-1}(y) dy \end{aligned} \quad (8)$$

Here, the second equality holds by $\int_{b^*}^\infty (x - b^*) f(x) dx = \int_{b^*}^\infty x f(x) dx - b^* \int_{b^*}^\infty f(x) dx$ and $\int_{b^*}^\infty f(x) dx = 1 - F(b^*) = 1 - \frac{\bar{p}}{\bar{p} + \underline{p}}$. The third equality holds by: $\int_{b^*}^\infty x f(x) dx = \int_0^\infty x f(x) dx - \int_0^{b^*} x f(x) dx$. The fourth equality holds by means of partial integration: $\int_0^{b^*} x f(x) dx = x F(x) \Big|_0^{b^*} - \int_0^{b^*} F(x) dx = b^* \cdot F(b^*) - 0 \cdot F(0) - \int_0^{b^*} F(x) dx$. The last equality follows from interchanging the axes of integration: $\int_0^{b^*} F(x) dx = b^* \cdot F(b^*) - 0 \cdot F(0) - \int_0^{F(b^*)} F^{-1}(y) dy$.

3.2.4 Expected payment under normal distributions

The expected payment with ex-post optimal baseline of Eq. 8 holds for a random variable that follows any class of probability distribution. Here, we are interested in deriving the expected payment in the particular case when the predicted consumption x follows a normal demand distribution $N(\mu, \sigma)$. For this case, we observe that the expected payment of Eq. 8 simplifies as:

$$\begin{aligned} \mathbb{E}[\psi^*(N(\mu, \sigma))] &= (\bar{p} + p) \mu + (-\underline{p} - \bar{p}) \int_0^{F(b^*)} F^{-1}(y) dy \\ &= (\bar{p} + p + (-\underline{p} - \bar{p}) r^*) \mu + (-\underline{p} - \bar{p}) \sigma \int_0^{r^*} \Phi^{-1}(y) dy \\ &= \underbrace{p \mu}_{\text{Consumption term}} + \underbrace{(-\underline{p} - \bar{p}) \sigma \int_0^{r^*} \Phi^{-1}(y) dy}_{\text{Penalty term}} \end{aligned} \quad (9)$$

The first equality follows from the expected payment of Eq. 8 after replacing the expected value of x by the mean of the normal distribution ($\mathbb{E}[x] = \mu$). We obtain the second equality by observing that in normal distributions $F^{-1}(y) = \mu + \sigma \Phi^{-1}(y)$ where $\Phi^{-1}(\cdot)$ is the inverse of the standard normal cdf. Finally, the third equality can be simplified and leads to a very intuitive equation composed of two terms: one depending on the expected consumption value (what we refer to as the *consumption term*) and one depending on its variance (what we refer to as *penalty term*). We refer to the first as the *consumption term* because it rates the expected consumption value at the best price you can achieve under a POU tariff, the baseline price. Then, we refer to the second as a *penalty term* because it increases the minimum payment per consumption depending on the variance of the customer.

We next prove that this penalty term is always greater or equal to 0. First, notice that by definition $\sigma \geq 0$. Thus, for the product in the penalty term of Eq. 9 to be ≥ 0 , it remains to show that:

$$(-\underline{p} - \bar{p}) \int_0^{r^*} \Phi^{-1}(y) dy \geq 0 \text{ for all } r^* \in [0, 1] \quad (10)$$

To prove that, notice that since $\bar{p} \geq 0$ and $p > 0$ then $(-\underline{p} - \bar{p}) \leq 0$. Hence in order for the product in Eq. 10 to be ≥ 0 , it remains to show that $\Phi^{-1}(y)dy \geq 0$ for all $r^* \in [0, 1]$. We will distinguish between two cases, $r^* \leq \frac{1}{2}$ and $r^* > \frac{1}{2}$. First, if $r^* \leq \frac{1}{2}$ we find directly, using that $\Phi^{-1}(r^*) \leq 0$ for all $0 \leq r^* \leq \frac{1}{2}$ that (10) holds. Secondly, suppose $r^* > \frac{1}{2}$. Then, $\int_0^{r^*} \Phi^{-1}(r^*) = \int_0^{\frac{1}{2}} \Phi^{-1}(r^*) + \int_{\frac{1}{2}}^{r^*} \Phi^{-1}(r^*)$ and since $\Phi^{-1}(r^*) \geq 0$ for all $\frac{1}{2} \leq r^* \leq 1$ and $\int_0^{\frac{1}{2}} \Phi^{-1}(r^*) - \int_{\frac{1}{2}}^1 \Phi^{-1}(r^*) \leq 0$ because of the symmetry of the standard normal distribution, hence inequality (10) holds and the penalty term in Eq. 9 is always greater or equal to 0.

4. THE PREDICTION-OF-USE GAME

The previous section studied the POU tariff from the perspective of an individual consumer. When applied to a set of consumers (aka agents), this tariff induces a cooperative game, as consumers can form coalitions to reduce their expected payments.

Consider a set of N customers that joined the group-tariff scheme. Given a subset of customers $S \subseteq N$, let $x_S = \sum_{i \in S} x_i$ be its aggregate error in demand prediction. We define the corresponding prediction-of-use (POU) game $G = \langle N, c \rangle$, where $c : 2^N \rightarrow \mathbb{R}$ is:

$$c(S) = \mathbb{E}^*[\psi(x_S)] \text{ for all } S \subseteq N. \quad (11)$$

where $\mathbb{E}^*[\psi(x_S)]$ is the expected joint payment for deviation demand prediction x_S defined as in Eq. 9.

Now, an important question raised by this grouping is whether it is beneficial for agents to form a coalition (i.e. to share the risk associated to their consumption imbalance) and whether the coalition that forms is core-stable (i.e. there is no individual agent or subset of agents that have an incentive to break away to form their own group). In this section, we answer these questions using the tools from cooperative game theory [3].

The following theorem states that POU games, assuming that deviations in consumption predictions follow independent normal distributions, are concave. Our proof uses a similar technique to the ones used in [10] to prove the convexity of newsvendor games.

THEOREM 2. *Assuming independent normal distributions for consumers prediction errors, $x_i \sim N(\mu_i, \sigma_i)$ for all $i \in N$, the POU game (N, c) is concave.*

PROOF. We need to prove that the following function (associated with the concavity condition given in Definition 3) is always negative, i.e.

$$con_{i,S,T} = c(T \cup \{i\}) - c(T) - (c(S \cup \{i\}) - c(S)) \leq 0 \quad (12)$$

for all $i \in N$ and all $S \subset T \subseteq N \setminus \{i\}$. Simplifying using Equation 9, we obtain:

$$con_{i,S,T} = p (\mu_{T \cup \{i\}} - \mu_T - \mu_{S \cup \{i\}} + \mu_S) + (-\underline{p} - \bar{p}) (\sigma_{T \cup \{i\}} - \sigma_T - \sigma_{S \cup \{i\}} + \sigma_S) \int_0^{r^*} \Phi^{-1}(y) dy$$

First, the difference in consumption term cancels out because $\mu_{T \cup \{i\}} - \mu_T = \mu_i = \mu_{S \cup \{i\}} - \mu_S$ (the mean of a sum of two normal distributions is the sum of the means of their individual distributions). For the second term, we know $(-\underline{p} - \bar{p}) \int_0^{r^*} \Phi^{-1}(y) dy \geq$

0 from Eq. 10. Finally, the following has been shown in [10] to hold for $\forall S \subset T \subseteq N \setminus \{i\}$: $\sigma_{T \cup \{i\}} - \sigma_T \leq \sigma_{S \cup \{i\}} - \sigma_S$. Thus it follows that $\sigma_{T \cup \{i\}} - \sigma_T - (\sigma_{S \cup \{i\}} - \sigma_S) \leq 0$ and therefore Eq. 12 holds. \square

Having established the POU game is concave, the following properties follow immediately.

COROLLARY 1. *The POU game with independent normal consumption prediction deviations in demand prediction is subadditive and totally balanced.*

COROLLARY 2. *The Shapley value payments are in the anti-core of the POU game with independent normal consumption prediction deviations.*

These properties follow immediately from the concavity result in Theorem 2 and standard cooperative game theory (c.f. [3]). Moreover if the game is convex (respectively concave, for cost games), for any ordering of the players $\pi : N \rightarrow N$ a core-stable allocation can be constructed in polynomial time by assigning to each agent its marginal contribution with respect to the ordering π ($\forall i \in N : \varphi(i) = \Delta^\pi(i)$). For a convex/concave game, all such marginal payments are in the core. The Shapley value, which computes the average marginal payments across all possible orderings can be seen as the “fairest” of such allocations, but is not the only core stable one. Computing the Shapley value for each agent can be computationally expensive and, as we discuss in the next section, in online settings we may like to focus on particular orders, that would guarantee the additional “no delay” property.

5. ONLINE SETTINGS

The previous section studied the properties of the POU game from a static perspective. Essentially, the property that a payment allocation scheme (e.g. *Shapley*) is in the *anticore* means that no subset of agents have an incentive to break away, once a coalition is formed. However, in real group buying situations, customers join dynamically, over a period of time until a critical mass is reached. In mechanism design terms, the group formation problem is an *online* problem [11]. An important challenge identified in the literature on group formation [5] is that strategic customers may delay joining the group, to see whether enough customers have already joined, or perhaps whether they can receive a better alternative in the meantime. Hence, in payment allocation schemes, it is important to prevent incentives for this behaviour. In this section, we propose several *computationally tractable* allocation schemes that, in addition to be core-stable, guarantee “no delay”.

Formally, we consider a set N agents (with $n = |N|$), arriving at times $1 \dots n$. Note that, w.l.o.g. we are only interested in the *order* of arrival of these agents, not in actual times. Here, a_t denotes the agent arriving at position t in the sequence of $1 \dots n$, assuming no two agents arrive exactly at the same time⁵. The order of arrival of agents in N is denoted by \mathcal{O} (where $\mathcal{O} : N \rightarrow N$). Each agent is expected a payment $\varphi^\mathcal{O}(a_t)$ if arriving at position t in \mathcal{O} . Moreover, let $\mathcal{O}_{<t}$ and $\mathcal{O}_{>t}$ refer to the suborderings of elements in \mathcal{O} before and after position t , and $S_{<a_t}^\mathcal{O}$ and $S_{>a_t}^\mathcal{O}$ denote respectively the subsets of predecessors and successors of agent a_t . Moreover, we will use \mathcal{O}^{-a} to denote ordering \mathcal{O} with agent a removed.

PROPERTY 2 (NO DELAY). *An agent a , whose earliest availability of arrival is t has no incentive to misreport a $t' > t$, because $\varphi^\mathcal{O}(a_t) \leq \varphi^{\mathcal{O}'}(a_{t'})$, where $\mathcal{O}' = \mathcal{O}_{<t'} \cup \{a\} \cup \mathcal{O}_{>t'}$*

Given the above notation, we are now ready to define some core-stable payment allocation schemes that guarantee no-delay.

⁵Ties can be broken at random, using a fair coin.

5.1 Allocation using a strict arrival order

A natural payment allocation scheme in this online domain is the *marginal cost* allocation with respect to the ordering \mathcal{O} , $\forall t : \varphi^{\mathcal{O}}(a_t) = \Delta^{\mathcal{O}}(a_t)$ (c.f. Definition 2). Such a allocation scheme is an imputation of N and core-stable, but, crucially, it does not satisfy the no delay property. This is because, due to the concavity of the offline game (c.f. Theorem 2), joining a larger group always has decreasing marginal costs than joining a smaller group (i.e. $\varphi^{\mathcal{O}}(a_{t+1}) \leq \varphi^{\mathcal{O}}(a_t)$). To rectify this, consider the *reverse marginal* cost of the agent (i.e. the marginal costs in the reverse order to \mathcal{O}):

$$\mathcal{R}\Delta^{\mathcal{O}}(a_t) = c(S_{>a_t} \cup \{a_t\}) - c(S_{>a_t}) \quad (13)$$

LEMMA 1. *The Reverse Marginal (RM) allocation scheme, which assigns $\varphi^{\mathcal{O}}(a_t) = \mathcal{R}\Delta^{\mathcal{O}}(a_t)$ is in the anticore of the POU game and satisfies the “no delay” property (Property 2).*

PROOF. The first proof is immediate since the marginal payments with respect to any ordering of N (including the reverse order of \mathcal{O}) are in the anticore of a concave game. Also due to concavity $\forall t, t' \in N$ such that $t < t'$, $\mathcal{R}\Delta^{\mathcal{O}}(a_t) \leq \mathcal{R}\Delta^{\mathcal{O}'}(a_t)$ where $\mathcal{O}' = \mathcal{O}_{<t'}^{-a} \cup \{a\} \cup \mathcal{O}_{>t'}^{-a}$ holds because $\mathcal{O}_{<t}^{-a}$ contains a subset of the agents in $\mathcal{O}_{<t'}^{-a}$. Hence Property 2 holds. ⁶ \square

EXAMPLE 4. Consider $n = 4$ agents, arriving in the order $\mathcal{O} = \{ABCD\}$ (for simplicity, say these 4 agents are equal in their expected consumptions $N(\mu, \sigma)$). Using the RM mechanism, the marginal payment of the first agent to arrive will be $\varphi^{\mathcal{O}}(A) = c(\{ABCD\}) - c(\{BCD\})$, which is the lowest of the 4 agents. The marginal payment of the last agent to arrive is $\varphi^{\mathcal{O}}(D) = c(\{D\}) = \psi(x_D, b_D)$, which is the highest of the 4 agents, and the same as if it participates in the market alone. Although the last agent is indifferent to joining, this is still his best strategy in expectation if there is some non-zero probability of later arrivals.

5.2 Allocation using arrival windows

As shown in the previous section, the RM allocation ensures the non-delay property by heavily penalizing customers based on their order of arrival. However, this penalization may be inappropriate as customers may see as unfair that payments vary significantly among near arrival positions. In contrast, in the Shapley value each agent receives a payment proportional to its contribution but it does not incentivize an early commitment since all agents are treated the same, regardless of their arrival time. Given this, in this section, we propose a new payment allocation that generalises both the RM and the Shapley value allocation.

In general, the way many group buying schemes work in practice is by promising a discount to the first w customers to register, and perhaps another discount level to agents in later arrival windows. For example, in the electricity domain, the first 100 customers to join a scheme are promised a certain price. When the next 100 customers join, these can also get a discount, but the designer may want to give a better deal to agents in earlier arrival windows. Inspired by this model, we can generalise both the reverse marginal (RM) and Shapley value as follows.

First, our mechanism divides the n agents into several arrival windows of size w . There are in total $\lceil \frac{n}{w} \rceil$ arrival windows, where the last window may contain less than w elements. Moreover, we

⁶While the proof is skipped here due to lack of space, it can also be shown that, out of the possible $n!$ orders of the n agents, the marginal payments in the RM mechanism are the *only* ones that guarantee no delay for any valuation setting. This is especially true for settings in which $\sigma_i > 0, \forall i \in N$, because in such settings marginal payments will all be different.

can factor each $t \in \{1 \dots n\}$ as $t = (k-1)w + l$, where $k = \lceil \frac{t}{w} \rceil$ is the window of arrival and $l = t \bmod w$ is the position of the agent within the window. To simplify the notation, we will refer to a_l^k as the agent that arrived at position l within window k . Thus, each window of arrival W_k , where $k = 1 \dots \lceil \frac{n}{w} \rceil$, will contain the agents: $W_k = \{a_1^k, \dots, a_w^k\}$.

The main intuition of our method is that, *between* the arrival windows W_t the reverse marginal mechanism will apply, rewarding earlier groups with a lesser marginal payment. But the agents *inside* each arrival group will be treated the same, and between them Shapley allocation scheme will apply.

Formally, let Π_{W_k} be the set of possible orderings of the agents inside the window of arrival k , and let $\omega \in \Pi_{W_k}$ be one such ordering. Given an order of arrival \mathcal{O} , for each ordering $\omega \in \Pi_{W_k}$, we construct an ordering $\pi(\omega) : N \rightarrow N$ as:

$$\pi(\omega) = \langle \mathcal{O}_{<(k-1) \cdot w+1}, \omega, \mathcal{O}_{>k \cdot w} \rangle$$

Hence, $\pi(\omega)$ merges ordering ω for those elements in W_k , and \mathcal{O} outside it. Then, the so-called Window-of-Arrival (WoA) allocation scheme assigns to each agent $i \in N$ the following payment:

$$\varphi(i) = \frac{1}{w!} \sum_{\omega \in \Pi_{W_k}} \mathcal{R}\Delta_{\pi(\omega)}(i) \quad (14)$$

where k is the window of arrival of i .

Note the similarity to the definition of the Shapley value (see Equation 1), except that we only average over all orderings of W_k , for the rest their arrival order is used.

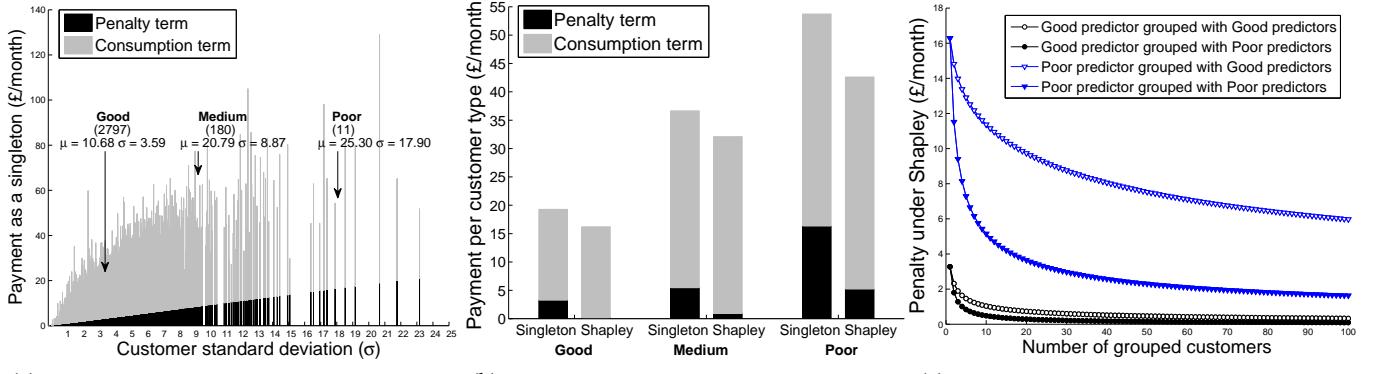
EXAMPLE 5. Consider an example with $n = 8$ agents, arriving in order $\mathcal{O} = \{ABCDEFGH\}$ and let the size of the arrival window be $w = 3$. Now, consider the payment that agent D (which is here a_1^2), belonging to arrival window $W_2 = \{D, E, F\}$. Then ω iterates over all orderings $\Pi_{\{D, E, F\}}$. The orders that need to be considered are $\pi = \{\langle ABCDEFGH \rangle, \langle ABCDFEGH \rangle, \langle ABCEDFGH \rangle, \langle ABCEF DGH \rangle, \langle ABCFDEGH \rangle, \langle ABCFEDGH \rangle\}$. The payment of agent D will be the marginal payments over the orders in π , each considered in reverse order.

Two mechanisms appear as special cases. For $w = 1$, the mechanism is the same as the RM mechanism, thus the strict order of arrival matters. For $w = n$, then $W_1 = N$ and the mechanism reduces to the Shapley scheme.

THEOREM 3. *The Window-of-Arrival (WoA) allocation scheme is in the anticore of the POU game.*

PROOF SKETCH. Recall that the POU game among a set of N agents is concave. Therefore, the anticore of this game is a closed, convex set. Mathematically, the anticore is a polytope with $n!$ vertices (or extrema). In each extrema, the payments are equivalent to the marginal costs in one of the $n!$ (where $n = |N|$) orderings of the set of N agents [14]. Given that the anticore is convex, points which are linear combinations of payments in these extrema are also in the anticore. The Shapley value, which averages over all $n!$ possible orderings is just the most well known such point, and the center of gravity of the anticore polytope. The payments described by Eq.14 can be shown to be averages over the marginal payments in $(w!)^{\lceil \frac{n}{w} \rceil}$ orders from the extrema, and thus are also in the anticore. \square

In terms of computational complexity, $O((w!)^{\lceil \frac{n}{w} \rceil}) \leq O(n!)$. But in practice, as shown in Eq. 14, because of redundancy only $w!$ distinct orderings need to be considered when computing the WoA payment of an agent with window size w . Thus, the computation of the WoA payment of an agent is tractable for small w .



(a) Singleton payments for customers ordered by predictability (increasing σ).

(b) Singleton payments vs Shapley payments for different types of customer predictability.

(c) Penalty under Shapley of good/poor types when grouped with good/poor types.

Figure 1: Expected payments under the P^+ redictive tariff for (a) individuals as singletons; (b) customers types as singletons and when grouping under Shapley; and c) good/poor customers types when joining a group of all good/poor types. Payments are split into the consumption term ($\mu \cdot p$) and the remaining penalty term.

6. EXPERIMENTAL EVALUATION

We analyse the incentives that a real set of domestic customers will have to collectively join our group buying scheme. The data used in our analysis comes from a large dataset of around 3000 households in the UK. For each household, the dataset included the electricity consumption of each consumer for every half an hour during a three-month period. For each consumer we take the sample mean over her consumption realizations as a point estimate for μ and the standard deviation as an unbiased estimator for σ . The evaluation considers the customers individual payments when joining altogether a P^+ redictive tariff (as defined in Table 1). The P^+ redictive tariff has been chosen to be the one with higher penalty for imbalances and hence, the one that provides higher benefits for grouping. In particular, the total payment reduction for grouping this set of households under the P^+ redictive tariff is £11113 per month (i.e. equivalent to an average saving of £3.7 per month per household).

In the next section, we evaluate how these group savings are allocated among individuals under a core-stable Shapley scheme. Afterwards, in Section 6.2, we do the same considering an online enrollment with savings distributed using arrival windows.

6.1 Payment allocation under Shapley

We analyse the incentives of customers to join our tariff scheme as a group with respect to joining alone, when using the Shapley value allocations. However, computing the Shapley value for each customer has a time exponential in the size of the group and thus it is not scalable to the entire dataset. Instead, we compute the Shapley value for types of customers which has a more tractable time complexity polynomial to the number of types [15].

To define such types, we order the 3000 customers from *good* predictors (with low standard deviation) to *poor* predictors (with high standard deviation). Then, Figure 1(a) plots the expected payment of each customer when joining the P^+ redictive tariff on her own (i.e. as a *singleton*) factored into two components: one corresponding to her consumption (in grey color) and another one corresponding to her penalty for imbalance (in black color). It is worth noting that the penalty term is the only one that can be reduced by grouping and hence it is the focus of our analysis. The results show that while plotting the consumption term w.r.t the standard deviation forms an irregular pattern, as expected, the amount of penalty that each consumer pays as a singleton linearly increases with her standard deviation (i.e. the better a customer predicts her consump-

tion, the lower the price per unit that she pays for it). Given this, we cluster our customer dataset into three types based on their standard deviation: *good* predictors ($\sigma = 3.59$); *medium* predictors ($\sigma = 8.87$), and *poor* predictor ($\sigma = 17.90$). The characteristics of each customer type, as well as its frequency in the dataset (the number between parenthesis) are depicted in Figure 1(a).

Then, Figure 1(b) shows the expected payment of each customer type as a singleton and when grouped with the 3000 customers by types (i.e. using type frequencies as listed in Figure 1(a)) under the Shapley allocation scheme. The first thing we observe is that customers have incentives for grouping independently of their type (i.e. the expected penalty is significantly reduced for all types). Thus, the expected penalty under Shapley for a customer of the *good* type is reduced from £3.27 to £0.19 per month. Similarly, the expected penalty under Shapley for *medium* and *poor* customer types are reduced from £5.47 to £0.91 per month and from £16.35 to £5.24 per month respectively. Second, we observe that due to the fairness of the Shapley scheme, the penalty of the *poor* predictable type is higher than the penalty of the *medium* type which, in turn is higher than those of the *good* type.

Now, whilst these results give us an idea of the incentives that each customer has to join our group-buying scheme in a real-world group sample, it also leads to the question of whether they are subject to have a minority of *poor* predictable customers within the group (i.e. 94% of customers in the real-world dataset are of the *good* predictable type). To answer this, we run a second set of experiments in which we study how the incentives of customers vary with the inclusion of the different customer types.

In more detail, Figure 1(c) shows how the penalty in the Shapley value of a *good* and a *poor* predictor changes as the group is enlarged (up to 100 customers) with customers of the same type or with customers of the opposite type. Interestingly, it shows not only that a *good* predictor benefits from joining other *poor* predictors, but also that the incentive is greater than when joining with other *good* predictors. For example, compare the £1.8 per month penalty expected for a *good* predictor when joining a *poor* one, with the £2.31 per month penalty expected when joining another *good* predictor. A similar trend is observed for *poor* predicting customers: the reduction on their penalties when joining a group of *poor* predictors is greater than when joining a group of *good* predictors (£1.63 instead of £5.97 per month, for a group of 100 customers). Intuitively, this is because joining an agent with poor

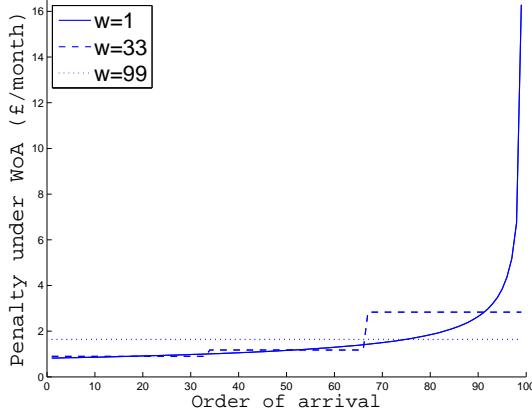


Figure 2: The expected penalty under the WoA scheme of poor predictable customers when joining a group of other poor predictable customers w.r.t. their order of joining.

predictability provides a greater aggregate reduction of the joint prediction error.

6.2 Payment allocation using arrivals

In this experiment, we consider customers (focusing, for ease of presentation, on the *poor* predictor type) joining the scheme over time and analyse the incentives that they have to commit early under a Window-of-Arrival (WoA) allocation scheme.

Figure 2 shows the expected payment of a *poor* predictor with respect to her order of arrival in a group composed of 99 *poor* predictor customers. We analyzed the results for three different sizes of arrival window. As expected, for a single window of arrival ($w = 99$) all customers are treated equally, regardless of their arrival order, and hence all are expected to have the same penalty payment (i.e. £1.64 per month). On the other end, for the case $w = 1$ (i.e. equal to the RM allocation), customers are heavily penalised for any delay in their arrival: the penalty goes from £0.82 per month for the first customer to £16.29 per month for the last (i.e. the penalty of the last customer is the same as the one she gets by joining alone). Finally, when dividing customers into three windows of arrival (i.e. $w = 33$) this sharp effect is smoothed by allocating payments more equally within windows. In particular, the expected monthly penalty of customers joining in the first window is £0.90, in the second £1.17 and in the third £2.83. We also observe that the difference in penalty due to the arrival is much more significant in the last windows (i.e. as more customers join the scheme the difference on the expected penalty between the first groups of arrival is reduced). For example, in the case $w = 1$, the penalty of the last customer is much higher than the penalty of the second-to-last one (compare £6.75 with £16.29 per month). This is a side-effect of concavity. The marginal cost contribution of a customer is always larger in smaller groups. Hence, as groups get bigger, the difference between the marginal contribution of a customer to one group with respect to another is almost insignificant.

7. CONCLUSIONS

This paper proposes a new tariff scheme that requires each consumer to predict a baseline for her future consumption, and computes the bill based on her actual consumption and on the deviation from the baseline. Our tariff has several advantages: it encourages users to provide reliable consumption predictions, as well as to join early in group buying initiatives. Moreover, while previ-

ous work has noted the potential of grouping consumers to reduce uncertainty in aggregate demand, ours is the first to characterise the incentives for joining electricity group buying initiatives, using principled concepts from cooperative game theory.

There are several open problems that are left for future work. On the theoretical side, while in this paper we studied incentives for cooperation w.r.t. a single tariff, it would be interesting to study the properties of our prediction-of-use game in a setting with multiple tariffs and competing providers (as is done in [8, 9] for volume-based discounts). On the practical side, it would be interesting to run trials with domestic customers to explore how consumer's adoption and behaviour are affected by the incentives given by our prediction-of-use group buying scheme.

8. ACKNOWLEDGMENTS

This work was supported by the iDEaS (www.ideasproject.info) and ORCHID (www.orchid.ac.uk) projects.

9. REFERENCES

- [1] C. Akasiadis and G. Chalkiadakis. Agent cooperatives for effective power consumption shifting. In *AAAI*, 2013.
- [2] E. Baeyens, E. Bitar, P. Khargonekar, and K. Poolla. Wind energy aggregation: A coalitional game approach. In *CDC-ECE*, pages 3000–’07, 2011.
- [3] G. Chalkiadakis, E. Elkind, and M. Wooldridge. *Computational Aspects of Cooperative Game Theory*. Morgan & Claypool Publishers, 2011.
- [4] M. Hazewinkel. *Theory of Errors. Encyclopedia of Mathematics*. Springer, 2001.
- [5] R. J. Kauffman, H. Lai, and C.-T. Ho. Incentive mechanisms, fairness and participation in online group-buying auctions. *ECRA*, 9:249–262, 2010.
- [6] R. Kota, G. Chalkiadakis, V. Robu, A. Rogers, and N. R. Jennings. Cooperatives for demand side management. In *ECAI*, pages 969–974, 2012.
- [7] C. Li, S. Chawla, U. Rajan, and K. Sycara. Mechanism design for coalition formation and cost sharing in group-buying markets. *ECRA*, 3:341–354, 2005.
- [8] T. Lu and C. Boutilier. Matching models for preference-sensitive group purchasing. In *ACM-EC*, pages 723–740, 2012.
- [9] R. Meir, T. Lu, M. Tennenholz, and C. Boutilier. On the value of using group discounts under price competition. In *AAAI*, 2013.
- [10] U. Ozen, H. Norde, and M. Slikker. On the convexity of newsvendor games. *International Journal of Production Economics*, 133:35–42, 2011.
- [11] D. C. Parkes. Online mechanisms. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, pages 411–439, 2007.
- [12] V. Robu, R. Kota, G. Chalkiadakis, A. Rogers, and N. Jennings. Cooperative virtual power plant formation using scoring rules. In *AAAI’12*, pages 370–376.
- [13] H. Rose, A. Rogers, and E. Gerding. A scoring rule-based mechanism for aggregate demand prediction in the smart grid. In *AAMAS*, pages 661–668, 2012.
- [14] L. S. Shapley. Cores of convex games. *International Journal of Game Theory*, 1:11–26, 1971.
- [15] T. Shrot, Y. Aumann, and S. Kraus. On agent types in coalition formation problems. In *AAMAS’10*, pages 757–764.
- [16] D. H. Steven Braithwait and M. O’Sheasy. *Retail Electricity Pricing And Rate Design In Evolving Markets*. Edison Electric Institute, 2007.
- [17] G. M. Team. Electricity and gas supply market report. Technical Report 176/11, The Office of Gas and Electricity Markets (Ofgem), December 2011.