Multiple-Profile Prediction-of-Use Games

(Extended Abstract)

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 were developed by Vinyals et al. [4] to address the mismatch between the costs of energy suppliers and the incentives imposed on consumers by a fixed-rate electricity tariff. However, the POU games framework does not address the question of how consumers should coordinate to maximize social welfare. To address this, we develop an extension to POU games, which we term *MPOU games*, where agents report *multiple acceptable electricity use profiles.* We show that MPOU games have many of the same properties that make 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; game theory; machine learning; economics

1. INTRODUCTION

Prediction-of-use games were developed by Vinyals et al. [4], 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 [3]. The cost to energy suppliers is thus highly dependent on their ability to predict future consumption, but consumers typically have no incentive to consume predictably. Suppliers use past behavior to predict consumption and pay for their uncertainty.

One way to improve supplier predictions of consumption 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, primarily directed at industrial consumers [1]. 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 major shortcomings. First, consumers cannot coordinate their consumption. The optimal *consumption profile*, which we think of as a random variable representing an acceptable behavior or electricity consumption pattern, depends on what profiles others use. In the POU model, consumer demands are represented by a *single* prediction, reflecting just one selected (or average) consumption profile, and the only choice that consumers make is what coalition to join. Thus, consumers must choose their consumption profile without knowing anything about the other consumers in the game. We show the the POU model provides social welfare gains when the profiles are selected optimally, but can result in social welfare losses when the profiles are selected badly.

Second, POU games have no concept of consumer utility. Clearly the motivation of POU tariffs is to induce consumers to consume more predictably, but the consequence of this change cannot be measured in the POU framework.

We introduce *multiple-profile POU (MPOU) games*, which extend POU games to admit *multiple* consumer profiles or "bids," allowing for the coordination of consumer behavior. Thus, the benefits of the POU model can be fully realized.

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 a general framework of *separating functions* to address the question of incentivizing agents to take a particular action when the action is only partially observable. While our system of coordination is centralized, separating functions could be used to address the problem of monitoring and enforcement in a decentralized mechanism as well.

We experimentally validate our techniques. To do this,

*Supported by OGS. We gratefully acknowledge the support of NSERC.

†Currently on leave at Google, Inc., Mountain View.

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

we learn household utility functions from publicly-available electricity use data using a structured prediction framework. We find that the MPOU framework provides a gain of 3-5% over a fixed-rate tariff over several test scenarios and that a POU tariff without consumer coordination can result in losses of up to 30% from a fixed-rate tariff. To our knowledge, these experiments represent the first study of the social welfare consequences of a POU tariff.

2. MULTIPLE-PROFILE POU GAMES

We extend POU games by allowing agents to report multiple profiles, each reflecting different behaviors or consumption patterns, and each associated 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 inherent utility. A *multiple-profile POU (MPOU) game* is a tuple $\langle N, \Pi, \tau, V \rangle$. 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* $V_i : \Pi_i \to \mathbb{R}$ indicates her value or relative preference (in dollars) for her demand profiles.[1] Admitting multiple profiles allows us to reason about an agent's response to the incentives that emerge with POU tariffs (and coalitional bargaining).

As in POU games, 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. The optimal demand that an agent reports depends on her consumption as well as the parameters of the tariff. In MPOU games, the optimal demand depends on her choice of consumption profile. 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.

As in VRRJ, we analyze 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.

3. THEORETICAL RESULTS

We begin by showing that, as with POU games, MPOU games are convex. However, MPOU games introduce a new coordination problem for coalitions. 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 be incentivized to use a profile that differs from the one agreed to. Typically, a penalty should be imposed for such a deviation to ensure that C's welfare in maximized. Unfortunately, i's profile cannot be directly observed, only her realized consumption x_i. 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* payment dubbed a *separating payment*, that is designed to make the expected value of using the declared profile higher than the others. Given such a D_i, we translate it to ensure that its expected value is 0, thus preserving the convexity of the MPOU game.

We show that separating functions of a certain form can be found using a linear program. While we lack proof that such a form of separating functions always exists, we are always able to find one in our experiments.

4. EMPIRICAL RESULTS

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. [2].[2] The main question we study empirically is what is the overall social welfare gain from using an MPOU vs. a POU vs. a fixed-rate model. In the full paper,[3] we also study the sensitivity of the results to agents' choices of profiles and the effect of separating functions on the agents' payment uncertainties.

We define the *uncoordinated POU setting* as the scenario where agents are subject to a POU tariff and form the grand coalition, but each agent uses the profile that individually maximizes his or her net utility relative to that POU tariff. The average social welfare achieved in the uncoordinated POU setting is less than that of the fixed rate setting, underscoring 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.

We find that the social welfare gain that can be achieved by a POU tariff when agents coordinate optimally under the MPOU framework is moderate, around 3-5%, under our experimental assumptions. Since we lack real data for several aspects of the problem, the gain could be higher or lower in practice. We note that these experiments are the first to study the resulting social welfare gain from a POU tariff.

5. FUTURE WORK

Interesting future directions for POU/MPOU games remain. First, 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. Second, 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.

[1]Such profiles and values may be explicitly elicited or estimated using past consumption data (see Sec. 4).

[2]Publicly available at `http://www.pecanstreet.org`.

[3]Available at `http://bit.ly/mpou-games`.

REFERENCES

[1] S. Braithwait, D. Hansen, and M. O'Sheasy. Retail electricity pricing and rate design in evolving markets. *Edison Electric Institute*, pages 1–57, 2007.

[2] J. D. Rhodes, C. R. Upshaw, C. B. Harris, C. M. Meehan, D. A. Walling, P. A. Navrátil, A. L. Beck, K. Nagasawa, R. L. Fares, W. J. Cole, et al. Experimental and data collection methods for a large-scale smart grid deployment: Methods and first results. *Energy*, 65:462–471, 2014.

[3] G. M. Team. Electricity and gas supply market report. Technical Report 176/11, The Office of Gas and Electricity Markets (Ofgem), December 2011.

[4] M. Vinyals, V. Robu, A. Rogers, and N. R. Jennings. 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)*, pages 829–836, Paris, 2014.