

Multiple-Profile Prediction-of-Use Games

Andrew Perrault and Craig Boutilier

Department of Computer Science
University of Toronto, Toronto, Canada
{perrault, cebly}@cs.toronto.edu

Abstract. *Prediction-of-use (POU) games* [18] address the mismatch between energy supplier costs and the incentives imposed on consumers by fixed-rate electricity tariffs. However, the framework does not address how consumers should coordinate to maximize social welfare. To address this, we develop *MPOU games*, an extension of POU games in which agents report *multiple acceptable electricity use profiles*. We show that MPOU games share many attractive properties with POU games attractive (e.g., convexity). Despite this, MPOU games introduce new incentive issues that prevent the consequences of convexity from being exploited directly, a problem we analyze and resolve. We validate our approach with experimental results using utility models learned from real electricity use data.

Keywords: Energy, cooperative game theory, machine learning, economics

1 Introduction

Prediction-of-use games were developed by Vinyals et al. [18], hereafter VRRJ, to address the mismatch between the cost structure of energy suppliers and the incentive structure induced by the fixed-rate tariff faced by consumers. In most countries, energy suppliers face a two-stage market, where they purchase energy at lower rates in anticipation of future consumer demand and then reconcile supply and demand exactly at a higher rate at the time of realization through a balancing market [17]. The cost to energy suppliers is thus highly dependent on their ability to predict future consumption. Since consumers typically have little incentive to consume predictably, suppliers use past behavior to predict consumption. The resulting prediction uncertainty incurs some additional cost for suppliers.

One way to improve supplier predictions is to incentivize consumers to report *predictions of their own consumption*, thus gaining access to the consumer's private information about the future. VRRJ analyze mechanisms where flat tariffs are replaced with *prediction-of-use (POU) tariffs*, in which consumers make a payment based on both their actual consumption and the accuracy of their prediction. Similar tariffs have been deployed in practice [3]. VRRJ analyze the *cooperative game* induced by POU tariffs, in which consumers form *buying coalitions* that reduce (aggregate) consumption uncertainty, and find that, under

normally-distributed prediction error, the game is *convex*, a powerful property that significantly reduces the complexity of important problems in cooperative games.

While attractive, the POU model has two significant shortcomings. First, it does not model how consumers coordinate their consumption. A consumer’s optimal *consumption profile*—a random variable representing possible behaviors or patterns of energy consumption—depends on the profiles others use. In POU games, the only consumer choice is what coalition to join—consumer demand is represented by a *single* prediction, reflecting just one selected (or average) consumption profile. In essence, consumers predict their behavior without knowing anything about the other consumers in the game. While the POU model can offer social welfare gains when the profiles are selected optimally, we show they can result in significant welfare loss when profiles are uncoordinated. Second, POU games have no concept of consumer utility. While POU tariffs are intended to induce consumers to consume more predictably, the full welfare consequences of such changes are not measured in the POU framework.

We introduce *multiple-profile POU (MPOU) games*, which extend POU games to admit *multiple* consumer profiles (or “bids”). This allows consumers to coordinate the behaviors that change their predictions, this allowing full realization of the benefits of the POU model. We show that MPOU games have many of the same properties that make the POU model tractable, e.g, convexity, which makes the stable distribution of the benefits of cooperation easy to compute. However, MPOU games also present a new challenge in coalitional allocation: since one can only observe an agent’s (stochastic) consumption—not their underlying behavior—determining stabilizing payments for coalitional coordination requires novel techniques. We introduce *separating functions*, which incentivize agents to take a specific action in settings where actions are only *partially observable*.

We experimentally validate our techniques, using household utility functions that we learn (via structured prediction) from publicly-available electricity use data. We find that the MPOU model provides a gain of 3-5% over a fixed-rate tariff across several test scenarios, while a POU tariff *without* consumer coordination can result in losses of up to 30%. These experiments represent the first study of the welfare consequences of POU tariffs.

Section 2 reviews cooperative games, the POU model and related work. Section 3 introduces MPOU games and Section 4 proves their convexity. Section 5 outlines the new class of incentive problems that arises when the mechanism designer cannot (directly) observe an agent’s selected profile, and develops a general solution to that problem. In Section 6, we describe an approach for learning consumer utility models from real-world electricity usage data, and experimentally validate the value of MPOU games using these learned models in Section 7.

2 Background

2.1 Cooperative Games

We briefly overview the relevant aspects of cooperative game theory. A prediction-of-use game is an instance of a *cooperative game* [9]. In a cooperative game, the set N of agents divides into a set of coalitions, i.e., a disjoint partitioning of the agents. In a *profit game* the *characteristic function* $v : 2^N \rightarrow \mathbb{R}$ represents the value that any subset of agents can achieve by cooperating. A profit game is a tuple $\langle N, v \rangle$.

The agents in a coalition can decide to distribute the benefits of cooperation however they choose. An *allocation* is a payment function $t : N \rightarrow \mathbb{R}$ that assigns some payment (which may be negative) to each agent. An allocation is *efficient* if it distributes the entire value, i.e., $\sum_{i \in N} t(i) = v(N)$. A major goal of cooperative game theory is to find allocations that prevent agents from *defecting* from their coalition, thus achieving *stability*. An allocation that stabilizes the *grand coalition* of all agents is in the *core*.

Definition 1 Allocation t is in the *core* of profit game $\langle N, v \rangle$ if it is efficient and $\sum_{i \in S} t(i) \geq v(S)$ for all $S \subseteq N$.

The core is a strong stability concept that may not exist for general games. Another central solution concept is the *Shapley value* $s_C(i)$ of an agent i in coalition C , which emphasizes fairness and always exists. It values each agent according to their marginal value contribution when averaged over all join orders over C (i.e., the order in which agents are added to C):

$$s_C(i) = \sum_{S \subseteq C \setminus \{i\}} \frac{|S|!(|C| - |S| - 1)!}{|N|!} (v(S \cup \{i\}) - v(S)) \quad (1)$$

A convex game is one where the value contributed by an agent to a coalition never decreases as more agents are added to that coalition:

Definition 2 Profit game $\langle N, v \rangle$ is *convex* if $v(T \cup \{i\}) - v(T) \geq v(S \cup \{i\}) - v(S)$, for all $i \in N$, $S \subseteq T \subseteq N \setminus \{i\}$.

Convex games have important properties [15]: the grand coalition maximizes social welfare, the Shapley value is in the core, and a core allocation must exist and is computable in poly-time in the number of agents.

2.2 Prediction-of-Use Games

A *prediction-of-use (POU) game* is a tuple $\langle N, \Pi, \tau \rangle$. N is a set of agents, where each $i \in N$ uses electricity according to a *consumption profile* in Π , a normal random variable with mean μ_i and standard deviation σ_i , say, in kilowatt-hours (kWh). Let x_i denote i 's realized consumption, $x_i \sim \mathcal{N}(\mu_i, \sigma_i)$. Consumption

profiles are assumed to be common knowledge (we do not address elicitation or estimation of consumption here, but see below).

A *POU tariff* has the form $\tau = \langle p, \underline{p}, \bar{p} \rangle$, and is intended to better align the incentives of the consumer and electricity supplier, whose costs are greatly influenced by how predictable demands are. Each agent i is asked to predict a baseline consumption b_i , and is charged p for each unit of x_i , plus a penalty that depends on the accuracy of their prediction: \bar{p} for each unit their realized x_i exceeds the baseline, and \underline{p} for each unit it falls short:

$$\psi(x_i, b_i, \tau) = \begin{cases} p \cdot x_i + \bar{p} \cdot (x_i - b_i) & \text{if } b_i \leq x_i \\ p \cdot x_i + \underline{p} \cdot (b_i - x_i) & \text{if } b_i > x_i \end{cases} \quad (2)$$

To ensure agents have no incentive to artificially inflate consumption, we require $0 \leq \bar{p}$ and $0 \leq \underline{p} \leq p$ [18]. An agent i should report a baseline that minimizes her expected payment. VRRJ show that i does this by predicting $b^* = \mu_i + \sigma_i \Phi^{-1}(\frac{\bar{p}}{\bar{p} + \underline{p}})$, where Φ^{-1} is the inverse normal CDF. They also show that i 's expected payment under the optimal baseline is $\mu_i p + \sigma_i L(\underline{p}, \bar{p})$ where $L(\underline{p}, \bar{p}) = \int_0^{\frac{\bar{p}}{\bar{p} + \underline{p}}} \Phi^{-1}(y) dy$.

In order to be more predictable in the *aggregate*, agents may form a coalition C , where C reports its aggregate demand and is charged as if it were a single agent. C 's aggregate consumption is the sum of the normal random variables corresponding to the members' profiles, itself normal with mean $\mu(C) = \sum_{i \in C} \mu_i$ and standard deviation $\sigma(C) = \sqrt{\sum_{i \in C} \sigma_i^2}$. This aggregate prediction will generally have lower variance, thus reducing the total penalty payments made by C when facing POU tariffs (compared to members acting individually). VRRJ analyze *ex-ante* POU games where agents make all decisions and internal payments are based on expected consumption (realized consumption plays no role). In this setting, the characteristic value of a coalition C is $v(C) = -\mu(C)p - \sigma(C)L(\underline{p}, \bar{p})$. They show that, in the ex-ante setting, the POU game is convex.¹

2.3 Related Work

In addition to POU games [18,12], others have proposed the formation of co-operatives or coalitions among electricity consumers. Rose et al. [13] develop a similar mechanism for truthfully eliciting consumer demand. Kota et al. [6] and Akasiadas & Chalkiadakis [2] propose using coalitions to improve reliability and shift peak power loads. Perrault et al. [10] focus on the formation of groups of consumers with multiple profiles to reduce peak loads. None of this work offers the theoretical guarantees of VRRJ.

Beyond electricity markets, several authors have studied the problem of group purchasing in an AI context. Lu and Boutilier [7] study a restrictive class of buyer preferences (unit demand, only the supplier affects utility) and seller price

¹ Technically, they define the game as a *cost game* and show that the game is concave, while we use a profit game, but results from the two perspectives translate directly.

functions (volume discounts), which has strong theoretical guarantees. Similarly, optimally matching a group of cooperative buyers to sellers has been studied [14,8].

3 Multiple-Profile POU Games

We extend POU games by allowing agents to report *multiple profiles*, each reflecting different behaviors or consumption patterns, and each with an inherent utility or value reflecting comfort, convenience, flexibility or other factors. This will allow an agent, when joining or bargaining with a coalition, to trade off cost—especially the cost of predictability—with her inherent utility. A *multiple-profile POU (MPOU) game* is a tuple $\langle N, \Pi, V, \tau \rangle$. Given set of agents N , each agent $i \in N$ has a non-empty *set of demand profiles* Π_i , where each profile $\pi_k = \langle \mu_k, \sigma_k \rangle \in \Pi_i$ reflects a consumption pattern (as in a POU model). Agent i 's *valuation function* $V_i : \Pi_i \rightarrow \mathbb{R}$ indicates her value or relative preference (in dollars) for her demand profiles.² Admitting multiple profiles allows us to reason about an agent's response to the incentives that emerge with POU tariffs and in coalitional bargaining.

We use the definition of POU tariffs and agent baselines as in POU games above. Notice that the optimal baseline report for an agent is now defined relative to the profile they use. As in POU games, agents are motivated to form coalitions to reduce the relative variance in their predictions. However, for a coalition C to accurately report its aggregate demand, its members must select and commit to a specific usage profile. We denote an *assignment of profiles to agents* as $A : N \rightarrow \times_{i \in N} \Pi_i$. Under such an assignment, C 's consumption is normal, with mean $\mu(C, A) = \sum_{i \in C} \mu(A(i))$ and standard deviation $\sigma(C, A) = \sqrt{\sum_{i \in C} \sigma^2(A(i))}$. The aggregate value accrued by the coalition (prior to supplier payments) is the sum of its members' values: $V(C, A) = \sum_{i \in C} V_i(A(i))$.

As in VRRJ, we begin by analyzing ex-ante MPOU games, where agents make decisions and payments before consumption is realized. The characteristic value v of a coalition C is the maximum value that coalition can achieve in expectation under full cooperation, i.e., an optimal profile assignment and baseline report, namely: $v(C) = \max_A v(C, A)$, where

$$v(C, A) = V(C, A) - \mu(C, A)p - \sigma(C, A)L(\underline{p}, \bar{p}) \quad (3)$$

Notice that profile selection does not arise in the POU setting. In our MPOU model, coalition value is non-concave, even if integrality of the assignment variables is relaxed, because the last term is a negative square root: $\sigma(C, A) = \sqrt{\sum_{i \in C} \sigma^2(A(i))}$. We can perform the optimization using a mixed integer program by replacing the negative square root with a piecewise linear upper bound, which requires two binary variables per segment. As in other matching problems, we can relax the assignment variables: in practice, relaxed solutions that are very close to integral.

² Such profiles and values may be explicitly elicited or estimated using past consumption data (see Section 6).

4 Convexity of MPOU Games

It is natural to ask whether, like POU games, ex-ante MPOU games are convex, since convexity simplifies the analysis of stability and fairness. We show that this is the case.

Theorem 1 *The ex-ante MPOU game is convex.*

Since the ex-ante MPOU game is convex, the Shapley value is in the core, hence we can compute a core allocation by averaging the payments from any number of join orders. In our experiments, we approximate the Shapley value by sampling [4].

5 Incentives in MPOU Games

MPOU games introduce a new coordination problem for coalitions that does not occur in POU games. In a fully-cooperative MPOU game, a coalition C agrees on a joint consumption profile prior to reporting its (aggregate) predicted demand. Despite this agreement, an agent $i \in C$ may have incentive to use a profile that differs from the one agreed to. For instance, suppose agent i has two profiles, π_0 and π_1 , with $V_i(\pi_0) > V_i(\pi_1)$, and that to maximize the social welfare of C , i should use π_1 (and receive coalitional payment t_i). By deviating from her agreed upon profile, i can increase her net utility (from $V_i(\pi_1) + t_i$ to $V_i(\pi_0) + t_i$).

Typically, a penalty should be imposed for such a deviation to ensure that C 's welfare is maximized. Unfortunately, i 's profile cannot be directly observed. Only her realized consumption x_i is observable, and it is related only *stochastically* to her underlying behavior (adopted profile). As such, any such transfer or penalty in the coalitional allocation must depend on x_i , showing that an ex-ante analysis is insufficient for MPOU games (in stark contrast to POU games). Furthermore, since x_i is stochastic, it could have arisen from i using either profile (i.e., we have no direct signal of the i 's chosen profile), which makes the design of such transfers even more difficult. Finally, the poor choice of a transfer function may compromise the convexity of the ex-ante game, undermining our ability to compute core payments.

To address these challenges, we use a *separating function* $D_i(x_i)$. For each agent i , D_i maps i 's realized consumption to an additional *ex-post separating payment*.

Definition 3 $D_i : \mathbb{R} \rightarrow \mathbb{R}$ is a *separating function* for i under assignment A if $\mathbb{E}_{x_i \sim A(i)}[D_i(x_i)] > \mathbb{E}_{x_i \sim \pi}[D_i(x_i)]$ for any $\pi \in \Pi_i$ such that $\pi \neq A(i)$.

Intuitively, given such a D_i , we scale it so that the separating payment is large enough to prevent an agent from using any profile other than the one she is assigned, while ensuring its expected value is 0 to preserve the convexity of the MPOU game. This has the added benefit of not requiring elicitation of exact

profile utilities—we need only an upper bound on the difference in utilities across the profiles.

When an agent has only two profiles, finding D_i is straightforward: we let D_i be the PDF of the assigned profile minus the PDF of the unassigned profile. The proof for this statement is algebraic, using the fact that $\mathcal{N}(x; \mu_0, \sigma_0)\mathcal{N}(x; \mu_1, \sigma_1)$ has a closed form that is proportional to a normal PDF in x . All proofs from this section are in the appendix.

Theorem 2 *Let i be an agent with two profiles π_0 and π_1 and let $A(i) = \pi_0$. Then, w.l.o.g., $D_i(x_i) = \mathcal{N}(x_i; \mu_0, \sigma_0) - \mathcal{N}(x_i; \mu_1, \sigma_1)$ is a separating function for i under A .*

With more than two profiles, this approach does not always work. Instead, we can use a linear program (LP) to find coefficients of a linear combination of the profile PDFs. Formally, denote the PDFs of the profiles as $\mathcal{N}_i(x_i) = \langle \mathcal{N}(x_i; \mu_0, \sigma_0), \dots, \mathcal{N}(x_i; \mu_{|\Pi_i|-1}, \sigma_{|\Pi_i|-1}) \rangle$, their weights as \mathbf{y}_i , and search over $\mathbf{y}_i \in \mathbb{R}^{|\Pi_i|}$ for a separating function of the form $D_i(x_i, \mathbf{y}_i) = \mathbf{y}_i \cdot \mathcal{N}_i(x_i)$. We use an LP that minimizes the L_1 -norm of \mathbf{y}_i subject to $\mathbb{E}_{x_i \sim A(i)}[D_i(x_i, \mathbf{y}_i)] > \mathbb{E}_{x_i \sim \pi}[D_i(x_i, \mathbf{y}_i)]$ for all $\pi \in \Pi_i, \pi \neq A(i)$. Ideally, we would also like to minimize the variance of the separating payment, giving agents maximal certainty w.r.t. this payment; however, this objective is not tractable in an LP (we leave this question to future work). In our experiments below, we do, however, assess the variance of the separating payment.

A feasible \mathbf{y}_i corresponds to a linear combination of vectors whose sum has only positive entries. We call these the *difference vectors* of D_i . While we cannot prove that a feasible \mathbf{y}_i always exists, viewing the problem in terms of difference vectors suggests why they exist in practice:

Definition 4 Let $A(i)$ be π_0 (w.l.o.g.). For each profile $\pi_k \in \Pi_i$ the *difference vector* $\mathbf{d}_k = \mathbb{E}_{x \sim \pi_k}[\mathcal{N}(x; \mu_0, \sigma_0)] - \langle \mathbb{E}_{x \sim \pi_k}[\mathcal{N}(x; \mu_1, \sigma_1)], \dots, \mathbb{E}_{x \sim \pi_k}[\mathcal{N}(x; \mu_{|\Pi_i|-1}, \sigma_{|\Pi_i|-1})] \rangle$.

Note that these vectors do not depend on \mathbf{y}_i . We can restate the LP constraints using difference vectors:

Theorem 3 *Let i be an agent with profiles Π_i and let A assign a profile to i . There exists $\mathbf{y}_i \in \mathbb{R}^{|\Pi_i|}$ that makes $D_i(x_i, \mathbf{y}_i)$ a separating function if and only if there is a linear combination of the difference vectors of $D_i(x_i, \mathbf{y}_i)$ that has only positive entries.*

Proof. First, we prove the forward direction. Let \mathbf{c} be the coefficients of the linear combination of the difference vectors that has only positive entries, i.e., $\sum_{k \in |\Pi_i|} \mathbf{c}_k \mathbf{d}_k = \mathbf{b}$ where \mathbf{b} is element-wise positive. Then, $\mathbb{E}_{x_i \sim A(i)}[D_i(x_i, \mathbf{c})] - \mathbb{E}_{x_i \sim \pi}[D_i(x_i, \mathbf{c})] = \mathbf{c} \mathbf{d}_k = \mathbf{b}_{k-1}$. Since \mathbf{b} is element-wise positive, letting $\mathbf{y}_i = \mathbf{c}$ makes $D_i(x_i, \mathbf{y}_i)$ a separating function.

The reverse direction is also straightforward. Suppose $D_i(x_i, \mathbf{y}_i)$ is a separating function. Then, let $\mathbf{b}_{k-1} = \mathbb{E}_{x_i \sim A(i)}[D_i(x_i, \mathbf{c})] - \mathbb{E}_{x_i \sim \pi}[D_i(x_i, \mathbf{c})] = \mathbf{y}_i \cdot \mathbf{d}_k$.

Thus, taking \mathbf{y}_i as the coefficients of the linear combination of difference vectors equals \mathbf{b} , which has only positive entries.

Corollary 1 *Let \mathbf{d}_k be the difference vectors for agent i . If the difference vectors are linearly independent, a setting of \mathbf{y}_i exists that makes $D_i(x_i, \mathbf{y}_i)$ a separating function.*

Proof. If the difference vectors are linearly independent, there exists a coefficient vector \mathbf{c} that makes $\sum_{k \in |\Pi_i|} \mathbf{c}_k \mathbf{d}_k$ elementwise positive. We can take $\mathbf{y}_i = \mathbf{c}$.

We generally expect a random set of vectors to be linearly independent when their entries are drawn from the reals. We have yet to encounter an instance where a separating function does not exist in our experiments. It is an open question as to whether a separating function of this form always exists.

6 Learning Utility Models




Fig. 1: The form of the learned valuation model. $NN(10)$ denotes a neural network with 10 hidden units.




Fig. 2: Translating the valuation function to pass through the origin.

To empirically test the MPOU framework and our separating functions, we require consumer utility functions. As we know of no data set with such utility functions, we learn household (agent) utility models from real electricity usage data from Pecan Street Inc. [11].³ We define our prediction period as 4-7 pm each day, when electricity usage typically peaks in Austin, Texas, where the data was collected. We decompose utility into two parts: $V_i^{(\mu)}(w, \mu)$ describes the value an agent i derives from her mean consumption given a vector w of weather conditions; and $V_i^{(\sigma)}(\sigma, \mu)$ represents utility derived from variance in consumption behavior. Agent i 's utility is $V_i(w, \mu, \sigma) = V_i^{(\mu)}(w, \mu)V_i^{(\sigma)}(\sigma, \mu)$.

³ Publicly available at pecanstreet.org.

Estimating $V_i^{(\mu)}$ is difficult, since we lack data for some aspects of the problem. Thus, we make some simplifying assumptions: (i) consuming 0 kWh yields value \$0; and (ii) $V_i^{(\mu)}(w, \mu)$ is concave and increasing. We learn a model for each of 25 households that have complete data from 2013–15 (about 1100 data points per household), using select weather conditions w and mean consumption between 4–7 pm as input, and outputting value (in dollars). We use this valuation function to predict consumption by maximizing an agent’s net utility under the observed price:

$$V_i^{(\mu)}(w, \mu) = z_i^{(0)}(w) \left(\mu - z_i^{(1)}(w) \right)^{z_i^{(2)}(w)} + z_i^{(3)}(w) \quad (4)$$

constraining $z_i^{(0)} > 0$, $z_i^{(1)} > 0$, $0 < z_i^{(2)} < 1$, $z_i^{(3)}(w) \geq 0$ (Figure 1 depicts the utility model). The term $z_i^{(3)}(w)$ has no influence on predictions: it can be viewed as inherent value due to weather, and is used to account for the flexibility provided by the $z_i^{(1)}$ term, which may create valuations where consumption 0 yields negative value (violating our assumptions). To prevent this, we set $z_i^{(3)}(w)$ to ensure the tangent at the predicted consumption for \$0.64 (the largest price in the data set) passes through (0,0). (see Figure 2). When this tangent crosses the y -axis above 0, we set $z_i^{(3)}(w) = 0$ and splice in an exponential function of the form ax^b that passes through (0,0) and matches the derivative at the splicing point.

For training, we use the model to predict consumption by solving the net utility maximization problem, $\max_{\mu}(V_i(w, \mu) - \mu p)$, yielding:

$$\hat{\mu}(w, p) = \frac{p}{z_i^{(0)}(w) z_i^{(2)}(w)} \frac{1}{z_i^{(2)}(w) - 1} + z_i^{(1)}(w) \quad (5)$$

We represent $z_i^{(0)}$, $z_i^{(1)}$ and $z_i^{(2)}$ in fully-connected single-layer neural networks, each with 10 hidden units and ReLU activations, and train the model with backpropagation. We implement the model in TensorFlow [1] using the squared error loss function and the Adam optimizer [5]. We use Dropout [16] with a probability of 0.7 on each hidden unit.

Model	Mean	Std. dev.	Mean	Std. dev.
	train	train	test	test
	RMSE	RMSE	RMSE	RMSE
Valuation	0.137	0.0168	0.148	0.0194
Unstructured	0.142	0.0226	0.144	0.0284
Constant	0.204	0.0345	0.205	0.0411

Table 1: Comparison of model prediction accuracy by root-mean-square error (RMSE). We divide each household’s consumption amounts by their largest observed consumption.




Fig. 3: Learned value models for three of the 25 households with consumption mean (kWh) on the x -axis and value (\$) on the y -axis. The red line represents the median weather conditions. The dotted line represents the median day with 90th percentile or higher temperature. The dashed and green lines are the same for sunshine and humidity, respectively.

We split the data into 80% train and 20% test for each household. Table 1 compares the prediction accuracy of our model (“valuation”) to (i) an unstructured neural network, and (ii) the best constant prediction for each household. The unstructured net learns a mapping from $\langle w, p \rangle$ to μ directly using 10 hidden units, without an intervening utility model.⁴ The best constant prediction disregards weather and price data, and simply predicts average consumption for that household. Table 1 shows that the valuation model overfits somewhat, but that predictive accuracy is on par with the unstructured model. This shows that our constraints on the form of the valuation function are not unduly restrictive and validates the value predictions produced by these learned models. However, we believe these value functions significantly underestimate value because we lack consumption observations when the price is higher than \$0.64.

Figure 3 shows the learned valuation functions for three of the 25 households, and the rest are shown in the appendix. Each line represents a household’s response to different weather conditions. While temperature is the most significant predictor of power usage, different households appear to exhibit sensitivity to different factors (e.g., the household on the right, is highly sensitive to humidity).

6.1 Modeling Unpredictable Consumption

Unfortunately, we do not have access to electricity usage data where consumers are charged differently depending on the accuracy of their predictions. Our model of the value of unpredictable consumption is thus speculative, but uses the Pecan Street data as a starting point. We assume that each household chooses the σ that maximizes its utility (since they are not being charged for σ), and that it has an optimal fraction β_i of σ/μ that does not depend on other conditions. We estimate β_i from the data by treating each data point as having an observed σ equal to the absolute error in consumption prediction made by the learned valuation model. We assume no value is gained by increasing σ above the optimal

⁴ Our other implementation choices are the same as the valuation model, except we use Dropout of 0.5.

ratio, and use an exponential to represent the loss in value when σ is reduced,

$$V_i^{(\sigma)}(\sigma, \mu) = \max\left(\frac{\mu/\sigma}{\beta_i}, 1\right)^{\gamma_i}, \quad (6)$$

where γ_i is a constant representing i 's cost for being predictable. A higher γ_i means that consumer i values variance more highly. In our experiments, we sample γ_i from the uniform distribution over the interval $[0.1, 2]$.

7 Experiments

The questions we study experimentally are: (i) how important is consumer coordination under POU tariffs; (ii) what is the overall social welfare gain from using an MPOU model vs. a flat tariff; (iii) how important is an agent's choice of reported profiles; and (iv) what are the variances of the payments introduced by the separating functions. We first describe the experimental setup: how we select agents, profiles and tariffs. For each trial, we select weather conditions w uniformly at random from the Pecan Street data. To generate agents, we sample from our 25 learned household utility models, using w as input and adding a small amount of zero mean noise to the model parameters. We sample γ_i from the uniform distribution $[0.1, 2]$ for each agent i . Each data point is an average of 100 trials with 5000 agents, unless otherwise noted. One of the goals of our experiments is to study the consequences of different choices of reported profile. To do this, we vary the way profiles are generated. Each agent has four profiles: a *base profile* (predicted to be optimal under a flat rate tariff with rate equal to the fixed-rate p of the POU tariff), and three others reflecting reduced consumption mean or variance. The first reduces the base profile *mean* by the amount required to reduce value by $u\%$, which we call the *profile spacing*. The second reduces *variance* to reduce value by $u\%$. The third reduces both. We vary u throughout the experiments.

To generate tariffs, we vary the amount of emphasis each puts on accurate predictions vs. the amount consumed. We let the *predictivity emphasis* (PE) of a tariff w.r.t. a group of agents be the fraction of the expected total cost paid for prediction penalties when each uses her base profile. In practice, PE should be set to match the properties of the reserve power generation capacity that is available: a higher PE corresponds to more expensive reserves. A tariff is *revenue-equivalent* to another with respect to a specific set of profiles if the revenue of the two is the same for that set. All of our tariffs will be revenue-equivalent with respect to the set of base profiles. To find a revenue-equivalent tariff with a certain PE, we use a numerical solver to find a tariff of the form $\langle p, r, r \rangle$ with the appropriate total cost. Intuitively, a higher PE should result in larger benefits from POU tariffs, and we find that to be the case in our experiments.

To generate Shapley values, we sampled a number of join orders equal to the logarithm of the number of agents in the instance. Shapley values were very close to linear in the standard deviation of the assigned profile. With 5000 agents and 100 trials, the average Shapley payment for prediction was \$0.41 per kWh of

uncertainty across trials with PE 10%, and \$0.82 per kWh with PE 20%.⁵ Within a single trial, the standard deviation of this ratio was less than 1.0×10^{-14} on average, suggesting that once Shapley payments for a small number of agents are calculated, those for the remainder can be approximated very accurately. There may also be an accurate closed-form approximation for the Shapley payment in MPOU games.

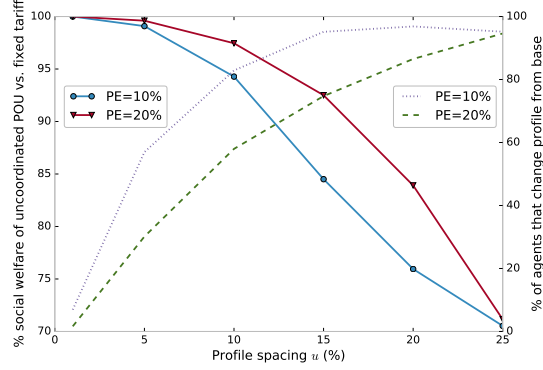


Fig. 4: Profile spacing vs. % of social welfare of fixed-rate tariff for uncoordinated POU setting and % of agents that change profile. 5000 agents, 100 trials.

7.1 Results

We first address the question of how important it is for agents to coordinate their consumption under a POU tariff. We define the *uncoordinated POU setting* as the scenario where agents are subject to a POU tariff, but do not coordinate their consumption behavior, i.e., each agent uses the profile that individually maximizes her net utility relative to that POU tariff. Figure 4 shows the social welfare derived by agents in the uncoordinated POU setting as a percentage of their social welfare under a revenue-equivalent fixed-rate tariff. We see that the average social welfare achieved in the uncoordinated POU setting is less than that of the fixed rate setting for all profile spacings. Individual agents react to the POU tariff by increasing their predictivity, and thus decreasing their realized value, but they do not account for the predictivity discount that results from being part of a coalition. As profile spacing increases, more agents shift away from their base profile and social welfare decreases, reaching 70% when spacing is 25%. These results underscore the need for a way for agents to coordinate their profile choices under POU tariffs and highlight one of the main challenges of successfully implementing a POU tariff in practice.

⁵ This and other tariffs in this section have $0.2 \leq p = \bar{p} \leq 1.5$.




Fig. 5: Profile spacing vs. social welfare % gain from fixed-rate tariff and % of agents that change profile. 5000 agents, 100 trials.

Next, we study the social welfare gain that can be achieved by a POU tariff when agents coordinate optimally under the MPOU framework. Figure 5 shows the effect of profile spacing (u) on the welfare gained by switching from a fixed-rate tariff to a revenue-equivalent POU tariff.⁶ Overall welfare gains are moderate, around 3.13% for PE of 10% and 4.4-4.9% for PE of 20%. A higher PE results in a larger social welfare gain because agents only benefit from cooperating when trading off predictivity for inherent utility. Profile spacing appears to have limited impact on social welfare gain, suggesting that most of the gain is achieved by the effective reduction in fixed-rate price under a POU tariff. We note that these experiments are the first to study the resulting social welfare gain from a POU tariff, as previous papers focused on non-welfare aspects.

Figure 5 appears to indicate that personalizing profile spacing based on each agent's value for predictivity would increase social welfare further. We can see this because increasing profile spacing increases welfare up to a spacing of 15% for both PE levels, but the number of agents that shift profiles decreases as spacing is increased (shown on the right-side axis). Thus, we hypothesize that welfare could be further increased if agents with higher γ spaced their profiles farther apart than those with lower.

Next, we address the question of uncertainty introduced by separating payments. Recall that while separating payments have expectation zero, they introduce additional uncertainty to agent payments. We find that the amount of uncertainty introduced is, in fact, minimal, and decreases with instance size and increased PE. The standard deviation of the separating payment is on average 15-20% of predictivity payment for PE of 10% and 7.5-10% for PE of 20%, and increases slightly as profile spacing increases. Note that only agents that actually require a separating function are taken into account, around 1-2% of all

⁶ Each instance took around 3 minutes on a single thread of 2.6 Ghz Intel i7, 8 GB RAM.




Fig. 6: Comparison of the standard deviation of the separating function payment to the ex-ante payment for prediction accuracy. Bars show one standard deviation. 5000 agents, 100 trials.

agents for PE of 10% and 5-10% for PE of 20%, on average. More agents require separating payments as PE increases, but the uncertainty introduced by each decreases. Note that these are uncertainties for a single instance of the game, and if the game is played repeatedly (e.g., every day), the aggregate uncertainty will decrease as the independent random variables are added.

Figure 7 shows the same uncertainty ratio for a single large instance versus the predictivity flexibility (γ) of each agent. This instance has PE of 20%, 100,000 agents, profile spacing of 15% and takes 90 min. to solve. The ratio is shown for the 4876 agents that require separating functions. The magnitude of the introduced uncertainty is smaller in this larger instance with an average of 2.07% (and not exceeding 3% for any agent). In addition, predictivity flexibility has little affect on the introduced uncertainty: the linear least-squares fit (red line) has slope of less than 10^{-4} .

8 Conclusion

We have introduced *multiple-profile POU (MPOU) games*, a framework for coordinating agent behavior under POU tariffs. MPOU games allow agents to express their consumption utility functions, while maintaining convexity of the basic POU model. MPOU games introduce a new class of incentive problems due to agent actions being partially observable: we introduce *separating payments* to restore proper incentives. Our experimental utility models are learned from historical electricity usage data in a novel way. Our experiments show that, while social welfare gained by introducing the MPOU model (w.r.t. a fixed-rate tariff) appear moderate, the gains relative to a POU tariff are substantial. The gains over a fixed-rate tariff may be worthwhile in a large system and may be further enhanced by more sophisticated agent utility and behavior profile models. They




Fig. 7: Comparison of the standard deviation of the separating function payment to the ex-ante payment for prediction accuracy. 100,000 agents of which 4876 required separating functions. Line shows linear best fit.

depend both on the predictivity emphasis (PE) of reserve generation and on consumers’ value for consuming unpredictably, which are both areas where more real-world data is needed. We find that the uncertainty introduced by separating payments decreases as instance size increases, and decreases in aggregate as more iterations of the game are played. Increased PE increases the number of agents that need separating functions, but the uncertainty introduced decreases.

Interesting future directions for POU/MPOU games remain. Greater access to household utility data, especially for variance of consumption, and data about the PE of generation mixes would allow us to more precisely test social welfare gain. In addition, it would be desirable to allow agents to make predictions contingent on intermediate predictions (e.g., of weather) thus reducing the need for agents to make accurate weather forecasts. While our discussion of POU and MPOU games has focused on electricity markets, we believe the approach may be more widely applicable in other cases where agents are contending with a scarce resource, e.g., cloud computing.

9 Acknowledgments

Perrault was supported by OGS. We gratefully acknowledge the support of NSERC. We thank Valentin Robu, Meritxell Vinyals, Marek Janicki, Jake Snell, and the anonymous reviewers for their suggestions.

References

1. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A.,

- Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-scale machine learning on heterogeneous systems (2015), <http://tensorflow.org/>, software available from tensorflow.org
2. Akasiadis, C., Chalkiadakis, G.: Agent cooperatives for effective power consumption shifting. In: Proceedings of the Twenty-seventh AAAI Conference on Artificial Intelligence (AAAI-13). pp. 1263–1269. Bellevue, WA (2013)
 3. Braithwait, S., Hansen, D., O’Sheasy, M.: Retail electricity pricing and rate design in evolving markets. Edison Electric Institute pp. 1–57 (2007)
 4. Castro, J., Gómez, D., Tejada, J.: Polynomial calculation of the Shapley value based on sampling. *Computers & Operations Research* 36(5), 1726–1730 (2009)
 5. Kingma, D.P., Ba, J.L.: Adam: A method for stochastic optimization. In: Proceedings of the 3rd ACM SIGKDD International Conference on Learning Representations (ICLR-15). San Diego (2015)
 6. Kota, R., Chalkiadakis, G., Robu, V., Rogers, A., Jennings, N.R.: Cooperatives for demand side management. In: Proceedings of the Twenty-First European Conference on Artificial Intelligence (ECAI-12). pp. 969–974. Montpellier, France (2012)
 7. Lu, T., Boutilier, C.: Matching models for preference-sensitive group purchasing. In: Proceedings of the Thirteenth ACM Conference on Electronic Commerce (EC’12). pp. 723–740. Valencia, Spain (2012)
 8. Manisterski, E., Sarne, D., Kraus, S.: Enhancing cooperative search with concurrent interactions. *Journal of Artificial Intelligence Research* 32(1), 1–36 (2008)
 9. Osborne, M.J., Rubinstein, A.: *A Course in Game Theory*. MIT Press, Cambridge (1994)
 10. Perrault, A., Boutilier, C.: Approximately stable pricing for coordinated purchasing of electricity. In: Proceedings of the Twenty-fourth International Joint Conference on Artificial Intelligence (IJCAI-15). Buenos Aires (2015)
 11. Rhodes, J.D., Upshaw, C.R., Harris, C.B., Meehan, C.M., Walling, D.A., Navrátil, P.A., Beck, A.L., Nagasawa, K., Fares, R.L., Cole, W.J., et al.: Experimental and data collection methods for a large-scale smart grid deployment: Methods and first results. *Energy* 65, 462–471 (2014)
 12. Robu, V., Vinyals, M., Rogers, A., Jennings, N.R.: Efficient buyer groups for prediction-of-use electricity tariffs. In: Proceedings of the Twenty-eighth AAAI Conference on Artificial Intelligence (AAAI-14). pp. 451–457. Quebec City (2014)
 13. Rose, H., Rogers, A., Gerding, E.H.: A scoring rule-based mechanism for aggregate demand prediction in the smart grid. In: Proceedings of the Eleventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-12). pp. 661–668. Valencia, Spain (2012)
 14. Sarne, D., Kraus, S.: Cooperative exploration in the electronic marketplace. In: Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI-05). pp. 158–163. Pittsburgh (2005)
 15. Shapley, L.S.: Cores of convex games. *International Journal of Game Theory* 1, 11–26 (1971)
 16. Srivastava, N., Hinton, G.E., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15(1), 1929–1958 (2014)
 17. Team, G.M.: Electricity and gas supply market report. Tech. Rep. 176/11, The Office of Gas and Electricity Markets (Ofgem) (December 2011)

18. Vinyals, M., Robu, V., Rogers, A., Jennings, N.R.: Prediction-of-use games: a cooperative game theory approach to sustainable energy tariffs. In: Proceedings of the Thirteenth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-14). pp. 829–836. Paris (2014)

10 Proofs

To prove convexity, we begin with a technical lemma.

Lemma 1. *Let $\langle N, \Pi, V, \tau \rangle$ be an MPOU game. Let $i \in N$ and $S \subset T \subseteq N \setminus \{i\}$ and $j \in T \setminus S$.*

$$v(S \cup \{i\}) - v(S) \leq v(S \cup \{i, j\}) - v(S \cup \{j\}) \quad (7)$$

Proof. We let $A^*(S)$ denote the assignment of profiles that maximizes the social welfare of S . In the case where there are multiple social welfare-maximizing configurations of S , we use the one with highest aggregate variance. We observe that $v(T, A^*(S)) \leq v(T)$ because $A^*(S)$ imposes a constraint on the behavior of S . For technical reasons, we break the proof into two cases based on whether it is more beneficial for i) i to join coalition S when S is configured to maximize $v(S \cup \{i\})$ or ii) i to join coalition $S \cup \{j\}$ when $S \cup \{j\}$ is configured to maximize $v(S \cup \{j\})$.

Case 1. $v(S \cup \{i\}) - v(S, A^*(S \cup \{i\})) > v(S \cup \{i, j\}, A^*(S \cup \{j\})) - v(S \cup \{j\})$

This inequality implies that $\sigma(S \cup \{j\}, A^*(S \cup \{j\})) < \sigma(S, A^*(S \cup \{i\}))$. Since j contributes a non-negative amount of variance, $\sigma(S, A^*(S \cup \{j\})) \leq \sigma(S \cup \{j\}, A^*(S \cup \{j\}))$, and likewise, $\sigma(S, A^*(S \cup \{i\})) \leq \sigma(S \cup \{i\}, A^*(S \cup \{i\}))$. Applying these inequalities yields $\sigma(S, A^*(S \cup \{j\})) < \sigma(S \cup \{i\}, A^*(S \cup \{i\}))$, implying:

$$\begin{aligned} v(S \cup \{j\}) - v(S, A^*(S \cup \{j\})) &< \\ v(S \cup \{i, j\}, A^*(S \cup \{i\})) - v(S \cup \{i\}) \end{aligned} \quad (8)$$

Then, applying the inequalities $v(S, A^*(S \cup \{j\})) \leq v(S)$ and $v(S \cup \{i, j\}, A^*(S \cup \{i\})) \leq v(S \cup \{i, j\})$, and rearranging terms:

$$v(S \cup \{i\}) - v(S) < v(S \cup \{i, j\}) - v(S \cup \{j\}) \quad (9)$$

which is a stronger version of the lemma.

Case 2. $v(S \cup \{i\}) - v(S, A^*(S \cup \{i\})) \leq v(S \cup \{i, j\}, A^*(S \cup \{j\})) - v(S \cup \{j\})$

Applying the inequality $v(S, A^*(S \cup \{i\})) \leq v(S)$ on the left side yields:

$$v(S \cup \{i\}) - v(S) \leq v(S \cup \{i, j\}, A^*(S \cup \{j\})) - v(S \cup \{j\}) \quad (10)$$

Applying on the right side $v(S \cup \{i, j\}, A^*(S \cup \{j\})) \leq v(S \cup \{i, j\})$ yields the lemma:

$$v(S \cup \{i\}) - v(S) \leq v(S \cup \{i, j\}) - v(S \cup \{j\}) \quad (11)$$

Theorem 1 *The ex-ante MPOU game is convex.*

Proof. If $S = T$, then $v(S \cup \{i\}) - v(S) = v(T \cup \{i\}) - v(T)$ since the welfare-maximizing configurations of S and T are the same. If $S \subset T$, we repeatedly apply Lemma 1 to “grow” S one agent a time, creating a series of inequalities, until we relate S and T .

Theorem 2 *Let i be an agent with two profiles π_0 and π_1 and let $A(i) = \pi_0$. Then, w.l.o.g., $D_i(x_i) = \mathcal{N}(x_i; \mu_0, \sigma_0) - \mathcal{N}(x_i; \mu_1, \sigma_1)$ is a separating function for i under A .*

Proof. We will show that the minimum of $\mathbb{E}_{x \sim \mathcal{N}(\mu_0, \sigma_0)}[\mathcal{N}(x; \mu_0, \sigma_0) - \mathcal{N}(x; \mu_1, \sigma_1)] - \mathbb{E}_{x \sim \mathcal{N}(\mu_1, \sigma_1)}[\mathcal{N}(x; \mu_0, \sigma_0) - \mathcal{N}(x; \mu_1, \sigma_1)]$ occurs when $\mu_1 = \mu_0$ and $\sigma_1 = \sigma_0$ and that the value of the expression at that point is positive.

We make use of the fact that $\mathcal{N}(x; \mu_1, \sigma_1)\mathcal{N}(x; \mu_2, \sigma_2)$ is a function proportional to the PDF of a normal distribution. Specifically,

$$\begin{aligned} & \mathcal{N}(x; \mu_0, \sigma_0)\mathcal{N}(x; \mu_1, \sigma_1) = \\ & \mathcal{N}\left(\mu_0; \mu_1, \sqrt{\sigma_0^2 + \sigma_1^2}\right) \mathcal{N}\left(x; \frac{\sigma_0^{-2}\mu_0 + \sigma_1^{-2}\mu_1}{\sigma_0^{-2} + \sigma_1^{-2}}, \frac{\sigma_0^2\sigma_1^2}{\sigma_0^2 + \sigma_1^2}\right) \end{aligned} \quad (12)$$

Then, by expanding terms and applying Equation 12:

$$\begin{aligned} & \mathbb{E}_{x \sim \mathcal{N}(\mu_0, \sigma_0)}[\mathcal{N}(x; \mu_0, \sigma_0) - \mathcal{N}(x; \mu_1, \sigma_1)] - \\ & \mathbb{E}_{x \sim \mathcal{N}(\mu_1, \sigma_1)}[\mathcal{N}(x; \mu_0, \sigma_0) - \mathcal{N}(x; \mu_1, \sigma_1)] \\ & = \frac{1}{2\sigma_0\sqrt{\pi}} - 2\mathcal{N}\left(\mu_1; \mu_0, \sqrt{\sigma_0^2 + \sigma_1^2}\right) + \frac{1}{2\sigma_1\sqrt{\pi}} \end{aligned} \quad (13)$$

We then minimize with respect to μ_1 and σ_1 . Since the middle term is the only one that contains μ_1 , we can minimize it separately:

$$- \frac{2}{\sqrt{2\pi(\sigma_0^2 + \sigma_1^2)}} \exp\left(-\frac{(\mu_0 - \mu_1)^2}{2(\sigma_0^2 + \sigma_1^2)}\right) \quad (14)$$

Since the argument of the exponent is always non-positive, it is maximized when it is zero, i.e., $\mu_1 = \mu_0$. Thus, we can make this substitution and rewrite the overall expression:

$$\frac{1}{2\sigma_0\sqrt{\pi}} - \frac{2}{\sqrt{2\pi(\sigma_0^2 + \sigma_1^2)}} + \frac{1}{2\sigma_1\sqrt{\pi}} \quad (15)$$

Taking the derivative with respect to σ_1^2 and setting it to zero yields two real roots of $\sigma_0 = \pm\sigma_1$. The second derivative at these points is positive. Thus, it is a minimum.

The value of the original expression at this point is 0 and positive otherwise.

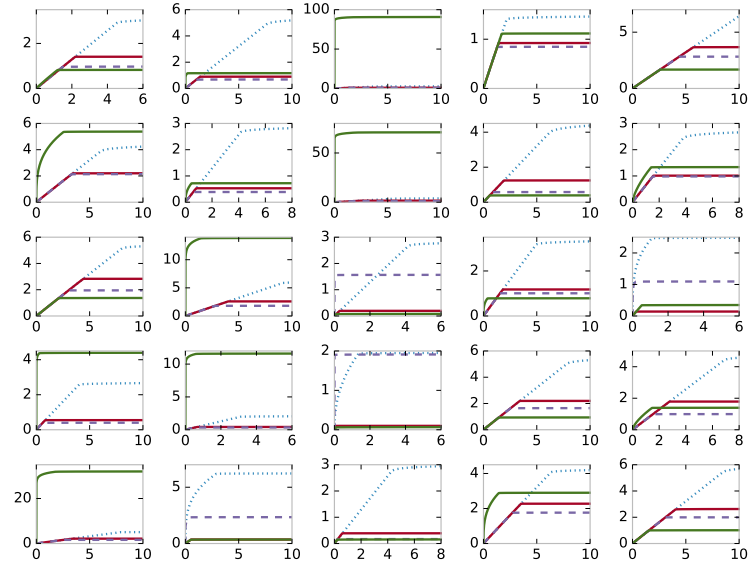


Fig.8: Learned value models for the 25 households with consumption mean (kWh) on the x -axis and value (\$) on the y -axis. The red line represents the median weather conditions. The dotted line represents the median day with 90th percentile or higher temperature. The dashed and green lines are the same for sunshine and humidity, respectively.