In *Proceedings of the 30th Conference on Artificial Intelligence* (AAAI 2016), Phoenix, AZ, USA February 2016

Autonomous Electricity Trading using Time-Of-Use Tariffs in a Competitive Market

Daniel Urieli and **Peter Stone**
Dept. of Computer Science
The University of Texas at Austin
Austin, TX 78712 USA
{urieli,pstone}@cs.utexas.edu

Abstract

This paper studies the impact of Time-Of-Use (TOU) tariffs in a competitive electricity market place. Specifically, it focuses on the question of how should an autonomous broker agent optimize TOU tariffs in a competitive retail market, and what is the impact of such tariffs on the economy. We formalize the problem of TOU tariff optimization and propose an algorithm for approximating its solution. We extensively experiment with our algorithm in a large-scale, detailed electricity retail markets simulation of the Power Trading Agent Competition (Power TAC) and: 1) find that our algorithm results in 15% peak-demand reduction, 2) find that its peak-flattening results in greater profit and/or profit-share for the broker and allows it to win against the 1st and 2nd place brokers from the Power TAC 2014 finals, and 3) analyze several economic implications of using TOU tariffs in competitive retail markets.

1 Introduction

The smart electricity grid, known as the *smart-grid*, is expected to be a main enabler of sustainable, clean, efficient energy supply. One of the milestones in the smart-grid vision is "customer participation in power markets through demand-side-management" (U.S 2003). Demand-side management refers to adapting customer demand to supply conditions, and may be implemented using new power market structures. Due to the high cost of failure in the real world (Borenstein 2002), it is important to test new power market structures in simulation (Weidlich and Veit 2008). This is the focus of the Power Trading Agent Competition (Power TAC) (Ketter, Peters, and Collins 2013), and of this paper.

In Power TAC, autonomous broker agents compete to make profit in a large-scale, realistic power markets simulator. Since wholesale power markets are not designed for individual customer participation (Kirschen 2003), such brokers can represent customer populations, and make profit while reducing customer costs and stabilizing the electricity grid (Ketter, Peters, and Collins 2013). Power TAC's simulator models real-world markets (like ERCOT's) with components of future markets, such as autonomous agents optimizing customer consumption.

One of the primary goals of demand-side management (DSM) is *peak-flattening*, i.e. distributing consumption more evenly throughout the day (Schweppe, Daryanian, and Tabors 1989). *Time-Of-Use (TOU) tariffs* are energy selling contracts that were proposed for implementing DSM (Joskow and Tirole 2006; Kirschen 2003). TOU tariffs specify time-based pricing, in contrast to *fixed-rate* contracts that currently dominate power markets, and thus incentivize customers to adapt their consumption to reduce costs.

TOU tariffs in competitive markets present at least three important challenges. First, in monopolistic retail markets the surplus resulting from reduced wholesale energy costs directly benefits the monopoly and possibly the customers. In contrast, in competitive markets some or all of this surplus might benefit the competitors (even if they do not use TOU tariffs), since wholesale energy costs are typically a function of the *total* amount bought, due to the wholesale auction structure. As a result, brokers using fixed-rate tariffs can enjoy the reduced prices resulting from peak-flattening by another broker using TOU tariffs, while at the same time gaining market share from this TOU broker due to the extra discomfort that TOU tariffs incur on customers. Second, TOU tariffs may cause a *herding* phenomenon, where customers shift their consumption to low-price times, thus creating new peaks (Ramchurn et al. 2011). The third challenge is addressing the aforementioned challenges of flattening demand and increasing the broker's surplus in a tractable way in a complex, realistic, real-time environment.

This paper focuses on two questions. First, how should an autonomous broker optimize TOU tariffs that are both 1) attractive to customers in a competitive retail market with fixed-rate tariffs and 2) more profitable for the broker than the best fixed-rate tariffs? Second, what is the economic impact of TOU tariffs in a competitive market? This paper's primary contributions are:

- We formalize the problem of optimizing TOU tariffs in competitive markets, show that it is intractable, and propose an efficient optimization algorithm that approximates its solution. Our algorithm is fully implemented in our broker agent.
- Our algorithm leads to 15% peak-demand reduction in a complex, large-scale simulation of competitive power markets (Power TAC). To the best of our knowledge, our

work is the first to show that TOU can achieve the primary goal of peak-flattening in competitive markets in such a large-scale, realistic simulation.

- Our agent's peak-flattening results in greater profit and/or profit-share and allows it to beat fixed-rate brokers, specifically the 1st and 2nd place agents from the 2014 Power TAC finals, while reducing the energy costs of both its customers and its competitors' customers.
- Using extensive experimentation, we analyze several economic implications of using TOU in competitive retail markets. For instance, while TOU tariffs can induce *customer-herding*, our TOU broker prevented it by implicitly coordinating flattening through profit-maximizing tariffs. This underlines a potential benefit of employing autonomous TOU brokers in competitive power markets.

2 Testbed Domain: Power TAC

Our testbed domain is the Power Trading Agent Competition (Power TAC) simulation environment. Power TAC is an annual competition in which the competitors are autonomous brokers programmed by participants from around the world. Participants release their broker binaries after the competition, and this allows for running controlled experiments against state-of-the-art brokers. Power TAC models competitive retail and wholesale power markets in a smart-grid environment of a medium-sized city, with more than 50,000 simulated customers. Power TAC's customers are autonomous agents that optimize the electricity-costs and comfort of their human-owners (Reddy and Veloso 2012). The simulator simulates 60 days in discrete, one-hour time-slots. One simulation takes about two hours to complete.

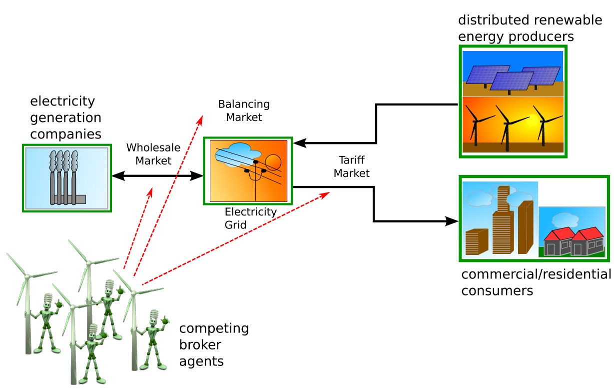

Figure 1: High-level structure of the Power TAC simulation

Figure 1 shows the structure of the Power TAC simulation environment. At a high level, autonomous broker agents compete with each other by acting in three markets: (1) a *wholesale market*, in which energy is traded with traditional generation companies, (2) a *tariff market*, which is a retail market in which energy is traded with consumers and distributed renewable energy producers, and (3) a *balancing market*, which serves to ensure that supply and demand are balanced at all times. More specifically, local customers such as office buildings, residential houses, and solar farms consume/produce energy according to real-world patterns, based on weather conditions and calendar factors such as day/hour. Power TAC uses state-of-the-art customer models, which consume/produce using time-series generators based on real-world data. Customers are equipped with smart-meters, which report consumption and production every hour. Autonomous brokers compete on gaining market share and maximizing profit by trading electricity. Brokers interact with local consumers and producers in the tariff market by publishing tariff contracts for energy consumption/production, which may include fixed and varying prices and possibly bonuses and/or fees. Customers subscribe to tariffs they find attractive. Brokers typically balance their portfolio's supply/demand by trading in the wholesale market. Full details can be found in The Power TAC Game Specification (Ketter et al. 2015).

3 Background - TOU Tariffs

A main motivator for demand-side management (DSM) is the variability in energy production prices. Figure 2 shows energy generation cost as a function of generated quantity in three large competitive wholesale power markets. The figure demonstrates how increased generation results in more sharply increasing costs. Typical daily customer demand has peaks, which thus result in high generation costs. One of the main goals of DSM is reducing these peaks by *flattening* customer demand throughout the day. This can reduce both generation costs, infrastructure costs, and CO_2 emissions.

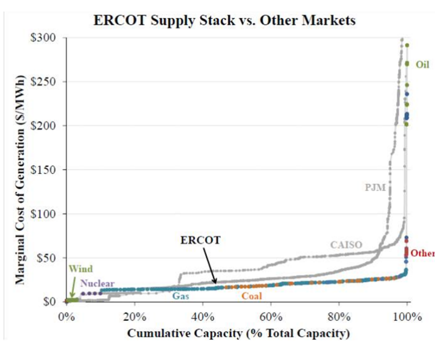

Figure 2: Generation-cost curves of three wholesale markets: ERCOT, PJM, CAISO. Source: (Newell et al. 2012)

One of the main methods proposed for implementing DSM is TOU tariffs, which specify different prices for different times of day. Here we define a TOU tariff τ to be a tuple $\tau := \langle p_0, p_1, \cdots, p_{23} \rangle$, where p_t is the energy price in cents/kWh during hour-of-day t. We refer to p_t as *hourly rate*. A TOU tariff with varying hourly rates incentivizes customers to adapt their consumption away from times of peak demand in order to reduce their energy costs. In Power TAC (Reddy and Veloso 2012) as in real-world markets (Albadi and El-Saadany 2008), when a tariff is published to the market, customers respond in 3 ways. The first two responses take place for any (fixed-rate or non-fixed-rate) tariff publication, while the third one takes place for non-fixed-rate tariffs, such as TOU, as follows:

1. *subscription changes*: a portion of the customer population may change their tariff subscriptions.

2. *consumption elasticity*: customers elastically adapt their *total* consumption based on prices.
3. *consumption shifting*: customers may shift consumption from expensive to cheap hours.

In competitive retail markets, TOU tariffs may need to compete with *fixed-rate* tariffs, which sell energy for a fixed price per unit. Fixed-rate tariffs do not affect customers' comfort, since customer payments are determined solely by the total energy consumed, regardless of *when* it is consumed. In contrast, under TOU tariffs customers face a trade-off between cost and comfort: to save costs, they may need to change their consumption patterns. Customers will subscribe to a TOU tariff and change consumption if the potential cost saving compared with competing fixed-rate tariffs is large enough to compensate for the extra discomfort.

Power TAC models this trade-off as follows (Reddy and Veloso 2012). A customer has a default *energy profile* e_H, which is a vector of desired consumption values up to some horizon H. Let $\bar{e}_H$ be a modified energy profile defined by some *admissible permutation* of e_H. Intuitively, an admissible permutation is a modified energy profile that satisfies the customer's constraints on how energy can be shifted, for instance not consuming below a customer's minimum required demand at any time, and shifting only portions of demand that are flexible. The discomfort implied by an admissible permutation $\bar{e}_H$ is quantified using a distance metric defined on profile vectors: $d\left(e_H, \bar{e}_H\right)$. Power TAC currently uses the L_2 distance metric $d\left(e_H, \bar{e}_H\right) := \sum_{t=1:H}\left(e_t - \bar{e}_t\right)^2$, and we find that it has desirable strategic effects, which we elaborate on later. Let $cost\left(\tau, \bar{e}_H\right)$ be the cost paid by a customer consuming energy according to $\bar{e}_H$ under a tariff τ. Let w be a constant weighting the importance of cost vs. discomfort. Then the customer's utility of subscribing to tariff τ and consuming according to $\bar{e}_H$ is $u_{cust}\left(\tau, \bar{e}_H\right) := -\left(cost\left(\tau, \bar{e}_H\right) + w \times d\left(e_H, \bar{e}_H\right)\right)$.

Customers optimize $\bar{e}_H$ to maximize their utility under a given tariff to which they are subscribed. This formulation assumes a baseline maximum utility of 0 corresponding to the customer using energy for free and consuming according to its desired energy profile. For a customer $cust$ subscribed to a tariff τ, the optimal consumption profile is $e^*_H := \arg\max_{\bar{e}_H} u_{cust}\left(\tau, \bar{e}_H\right)$, and the utility of tariff τ as (overloading notation) $u_{cust}\left(\tau\right) := u_{cust}\left(\tau, e^*_H\right)$. For any fixed-rate tariff τ_{fixed}, all permutations have the same price, so by the above definitions $e^*_H = e_H$, and $u_{cust}\left(\tau_{fixed}\right) := -cost\left(\tau_{fixed}, e_H\right)$. Therefore, for a given TOU tariff τ_{tou} and a fixed-rate tariff τ_{fixed}, the utility of τ_{tou} for a customer $cust$ is higher than that of τ_{fixed} (i.e. $u_{cust}\left(\tau_{tou}\right) > u_{cust}\left(\tau_{fixed}\right)$), when $-\left(cost\left(\tau_{tou}, e^*_H\right) + w \times d\left(e_H, e^*_H\right)\right) > -cost\left(\tau_{fixed}, e_H\right)$, i.e. when when it saves enough cost to overcome the extra discomfort.

4 Problem Formulation and Optimization

This section presents two of our main contributions: formalizing the TOU tariff optimization problem in competitive markets, and designing an algorithm for TOU optimization.

4.1 The TOU optimization problem

We frame the TOU optimization problem in terms of demand and cost curves. For simplicity, we start by considering a broker who has published no tariffs, and who is about to offer a candidate tariff τ with a single hourly-rate for a single hour in the future, trying to optimize its selling-price p. Let the *demand curve* of τ be the function $D : \mathbb{R} \to \mathbb{R}$ that maps energy-selling prices to the resulting energy demand from customers that will subscribe to τ. Typically, D is a decreasing function of price, since the higher the price, the lower the number of subscribers to τ as well as their consumption. Let the *unit-cost curve* be the function $C : \mathbb{R} \to \mathbb{R}$ that maps an energy amount to the unit-price for which the broker is able to procure it in the wholesale market. The energy unit-cost curve is typically an increasing function, due to the increasing generation costs illustrated in Figure 2. For a given selling price p, the profit, or *utility*, of the broker is $u(p) := D(p) \times p - D(p) \times C(D(p)) = D(p) \times (p - C(D(p))$, as illustrated in Figure 3. The optimal, profit-maximizing, selling price is

$$p^* := \arg\max_p u(p)) \tag{1}$$

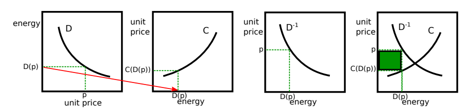

Figure 3: From left to right: (1) demand curve (2) unit-cost curve (3) inverse-demand curve (4) overlaying (2) and (3). The optimal selling-price p maximizes the rectangular area (profit).

Next, we generalize this example to tariff optimization for maximizing profit over some future horizon H. In this case there are potentially different demand and cost curves D_t and C_t for different future times $t = +1, \ldots, +H$ (using the notation '$+i$' to denote 'i timeslots into the future'). An optimal, profit-maximizing tariff might need to specify a vector of potentially different prices $P := (p_{+1}, \ldots, p_{+H})$. This can be implemented as a TOU tariff. The profit-maximizing price vector is $P^* := \arg\max_P \sum_{t=+1}^{+H} u_t(P_t)$, where $u_t(p)$ is defined similarly to u using D_t, C_t (D_t now depends on all prices in P due to subscription and shifting effects). In realistic scenarios, D_t and C_t are unknown and so is u_t, so an optimal price vector will be the one maximizing the predicted *expected* utility $P^* := \arg\max_P \sum_{t=+1}^{+H} E[u_t(P_t)]$.

Finally, we generalize to the case where the optimizing broker, as well as competing brokers, have possibly published tariffs in the market, so that publishing a new tariff may affect the demand for existing tariffs. Let $\mathcal{T}, \mathcal{T}'$ be the set of tariffs published by the optimizing broker and its opponents, respectively. Let $p^{(\tau)}_t$ be the hourly rate of some tariff τ for time t, let $P^{(\tau)} := (p^{(\tau)}_{+1}, \ldots, p^{(\tau)}_{+H})$, and let $\mathcal{P}^{(\mathcal{T})} := \{P^{(\tau)} | \tau \in \mathcal{T}\}$ the set of all vector prices of tariffs in the set $\mathcal{T}$. Let $D^{(\tau)}_t$ be the demand curve of τ at time t, given all other tariffs published in the market. The cost curve is not indexed by τ; it is a function of the total energy bought by all brokers. Using the notation $\mathcal{P}_t^{(\mathcal{T})}$ to stand for prices offered at time t by tariffs in a set $\mathcal{T}$, the utility at time t now becomes $u_t^{(\mathcal{T},\mathcal{T}')}(\mathcal{P}_t^{(\mathcal{T}\cup\mathcal{T}')}) = \left(\sum_{\tau\in\mathcal{T}} D_t^{(\tau)}(P^{(\tau)}) \times p_t^{(\tau)}\right) - \left(\sum_{\tau\in\mathcal{T}} D_t^{(\tau)}(P^{(\tau)})\right) \times C_t(\sum_{\tau\in\mathcal{T}\cup\mathcal{T}'} D_t^{(\tau)}(P^{(\tau)}))$ where the first term is income and the second term is costs (total demand times unit-cost). The utility of a candidate TOU tariff τ over a horizon H becomes $U_H(P^{(\tau)}) := \sum_{t=+1}^{+H} E[u_t^{(\mathcal{T}\cup\{\tau\},\mathcal{T}')}(\mathcal{P}_t^{(\mathcal{T}\cup\{\tau\}\cup\mathcal{T}')})]$. The TOU tariff optimization problem is defined in these terms as:

$$P^* := \arg\max_P U_H(P) \tag{2}$$

Note: this formulation can be extended with predicting opponent responses to publishing a new tariff. Due to the complexity of opponent modeling in our domain, we leave it as an avenue for future work, and focus on designing a best-response TOU optimization algorithm that approximates the solution to Equation 2. Our implementation compensates for the lack of opponent modeling by frequent replanning.

4.2 TOU Optimization Algorithm

Equation 2 provides a theoretical objective for solving an autonomous broker's TOU optimization problem. In practice, there are at least two main obstacles to solving this equation. First, future demand and cost-curves are typically unknown. Second, even if they were known, they would have a complex representation so that the optimal solution might not have a closed form and would be hard to find. We start with the second obstacle, assuming that the broker can predict the future demand and cost curves, and then we describe how curve predictions are implemented in our broker.

The second obstacle to coming up with a practical solution to Equation 2 is that of finding P^* in a non-convex, H-dimensional space (typically $H = 24$): effects such as subscription-changes and consumption-shifting create local maxima and discontinuities in $U_H(P)$, so that finding a global optimum is generally intractable. Therefore, we resort to finding a local optimum. In the simplified case of Equation 1 (in which $H = 1$), p^* may be approximated efficiently using 1-dimensional search over prices. However, it would be incorrect to apply such a method independently for each $i = +1, \ldots, +H$ for solving Equation 2. The reason is that changing a single price p_{+i} can create subscription-changes and consumption-shifting that affect all demand-curves in the horizon. In other words, hourly rates should be optimized in conjunction. Several well-known local optimization methods (Amoeba, BOBYQA, Powell's method) failed to achieve reasonable performance when approximating the solution to Equation 2. Therefore, we designed the following empirical gradient-based optimization algorithm for efficiently finding a local optimum to Equation 2.

Our algorithm both builds upon, and generalizes a fixed-rate tariff optimization algorithm proposed by (Urieli and Stone 2014). Their algorithm optimizes *fixed-rate* tariffs, and can be viewed as solving a limited version of Equation 2, using 1-dimensional search where P is only allowed to have identical entries. Another simplification of their algorithm compared to ours is that they estimated cost-curves as average-price, flat curves, while we estimate curves more accurately using online regression, since flat cost-curves would not correctly reflect the need for TOU tariffs. Algorithm 1 shows a pseudocode of our TOU optimization algorithm, which has two phases. In a preliminary first phase (line 1) it uses the fixed-rate optimization algorithm of Urieli and Stone to get a reasonable fixed-rate seed for phase 2 (line 2-11). Phase 2 optimizes a TOU tariff using empirical gradient-ascent. First, the fixed-rate seed is converted into a TOU tariff with H identical entries (line 2). Next (lines 3-5), each of the entries is perturbed by $\pm\epsilon$ and the utility $U_H(P)$ of the perturbed price vector is predicted using the predicted demand and cost curves (in our experiments $\epsilon = 0.5 cents/kWh$). Next, a 2-sided gradient is computed from the $2H$ perturbations and is normalized to a length of ϵ (line 6). Finally (lines 7-11), starting with the fixed-rate seed, the algorithm repeatedly takes steps in the direction of the gradient, as long as it has computation time and has not reached a local maximum. It returns the TOU tariff with the maximum predicted utility. To compensate for the lack of opponent modeling, Algorithm 1 is executed frequently by our broker, and thus responds to opponent actions. In our experiments, Algorithm 1 outperformed optimization algorithms like BOBYQA, Amoeba, and Powell's (these results were left out due to space constraints).

Algorithm 1 Gradient-Based TOU Tariff Optimization

1: fixedRateSeed $\leftarrow$ FindBestFixedRateTariff()
2: $\langle p, p, \cdots, p\rangle \leftarrow$ ConvertToTOUTariff(fixedRateSeed) //vector of length H
3: **for** i in $1, \ldots, H$ **do**
4: $\quad u_i^+ \leftarrow$ PredictUtility($\langle p, .., p, p+\epsilon, p, .., p\rangle$) //$\epsilon$ added to i'th entry
5: $\quad u_i^- \leftarrow$ PredictUtility($\langle p, .., p, p-\epsilon, p, .., p\rangle$) //$\epsilon$ subtracted from i'th entry
6: $\langle \epsilon_1, \epsilon_2, .., \epsilon_H\rangle \leftarrow$ NormalizeGradient($\langle \frac{u_1^+ - u_1^-}{2}, \frac{u_2^+ - u_2^-}{2}, .., \frac{u_H^+ - u_H^-}{2}\rangle$)
7: $P \leftarrow \langle p, p, \cdots, p\rangle$; currentUtility $\leftarrow$ PredictUtility(P); prev $\leftarrow$ null
8: **while** hasTime and notConverged(currentUtility, prev) **do**
9: $\quad P \leftarrow P + \langle \epsilon_1, \epsilon_2, \cdots, \epsilon_H\rangle$
10: $\quad$ prev $\leftarrow$ currentUtility; currentUtility $\leftarrow$ PredictUtility(P);
11: **return** P with highest utility

Back to the first obstacle, a broker agent needs to predict the future demand and cost curves $D_{+1}, \ldots, D_{+H}$ and $C_{+1}, \ldots, C_{+H}$. Our broker learns the cost curves C_{+i} from past data and adapts its predictions while trading as follows. First, it makes a simplifying assumption, that $C_{+1} = C_{+i}$ for $i = 1, \ldots, H$. Second, it uses L_2 regularized linear regression to create an initial curve estimation from past wholesale trading data. Third, at trading time our broker continually monitors its past prediction errors and adds a correction factor to future predictions, which is the average prediction error in the last 24 hours. This reduces the bias in future cost-curve predictions.

D_{+i} is determined by the combined effect of the three customer-behaviors of subscription-changes, consumption-elasticity and consumption-shifting, described in Section 3. Demand curve prediction can now be broken into three simpler (though still complex) prediction problems. If needed, each one of these behaviors could be learned from past data. In our case, in contrast to cost-curve which must be learned online (since they depend on competitor behaviors), the

Table 1: FixedRate, TOUNaive and TOU playing against AgentUDE (top table) and CWIBroker (bottom table).

	(a) Fixed-Rate-vs-UDE	(b) TOUNaive-vs-UDE	(c) TOU-vs-UDE	Change (c)/(a)
score: our-agent (M$)	**1.893**	**1.689**	**1.922**	1.016 **(+1.6%)**
score: UDE (M$)	**0.895**	**0.578**	**1.122**	1.253 **(+25.3%)**
market-share: our-agent (%)	64.0	73.3	61.4	0.959 **(-4.1%)**
(our) avg energy-buy price	0.053	0.051	0.051	0.963 **(-3.7%)**
(our) avg energy-sell price	0.105	0.098	0.105	1.000 **(-0.0%)**
(all) avg energy-buy price	0.051	0.049	0.049	0.961 **(-3.9%)**
(all) avg energy-sell price	0.105	0.099	0.104	0.990 **(-1.0%)**
peak-demand (MW)	**86.771**	**71.882**	**73.519**	0.847 **(-15.3%)**

	(a) Fixed-Rate-vs-CWI	(b) TOUNaive-vs-CWI	(c) TOU-vs-CWI	Change: (c)/(a)
score: our-agent (M$)	**0.677**	**0.524**	**0.622**	0.919 **(-8.1%)**
score: CWI (M$)	**0.771**	**0.620**	**0.558**	0.724 **(-27.6%)**
market-share: our-agent (%)	44.2	54.3	54.7	1.238 **(+23.8%)**
(our) avg energy-buy price	0.057	0.054	0.054	0.947 **(-5.3%)**
(our) avg energy-sell price	0.095	0.087	0.086	0.905 **(-9.5%)**
(all) avg energy-buy price	0.057	0.055	0.053	0.930 **(-7.0%)**
(all) avg energy-sell price	0.094	0.086	0.086	0.915 **(-8.5%)**
peak-demand (MW)	**86.701**	**74.720**	**73.651**	0.849 **(-15.1%)**

Table 2: Self-play of TOU, compared with TOU playing against AgentUDE, and against CWIBroker.

	(d) TOU-vs-TOU	Change (d)/(c) (UDE)	Change (d)/(c) (CWI)
score: our-agent (M$)	**0.493**	0.257 **(-74.3%)**	0.791 **(-20.9%)**
score: agent-copy (M$)	**0.482**	–	–
market-share: our-agent (%)	50.5	0.823 **(-17.7%)**	0.927 **(-7.3%)**
(our) avg energy-buy price	0.051	1.000 **(-0.0%)**	0.944 **(-5.6%)**
(our) avg energy-sell price	0.083	0.790 **(-21.0%)**	0.954 **(-4.6%)**
(all) avg energy-buy price	0.051	1.041 **(+4.1%)**	0.944 **(-5.6%)**
(all) avg energy-sell price	0.083	0.798 **(-20.2%)**	0.954 **(-4.6%)**
peak-demand (MW)	**70.101**	0.954 **(-4.6%)**	0.947 **(-5.3%)**

demand-curves are determined by stochastic models inside the simulator. Therefore, in this paper we equip the broker with the simulator's models of customer behaviors (though it still needs to estimate the models' parameters since they are drawn from a distribution at run time, and not reported to the broker; it does so by setting the parameters to their mean values). In the results section we ablate predictions modules and test their impact on performance.

5 Experimental Results

We evaluated our TOU broker using paired tests. We measured the impact of modifying a component of the broker by testing the original and the modified version in a set of games, in which the opponents and most random factors in the simulation were held fixed (random seeds, weather conditions). Paired testing improves our ability to evaluate the statistical significance of the results, by allowing us to use the Wilcoxon matched-pairs signed-ranks test instead of an unpaired test. To fix weather conditions, we used weather files containing 3 months of real-world weather. To cover year-round weather conditions we used 8 weather files (each file used by 1/8 of the games) with start-dates of January, April, July, October of 2009 and 2010. Opponents were fixed to be one of the top brokers played in the Power TAC 2014 finals: AgentUDE (Ozdemir and Unland 2015) (1st place), and CWIBroker (Hoogland and Poutre 2015) (2nd place).

5.1 Impact of Gradient-Based TOU on Broker's Performance and on the Economy

We tested how using TOU tariffs optimized with Algorithm 1 affected 1) the broker's performance, and 2) the economy. We compared a TOU Broker using Algorithm 1 with two variations: one that uses fixed-rate tariffs and another that uses a naive TOU tariff optimization. We refer to these brokers as TOU, FixedRate, and TOUNaive. FixedRate was created from TOU by disabling phase 2 of Algorithm 1, and using the fixed-rate tariff returned by phase 1 (line 1). TOUNaive was created from TOU as follows. Phase 2 of Algorithm 1 was replaced with a phase that naively assigns higher rates to hours with higher predicted costs, by adding a fixed margin to these predicted costs. Specifically, given a fixed-rate tariff with rate p returned by phase 1, and given a predicted cost vector $(c_{+1}, \ldots, c_{+H})$, the naive algorithm computes an average margin $m := \frac{1}{H}\sum_{i=+1}^{+H}(p - c_i)$, and publishes a TOU tariff with the price vector $P = (c_{+1} + m, ..., c_{+H} + m)$. All other broker components remained identical between the three brokers. We compared these three brokers in 2 different experiments, in which they played 200 games against (1) AgentUDE, and (2) CWIBroker, both of which use only fixed-rate tariffs.

Table 1 shows the results of these two experiments. Each row shows a measured quantity averaged over games played by FixedRate, TOUNaive and TOU, as well as the relative change in this quantity when using TOU instead of FixedRate. All results are statistically significant with $p = 0.01$ (many with $p \ll 0.01$), using the Wilcoxon matched-pairs signed-ranks test. TOU was the only agent that earned a higher score than both competitors (by 9% and 70% gaps against CWIBroker and AgentUDE respectively). TOUNaive was dominated by TOU in the sense that it made less profit against both opponents, and lost to CWIBroker.

Compared with FixedRate, TOU either earned more profit (against AgentUDE), or increased its profit-share from losing to winning (against CWIBroker, although with lower profit), while reducing peak demand by around 15%. TOU's peak reduction reduced the energy costs for both brokers and customers (including competitors'), and therefore increased social welfare. The surplus resulting from peak-reduction benefited either brokers or customers, depending on broker strategies. When playing against CWIBroker, customers enjoyed an 9.5% cost reduction, and brokers' suffered profit reduction, due to a fierce price-reduction competition. On the other hand, when playing against AgentUDE, brokers did not reduce prices as much; customers' cost reduction was only 1%, while brokers' profits increased. Since our broker plays a best-response strategy, the difference depends on how cooperative the other broker is.

Table 2 shows the results of running TOU against itself. It achieved the best flattening (around 20% peak reduction compared with FixedRate), and the lowest price for customers (around 5%-20% savings compared with when TOU played against AgentUDE, CWIBroker). However, in this case TOU achieved the lowest profit of all brokers due to a fierce price-reduction competition. TOU's best-response self-play benefited customers but not the broker. This illustrates game-theoretic issues pointed out by (Liefers, Hoogland, and Poutre 2014), whereby cooperative brokers could make higher profits, in this case by enjoying more of the surplus created by peak-reduction, at the expense of customers.

Figure 4 shows how the market power of a TOU broker affects its ability to flatten demand. All plots show consumption over 24 simulated hours. The left, middle, right columns show FixedRate, TOUNaive and TOU playing against CWIBroker. Peak demand is around 90MW, 80MW, and 70MW respectively (top row). TOU's large market share allowed it to counter-balance CWIBroker's customers'

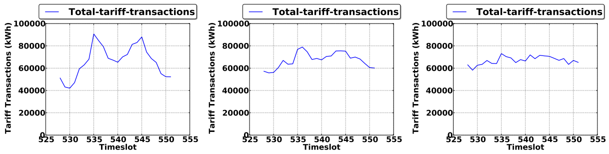

(a) Total consumption over 24 hours.

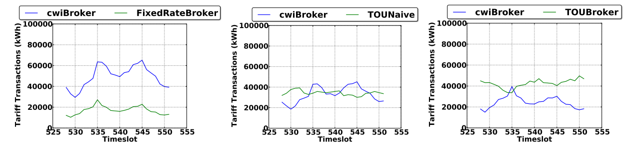

(b) Per-broker consumption over 24 hours.

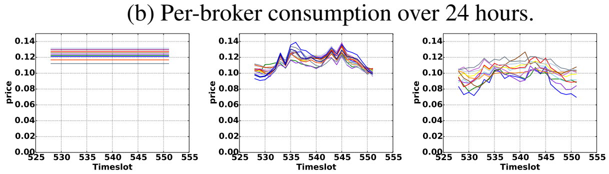

(c) Active tariffs over 24 hours.

Figure 4: Consumption flattening: FixedRate (left column), TOUNaive (middle), TOU (right).

Table 3: Ablation analysis: erratic-predictions

	TOU	NoShift	FlatCost
score: our-agent (M$)	0.622	0.507	-0.007
score: CWI (M$)	0.558	0.550	0.210
peak-demand (MW)	73.651	83.728	82.779

peaked-demand, while TOUNaive was only partially successful in doing so due to both lower market share (middle row), and suboptimal TOU pricing (bottom row).

TOU's frequent replanning using Algorithm 1 prevented customer herding (many customers shifting consumption to lowest-price times, causing a new peak (Ramchurn et al. 2011)). Even though we disabled customer-components for addressing herding (bundle-based optimization and stochastic shifting (Reddy and Veloso 2012)) and let customers shift greedily to their utility-maximizing energy profile, no herding was observed, due to a combination of (1) a TOU broker that implicitly coordinated flattening through profit-maximizing tariffs (Figure 4c, right), with (2) a smooth discomfort metric $dist\,(e_H, e^*_H)$. This underlines a potential benefit of employing TOU brokers in competitive markets.

5.2 Robustness of TOU to Prediction Errors

We tested the robustness of TOU to errors in its consumption-shifting predictions. Table 3 compares profits and peak-demand when testing 2 variations of TOU against CWIBroker. We chose CWIBroker as an opponent against which TOU had smaller profit margins (see Table 1), so accurate predictions seemed important. The left column shows the results of TOU, copied from Table 1, as a reference. The NoShift broker was created from TOU by disabling the consumption-shifting prediction module, and the FlatCost broker was created by adding noise to cost-prediction, making it predict a flatter cost-curve slope. Based on Table 3, both the consumption-shifting and cost prediction modules are critical for both profit and peak-flattening: without them TOU lost and peak-demand was barely reduced.

6 Related Work

DSM is viewed as an important component of future smart-grids. (Palensky and Dietrich 2011) provides a taxonomy for DSM. Existing work on TOU tariffs either has not considered competitive retail markets or has used more abstract, smaller-scale simulations (Yang, Tang, and Nehorai 2013; Wu, Wang, and Cheng 2004; Celebi and Fuller 2007; Datchanamoorthy et al. 2011; Veit et al. 2014; Albadi and El-Saadany 2008; Triki and Violi 2009; Yousefi, Moghaddam, and Majd 2011). To the best of our knowledge, this paper is the first to investigate the usage of TOU tariffs by autonomous brokers in a large-scale, detailed, realistic simulation of competitive power markets with autonomous customer agents. Recent work on Power TAC agents used fixed-rate tariffs (Ozdemir and Unland 2015; Hoogland and Poutre 2015; Urieli and Stone 2014; Liefers, Hoogland, and Poutre 2014; Babic and Podobnik 2014; Kuate, Chli, and Wang 2014; Kuate et al. 2013), except one (Ntagka, Chrysopoulos, and Mitkas 2014) who reports using TOU tariffs with 2 or 3 daily rates, however at that time Power TAC included only non-shifting customers, so that the impact of TOU tariffs in presence of demand-shifting customers could not be tested. In prior trading agent competitions, utility-optimization approaches were used in different domains (Pardoe 2011; Stone et al. 2003).

7 Conclusion

We formalized the problem of TOU tariff optimization in competitive retail markets, and proposed a real-time gradient-based, utility-optimization (profit-maximization) algorithm that approximates its solution. Our algorithm is fully implemented and tested extensively in the Power TAC simulator. Our gradient algorithm is currently the only TOU algorithm that performs robustly in Power TAC's complex, realistic environment: both a naive approach (TOUNaive) and well-known optimization algorithms failed to outperform fixed-rate brokers. We have shown that TOU tariffs can compete successfully with fixed-rate tariffs: our TOU agent outperformed the top 2 agents of the Power TAC 2014 finals, reduced peak-demand by 15% compared with using only fixed-rate tariffs, increased its profit and/or profit-share, and saved costs for all customers (including competitors').
Our ablation analysis showed the importance of having accurate customer-shifting and cost-curve predictions.

While TOU tariffs can induce *customer-herding*, our TOU broker prevented it by implicitly coordinating flattening through profit-maximizing tariffs. This underlines a potential benefit of employing autonomous TOU brokers in competitive power markets. In addition, we have seen that a TOU broker's customer share is an important factor in its ability to flatten demand: to counter-balance peaked consumption of fixed-rate brokers' customers, it needs to gain large customer-share by creating attractive TOU tariffs that are still profitable. Finally, our experiments demonstrated game-theoretic issues that affect the distribution of surplus created by reduced costs. An important direction for future work is exploring the market efficiency when many agents are competing against each other.

8 Acknowledgments

This work has taken place in the Learning Agents Research Group (LARG) at UT Austin. LARG research is supported in part by NSF (CNS-1330072, CNS-1305287), ONR (21C184-01), AFRL (FA8750-14-1-0070), and AFOSR (FA9550-14-1-0087).

References

Albadi, M., and El-Saadany, E. 2008. A summary of demand response in electricity markets. *Electric Power Systems Research* 78(11):1989 – 1996.

Babic, J., and Podobnik, V. 2014. Adaptive bidding for electricity wholesale markets in a smart grid. In *AAMAS Workshop on Agent-Mediated Electronic Commerce and Trading Agents Design and Analysis (AMEC/TADA 2014)*.

Borenstein, S. 2002. The trouble with electricity markets: Understanding california's restructuring disaster. *Journal of Economic Perspectives* 16(1):191–211.

Celebi, E., and Fuller, J. 2007. A model for efficient consumer pricing schemes in electricity markets. *Power Systems, IEEE Transactions on* 22(1):60–67.

Datchanamoorthy, S.; Kumar, S.; Ozturk, Y.; and Lee, G. 2011. Optimal time-of-use pricing for residential load control. In *Smart Grid Communications (SmartGridComm), 2011 IEEE International Conference on*, 375–380.

Hoogland, J., and Poutre, H. L. 2015. An effective broker for the power tac 2014. In *AAMAS Workshop on Agent-Mediated Electronic Commerce and Trading Agents Design and Analysis (AMEC/TADA 2015)*.

Joskow, P., and Tirole, J. 2006. Retail electricity competition. *The RAND Journal of Economics* 37(4):799–815.

Ketter, W.; Collins, J.; Reddy, P. P.; and Weerdt, M. D. 2015. The 2015 Power Trading Agent Competition. *ERIM Report Series Reference No. ERS-2015-001-LIS*.

Ketter, W.; Peters, M.; and Collins, J. 2013. Autonomous agents in future energy markets: The 2012 Power Trading Agent Competition. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*. AAAI.

Kirschen, D. 2003. Demand-side view of electricity markets. *Power Systems, IEEE Transactions on* 18(2):520–527.

Kuate, R. T.; He, M.; Chli, M.; and Wang, H. H. 2013. An intelligent broker agent for energy trading: An mdp approach. In *The 23rd International Joint Conference on Artificial Intelligence*.

Kuate, R. T.; Chli, M.; and Wang, H. H. 2014. Optimising market share and profit margin: Smdp-based tariff pricing under the smart grid paradigm. In *Innovative Smart Grid Technologies Conference Europe (ISGT-Europe), 2014 IEEE PES*, 1–6.

Liefers, B.; Hoogland, J.; and Poutre, H. L. 2014. A successful broker agent for power tac. In *AAMAS Workshop on Agent-Mediated Electronic Commerce and Trading Agents Design and Analysis (AMEC/TADA 2014)*.

Newell, S.; Spees, K.; Pfeifenberger, J.; Mudge, R.; DeLucia, M.; and Carlton, R. 2012. Ercot investment incentives and resource adequacy.

Ntagka, E.; Chrysopoulos, A.; and Mitkas, P. A. 2014. Designing tariffs in a competitive energy market using particle swarm optimization techniques. In *AAMAS Workshop on Agent-Mediated Electronic Commerce and Trading Agents Design and Analysis (AMEC/TADA 2014)*.

Ozdemir, S., and Unland, R. 2015. Agentude: The success story of the power tac 2014's champion. In *AAMAS Workshop on Agent-Mediated Electronic Commerce and Trading Agents Design and Analysis (AMEC/TADA 2015)*.

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

Pardoe, D. M. 2011. *Adaptive Trading Agent Strategies Using Market Experience*. Ph.D. Dissertation, The University of Texas at Austin.

Ramchurn, S.; Vytelingum, P.; Rogers, A.; and Jennings, N. 2011. Agent-based control for decentralised demand side management in the smart grid. In *The Tenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, 5–12.

Reddy, P. P., and Veloso, M. M. 2012. Factored Models for Multiscale Decision Making in Smart Grid Customers. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI-12)*.

Schweppe, F.; Daryanian, B.; and Tabors, R. 1989. Algorithms for a spot price responding residential load controller. *Power Systems, IEEE Transactions on* 4(2):507–516.

Stone, P.; Schapire, R. E.; Littman, M. L.; Csirik, J. A.; and McAllester, D. 2003. Decision-theoretic bidding based on learned density models in simultaneous, interacting auctions. *Journal of Artificial Intelligence Research* 19:209–242.

Triki, C., and Violi, A. 2009. Dynamic pricing of electricity in retail markets. *4OR* 7(1):21–36.

Urieli, D., and Stone, P. 2014. Tactex'13: A champion adaptive power trading agent. In *Proceedings of the Twenty-Eighth Conference on Artificial Intelligence (AAAI 2014)*.

U.S. Department of Energy. 2003. *"Grid 2030" A National Vision For Electricity's Second 100 Years*.

Veit, A.; Xu, Y.; Zheng, R.; Chakraborty, N.; and Sycara, K. 2014. Demand side energy management via multiagent coordination in consumer cooperatives. *Journal of Artificial Intelligence Research* 50:885–922.

Weidlich, A., and Veit, D. 2008. A critical survey of agent-based wholesale electricity market models. *Energy Economics* 30(4):1728 – 1759.

Wu, Q.; Wang, L.; and Cheng, H. 2004. Research of tou power price based on multi-objective optimization of dsm and costs of power consumers. In *Electric Utility Deregulation, Restructuring and Power Technologies, 2004. (DRPT 2004). Proceedings of the 2004 IEEE International Conference on*, volume 1, 343–348 Vol.1.

Yang, P.; Tang, G.; and Nehorai, A. 2013. A game-theoretic approach for optimal time-of-use electricity pricing. *Power Systems, IEEE Transactions on* 28(2):884–892.

Yousefi, S.; Moghaddam, M. P.; and Majd, V. J. 2011. Optimal real time pricing in an agent-based retail market using a comprehensive demand response model. *Energy* 36(9):5716 – 5727.