Assessing Fairness of Residential Dynamic Pricing for Electricity using Active Learning with Agent-based Simulation

Swapna Thorve
University of Virginia
Charlottesville, United States
st6ua@virginia.edu

Henning Mortveit
University of Virginia
Charlottesville, United States
henning.mortveit@virginia.edu

Anil Vullikanti
University of Virginia
Charlottesville, United States
vsakumar@virginia.edu

Madhav Marathe
University of Virginia
Charlottesville, United States
marathe@virginia.edu

Samarth Swarup
University of Virginia
Charlottesville, United States
swarup@virginia.edu

ABSTRACT

Extreme weather events and fast-paced adoption of green energy technologies have led to new challenges in demand-side management, maintaining grid reliability, and fulfilling variable consumer demands One of the effective ways to address these difficulties is by introducing economic incentives – replacing the flat rate tariffs with dynamic tariffs. However, dynamic pricing schemes need to be designed carefully to consider fairness and benefits for consumers as well as power companies. This paper describes an ML-based simulation framework for exploring two fairness constructs of dynamic pricing for residential electricity with behavioral agent-based models based on social theory combined with active learning. As an example, we simulate behavior adaptations in response to changes in electricity prices to study cost savings through monthly bills and peak demand reduction in synthetic household agents in a Time Of Use (TOU) pricing scheme in Virginia, USA. Further, we can show that there exists a region in the parameter space that corresponds to a fair TOU pricing scheme for both entities: all income-stratified communities and power companies.

KEYWORDS

fairness; active learning; agent behavior simulation; UN SDG; dynamic pricing; complex systems; smart grid

ACM Reference Format:
Swapna Thorve, Henning Mortveit, Anil Vullikanti, Madhav Marathe, and Samarth Swarup. 2024. Assessing Fairness of Residential Dynamic Pricing for Electricity using Active Learning with Agent-based Simulation. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 10 pages.

1 INTRODUCTION

Efforts for meeting climate change goals through the residential energy sector are increasingly focusing on incentivizing green energy adoption, decreasing GHG emissions, and conserving overall usage. In this context, the U.S. Energy Information Administration (EIA) and the U.S. Agency for International Development (USAID)[1] describe three popular demand-side energy use management strategies – peak shaving, load shifting, and promoting conservation/efficiency. One way of incentivizing the reduction and/or shifting of household energy use from peak demand times (e.g., evening hours) to non-peak hours is by varying the price of electricity at short intervals throughout the day. This type of pricing scheme is 'dynamic grid tariff' and can be beneficial to electricity providers[2] as well as the consumers. Power companies are interested in reducing the stress on the grid during peak times. They nudge households to alter their energy use patterns for reducing peak time demand by using demand-response strategies such as dynamic pricing [3, 6, 37].

Power companies and economists have been studying the importance and effects of rate/tariff design through the lens of economic & social theory [3, 33, 35] as well as field trials [1, 7, 17, 25]. Recent consumer trials study the effects of different types of dynamic pricing schemes such as real-time pricing (RTP; e.g., [10]), time of use (TOU; e.g., [37, 39]), and critical peak pricing (CPP; e.g., [21]). These are usually longitudinal experiments and support the analysis of demand changes throughout the day as well as document long-term benefits for the utilities in terms of demand reduction and costs. Findings from these trials have reported that many residential consumers achieve a reduction in peak-time energy use and sometimes observe a reduction in their monthly electricity bills. Such studies reveal that the potential for these savings varies by geography, household practices, occupancy patterns, weather variables, affordability, household demographic composition, and finally household demand elasticity [14, 17, 36, 46].

A few pricing trials have reported disparities in benefits after adopting dynamic pricing in different communities. For example, dynamic tariffs rendered vulnerable consumers unable to afford adequate cooling or heating of their homes [13, 42], thereby having adverse health consequences. Other instances show disproportionately increased bills for households with elderly and disabled occupants [8, 43]. Apart from income-related inequities [46], embracing dynamic tariffs has also predicted worse health outcomes for households with disabled and ethnic minority occupants [42].

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

[1] https://www.usaid.gov/energy/efficiency/basics
[2] We use the terms utility and power company interchangeably to represent electricity provider.

Thus, it is important to design a tariff that not only benefits power companies but also is fair to all communities in the region. In this paper, we propose a framework for designing a fair Time Of Use (TOU) dynamic pricing scheme with agent-based simulation of energy use behaviors and machine learning. Our contributions are outlined below:

- Our work lies at the intersection of three sets of research: ABM+ML, fairness, and energy economics. Two fairness metrics are proposed for designing a Time of Use (TOU) dynamic pricing scheme based on principles of tariff design, and economic and behavior theory. The first fairness criterion benefits household agents by either decreasing the monthly bills or keeping them the same as under a flat rate tariff. The second fairness constraint considers the economics of tariff design principles and satisfies the power company's main goal of reducing peak time demand.
- A novel framework based on active learning and agent-based simulations is presented to design a fair dynamic pricing scheme. The agent-based simulations play a two-fold role in the framework. First, it simulates elasticity in household behaviors in response to changes in peak price at the appliance level. Second, the simulation acts as an *oracle* (in the active learning setup) to create and annotate training data for the active learner (e.g., classifier). This setup efficiently navigates the parameter space to find a feasible region that represents a fair dynamic pricing region.
- Given a reasonable range of peak and non-peak prices, our most important result shows that it is possible to design a fair TOU pricing scheme such that (a) all households achieve the threshold peak-time energy demand reduction; (b) the pricing is fair to all population groups (based on income as a sensitive attribute) in terms of observing benefits on monthly energy bills.

Paper organization. We first provide background on electricity pricing principles and related literature. This is followed by our proposed framework, experiments, and results. We conclude with a discussion and avenues for future work.

2 BACKGROUND

This section will provide background on the principles of designing energy pricing, appliance scheduling techniques in the literature, and the synthetic energy use dataset used in this paper.

2.1 Residential energy pricing

Seminal work on rate design goals was proposed by Bonbright [5] in 1960 and was further expanded by Public Utility Regulatory Policies Act of 1978 (PURPA) and American Council for an Energy-Efficient Economy (ACEEE) [3]. There exist many competing policy objectives in designing residential pricing w.r.t. two entities: utilities and consumers. Many of these principles focus on fairness of rate design and promoting energy efficiency – (i) fair return and revenue stability to maintain grid health and utility profits; (ii) rate recovery is evenly distributed among all customer classes; (iii) rates will be designed to discourage wasteful use of public utility services; (iv) rates should be easy to understand and respond to; (v) revenue stability to maintain grid health. Further explanation is provided in Appendix B.

Recently, *fairness* has gained importance in discussions about energy pricing as a rate design principle. ACEEE's preliminary analyses of dynamic pricing trials have found that adopting time-varying rates, specifically, TOU rate design with Critical Peak Pricing (CPP) or Peak Time Rebate (PTR) shows the greatest promise of satisfying the principles of rate design described above. While this is true, the fairness of these schemes still remains an open question given that some trials reported discrimination and adverse effects on consumers. Apart from economic and social theory works in the fairness of dynamic tariffs [3, 22, 24, 33, 35], there is limited data-driven AI/ML literature on the design of fair dynamic tariffs for the residential sector (e.g., [2, 27]). We summarize three essential qualities that a fair dynamic tariff should possess: (i) Fluctuations in tariffs should be transparent and easy to understand; (ii) Reference dependency and fairness of distribution – transition to the new tariff should be fair to all consumers and the new cost should be the same or less than the initial cost; (iii) Appropriate controls should be in place to prevent consumer exploitation after transition. We use this background knowledge in identifying fairness constructs in our work.

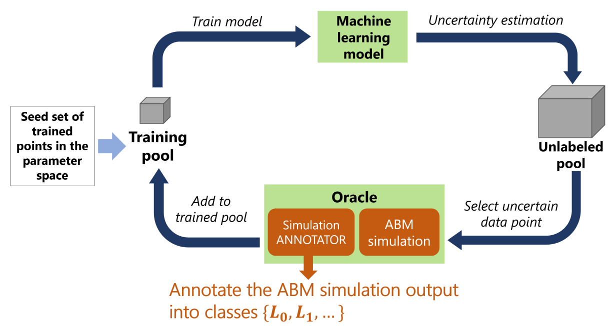

Figure 1: Conceptual overview - Active learning framework for learning fairness decision boundaries in residential dynamic pricing using ABM simulation.

2.2 Methods in literature

While much of the work in energy economics and public policy focuses on rate design, statistical optimization techniques have been employed for optimal scheduling of household appliances under dynamic tariffs [23]. They optimize competing objectives such as maximizing utilities' profits vs. minimizing users' costs under dynamic tariffs. A few game-theoretic approaches have modeled appliance scheduling under dynamic tariffs as a game between utilities and consumers [9, 11, 20, 32, 45]. However, most of these methods do not guarantee fairness.

Other works develop fair allocation strategies for solving different social good problems [26, 29, 30, 34]. For example, Oluwasuji et al. [34] develop a suite of optimization models for studying the fair allocation of scarce resources in agents in the context of enabling fairer and more efficient load shedding. They evaluate the results with utilitarian, egalitarian, and envy-freeness social welfare metrics which are standard in the fair allocation literature.

A broader class of fairness metrics such as demographic parity, equalized odds, disparate impact, and others have been used in many applications as fairness constraints [15, 16, 28, 31]. For example, Biswas et al. [4] use a combination of machine learning and fairness, to propose "proportional equality" as a fairness metric to evaluate the fairness of classifiers employed in applications involving societal decisions (e.g., hiring, loan allocation).

When considering ABM/ML and fairness, most work is focused on making ABMs fairer [12, 44]. They emphasize that fairness-related representations have not received due attention in ABMs. We consider our work here a step in this direction. We combine fairness and energy economics to develop our constraints. Our fairness metrics have been defined based on the behavior theory construct of 'reference dependency' and rate design principle of 'promotion of energy conservation and energy efficiency'. Our goal with this is to bring fairness constructs into computational modeling. We further integrate the ABM with machine learning (active/semi-supervised learning) to scale the entire framework to work with large-scale high-resolution agents and save on computation times.

Table 1: Entities and attributes used in the simulation

Entities	Attributes
Household agent	*Intrinsic attributes*: location, income, occupancy. *Energy-related attributes*: presence of major appliances, energy use (kWh), energy use bill ($) *Derived attributes*: flexibility factor, appliance shift flexibility
Energy data	Hourly appliance-level energy use time-series for each agent
Price signal	Peak price, non-peak price

3 PROBLEM DESCRIPTION

Let there be a valid pricing range for residential electricity. In this instance, the pricing range is defined by a 2D parameter space of peak and non-peak pricing and is considered as the region of interest. Let c_{peak} and c_{npeak} be an instance of peak and non-peak prices defined under the TOU pricing strategy. Our goal is to find a feasible region in this parameter space that is fair to all agents – consumers and electricity providers. Due to complex agent behaviors and dependence on exogenous variables, it is hard to obtain an analytical form of the *fair* pricing region. But, it is possible to **query** the parameter space and **evaluate** individual points in the 2D space. Such points can be assessed by a function $F(\Theta)$ and $c_{\text{peak}}, c_{\text{npeak}} \in \Theta$. Let y be the output of F. y can be evaluated to represent if the TOU price vector was a fair price or not. If y can be labeled, then, we can treat it as a classification problem where one can learn the decision boundary in the pricing parameter space that partitions it into fair and unfair regions. Next, we will describe our particular problem instance and its proposed solution.

Let $\mathscr{H}$ be the set of households serviced by a power company. Let $\tilde{c}$ be the flat rate tariff ($/kWh) indicating that the price is the same for all T time periods. For household $i \in \mathscr{H}$, let $\tilde{e}_i$ be the energy (kWh) used and $\tilde{b}_i$ be the monthly energy use bill ($) such that $\tilde{b}_i = \tilde{e}_i \times \tilde{c}$. An agent i responds to a TOU vector $[c_{\text{peak}}, c_{\text{npeak}}]$ by changing appliance use behavior subject to household constraints. The resultant peak energy use ($e_{i,\text{peak}}$), non-peak demand ($e_{i,\text{npeak}}$), and monthly energy bill (b_i) are computed for every agent i. Note that each i responds differently to a given pricing vector.

Our goal is to find a range of peak and non-peak prices that are **fair to all agents** (in terms of monthly bills) under TOU pricing and that also satisfy the power company's goal of **peak-time demand reduction**. The feasible region for TOU prices will satisfy the following constraints –

$$c_{\text{peak}} > 0, c_{\text{npeak}} > 0 \tag{1}$$

$$c_{\text{peak}} > c_{\text{npeak}} \tag{2}$$

$$c_{\text{npeak}} \leq \tilde{c} \tag{3}$$

$$\sum_{i=1}^{n} \tilde{e}_{i,\text{peak}} - \sum_{i=1}^{n} e_{i,\text{peak}} \geq E \tag{4}$$

$$(e_{i,\text{peak}} \times c_{\text{peak}} + e_{i,\text{npeak}} \times c_{\text{npeak}}) \leq \tilde{e}_i \times \tilde{c} \tag{5}$$

where E is the total energy reduction in peak time demand expected by the power company. Inequalities 4 – 5 are examples of fairness constraints for the TOU scheme. They are complex to model in an analytical form since they *strongly depend on environmental factors as well as agent's affordability and behaviors*. Examples of environmental factors are building stock attributes or weather conditions such as irradiance, temperature. One example of consumer routines related to appliance use and comfort preferences is while using heating/air conditioner systems, since it depends upon efficiency of the equipment, outside temperature, does the agent have a habit of changing thermostat setting when at home and not at home. All of these play an important role in determining if and when they intend to perform energy-consuming activities. Similar behaviors apply to other major appliance use. These constraints have been modeled as an ABM simulation in this work.

Let the ABM simulation be described as a stochastic function $F(\Theta)$ where Θ is the set of k parameters. $c_{\text{peak}}, c_{\text{npeak}} \in \Theta$ will be one of the input parameters to the simulation. y is the output of the simulation for a setting of Θ. Each agent i responds to a setting of Θ through the behavior models. At every iteration, the simulation computes y and classifies it into two classes: *fair* or *unfair*, corresponding to pre-defined fairness constructs. The representation of y for this work is described in the fairness criteria definitions below. We attempt to iteratively characterize the behavior of the ABM in terms of the probability of seeing a fair output with different Θ settings. In this work, we define two classes L_0 indicating the ABM output lies in the infeasible region or is unfair, and L_1 suggesting the ABM output lies in the feasible region (i.e., fair).

This setup makes it suitable to treat this as a classification problem and find a decision boundary that separates the fair pricing region from the unfair pricing region in the 2D space of c_{peak} and c_{npeak} parameters in terms of the ABM output y such that it satisfies a set of constraint(s) listed in Equations 1–5. To navigate this unlabeled parameter space efficiently and reduce the number of expensive simulation runs, we use active learning. The general idea is to train a multi-class classifier on the data points labeled by the ABM simulation. The classifiers estimate the area under each class, thus giving us the functional representation of the ABM outputs over the parameter space. Thus, our ABM simulation is plugged as an *oracle* in the active learning framework.

Once the ABM output is annotated, the active learning process trains a binary classifier with the updated training data in each iteration. Then, we use the uncertainty margin sampling function to find the next most indecisive price point. The selected point generates new parameter settings on which the simulation is evaluated in the next iteration. Once termination condition is reached, the classifier has successfully learned the *fairness* boundary in the parameter space. Figure 1 provides a conceptual overview of the problem setup.

Fairness criterion 1. This criterion is defined based on the behavior theory construct of 'Reference Dependency' (described in Background). A TOU price vector is fair only if the new monthly bill b_i is the same or less than the baseline monthly bill $\tilde{b}_i$ under the flat rate tariff with no behavior change. In the proposed active learning framework, the *oracle* is our ABM simulation. We define two bins/classes L_0 and L_1 for the active learner. The last step of the ABM simulation is to analyze and annotate the output. The output lies in bin L_0 when the price vector (selected by active learning) induces an unfair outcome and in L_1 when the price vector induces a fair outcome. Let the simulation run for n agents for a TOU price vector $[c_{\text{peak}}, c_{\text{npeak}}]$. Then, the label of the simulation output is decided by the following equation

$$B = \frac{\sum_{i=1}^{n} b_i}{n}, \qquad \tilde{B} = \frac{\sum_{i=1}^{n} \tilde{b}_i}{n}$$
$$\text{class} = \begin{cases} L_1, & \text{if } \tilde{B} - B \geq 0 \\ L_0, & \text{otherwise} \end{cases} \tag{6}$$

where e_i and b_i are calculated as follows –

$$\begin{aligned} e_i &= e_{i,\text{peak}} + e_{i,\text{npeak}} \\ b_i &= e_{i,\text{peak}} \times c_{\text{peak}} + e_{i,\text{npeak}} \times c_{\text{npeak}} \end{aligned} \tag{7}$$

Fairness criterion 2. This criterion is defined in reference to a rate design principles i.e. 'promotion of energy conservation and energy efficiency'. In our case, we quantify this as a minimum average amount of peak demand reduction achieved by agents for a TOU price vector to be fair (from the power company's perspective). Let the average minimum peak demand reduction (in kWh) be Δ_{peak} across n agents. Two classes L_0 and L_1 as defined similar to fairness criterion 1. The ABM output is categorized as unfair when the average minimum peak demand reduction is not achieved (thus, unfair from the utility's perspective), thus labeled as L_0. If the minimum peak demand reduction is achieved, then the oracle labels this point as fair and is assigned to L_1. Let the simulation run for n households for a TOU price vector $[c_{\text{peak}}, c_{\text{npeak}}]$. The label for the simulation output is computed using the following condition

$$E_{\text{peak}} = \frac{\sum_{i=1}^{n} e_{i,\text{peak}}}{n}, \qquad \tilde{E}_{\text{peak}} = \frac{\sum_{i=1}^{n} \tilde{e}_{i,\text{peak}}}{n}$$
$$\text{class} = \begin{cases} L_1, & \text{if } \tilde{E}_{\text{peak}} - E_{\text{peak}} \geq \Delta_{\text{peak}} \\ L_0, & \text{otherwise} \end{cases} \tag{8}$$

For evaluating fairness criteria 1, y is defined in terms of B and then classified as shown in Eq 6. For evaluating fairness criteria 2, y is computed as E_{peak} and then labeled by the condition in Eq 8.

4 SIMULATION FRAMEWORK

We describe our simulation using the general skeleton of the ODD protocol [19]. The purpose of the simulation is to compute the changes in energy use by households due to the given price signal, and the consequent change in the energy bill. An ML wrapper for the simulation is then used to find a range of fair prices (in terms of peak and non-peak prices) for residential electricity consumers.

4.1 Entities

The primary agent in the simulation is a household denoted by $i \in \mathcal{H}$. We use a synthetic representation of households from an openly available digital twin of the U.S. population in Gallagher et al [18] and supporting data from [40]. The entities and attributes in the simulation are listed in Table 1. In the following, the household agent's energy use (in kWh) is denoted as $\tilde{e}_i$ and the agent's monthly energy use bill (in \$) as $\tilde{b}_i$. The household agent interacts with an external factor – the TOU pricing variable which is defined as a vector $[c_{\text{peak}}, c_{\text{npeak}}]$. TOU prices are typically defined on an hourly basis. Prices are typically lower early in the day, overnight but are higher during peak evening hours. An example of TOU implementation by Southern California Edison is illustrated in the Appendix. Section 5 describes the setup of TOU for our experiments.

4.2 Inputs and initialization

We consider a valid price range for peak and non-peak prices for the region under consideration. Non-peak price ranges from 0-0.11 \$/kWh. Peak price ranges from 0-1.5 \$/kWh. The flat rate tariff for the region under consideration is 0.11 \$/kWh. Note that behavior attributes such as the flexibility factors of individual agents are derived from occupancy pattern, affordability, and their interactions with the TOU vector with the goals of reducing peak time consumption and overall monthly costs.

Second, the preference data for appliance shifting is obtained from a survey conducted by Stelmach et al. [39] for over 337 households in a Northern California city slated for TOU rates. The survey analyses show that agents' "willingness-to-shift" is highest for activities such as dishwasher, and laundry activities followed by EV charging and personal grooming activities such as showering. Further, the survey illustrates that agents' are hesitant to shift peak-time activities such as cooking, resetting thermostat controls, and watching TV. We quantify this data for "willingness-to-shift" out of peak time given a 30% price increase for all activities in the survey and use it as a simulation input. This data supports calculating flexibility factors for individual agents in our work. Note that this data is already made available by Stelmach et al. [39].

Digital twin of household-level energy use: The final major input to the simulation is the detailed energy use data at the household appliance level for 24 hours. Energy demand literature has shown that there is a general lack of large-scale high-resolution comprehensive datasets on energy use in households mainly due to consumer privacy and lack of infrastructure. In this work, we use openly available synthetic energy data and code provided by [40]. This digital twin is generated using data-centric AI methods and synthetic populations [18]. A mixture of stochastic, machine learning, physics-based engineering methods are used to model the different end-uses. It comprises hourly energy demand profiles for multiple appliances at the household level – heating/air-conditioning (HVAC), lighting, dishwashing, cooktops/oven, clothes washer and dryer, refrigerator, vacuum, computers, and TV. Household-level metadata (e.g., income) is also available. The published data is shown to be representative of real-world energy use data.

Thanks to the availability of detailed datasets that preserve consumer privacy, we now have the ability to combine Stelamch et al. [39] survey conclusions with the digital twin [40]. This is an important step towards micro-level policy recommendations for social good problems supported by large-scale simulations and ML.

Initialization of the simulation: We run the simulation for flat rate tariff with no behavior changes, which we call the Business as Usual (BAU) scenario. In addition, we also execute four initial simulations to provide a seed set to the active learning algorithm.

4.3 Process overview

While Figure 1 provides a conceptual overview of our framework, Figure 2 provides details of the simulation process and the iterative steps. Remember that the goal of the simulation is to find a range of fair prices in a 2D parameter space of peak and non-peak prices for residential electricity consumers. The simulation takes the inputs of energy data and a TOU price signal. Each agent responds to the TOU price signal and generates a new energy demand profile. If the collective simulation output for all agents satisfies fairness constraints, we label the TOU price signal as 'fair' or 'unfair'. One can observe that this particular problem instance can be formulated as a classification problem. Thus, we can learn a decision boundary in the 2D pricing parameter space that can separate fair and unfair prices. Also, note that such social simulations take a long time to execute. Thus, it will take a long time to find all possible fair and unfair pricing signals if the parameter space is not navigated intelligently. Therefore, we employ active learning to facilitate efficient navigation of the parameter space and use the simulation as an 'oracle' in the active learning process, as follows (Fig. 2).

1. Create a small "seed-set" of simulation runs to initialize the active learning process.
2. Label the simulation output and add it to the training set.
3. Train the active learning model on the simulation outputs.
4. Choose the next TOU pricing signal using the active learning acquisition function.
5. Execute the simulation with the new pricing signal as input.
6. Compute the flexibility factor for each agent that determines the household's ability to shift appliances' use outside peak hours. (Refer to models described in Section 4.4).
7. Execute the individual agents' decision for shifting each peak-time appliance use occurrence to non-peak hours subject to household constraints.
8. Generate the new energy demand profiles for all agents.
9. Calculate monthly bills and peak time metrics for agents.
10. Go to step 2.

4.4 Simulation design

In order to find the decision boundary that transitions from fair to unfair pricing, we propose an active learning based framework that embeds computational agent-based simulation to model agent behaviors for shifting activities from peak hours to non-peak hours and to navigate the parameter space with ease by minimizing the number of executions of expensive simulations. This is made possible by using a high-resolution digital twin of household appliance energy data from the literature [40]. Agent-based models (ABM) and detailed synthetic data provide us the flexibility to represent a household as an independent agent and model their inclination towards shifting individual appliance demands outside of peak hours. This also gives us the freedom to design fairness constraints at agent-level. Figure 2 describes the detailed steps for learning the fairness boundary in the TOU pricing parameter space[3].

Let $c_{\text{peak}}, c_{\text{npeak}}$ be the price vector selected by active learning that is input to the simulation. The simulation engine models agent behaviors in response to a price vector and household constraints. First, a *flexibility factor* is computed for each household i that quantifies the agent's ability to shift peak-time activities outside peak hours. Next, based on an activity shift priority (adopted from Stelmach et al. [39]) and the TOU vector, we calculate the *probability of shifting* an activity from peak to non-peak hours. Depending upon the occupancy schedule (at 15-minute intervals), *flexibility factor*, the presence of smart technologies, household (i.e. agent) preferences, and *probability of shifting* activity, the new activity schedule is generated by our appliance scheduling model. Then, the energy models simulate new hourly residential energy demand profiles. Thus, the simulation generates new energy demand profiles with appliance shifts and peak demands and updated monthly energy bills. The classifier is trained on labeled data generated at the r^{th} iteration of active learning. Here, we use a random forest classifier. Active learning uses the updated classifier to select the most informative price point using uncertainty sampling acquisition function (e.g., "smallest-margin" uncertainty sampling [38]). This process is run a fixed number of times or until there is no change in the learned boundary between two rounds.

4.4.1 Agent flexibility - ability to shift peak activities.

Many external variables (e.g., temperature, irradiance), household behaviors (e.g., cooking every day at 5 pm), building characteristics, demographics, and socioeconomic indicators determine how energy is used in a household. Nudging agents to modify when and how to use electricity in response to price fluctuations is challenging because all households may not respond to the change to the same degree. In addition, it may require negotiating with the agent and/or frequently compromising household practices. Thus, we can say that an agent's flexibility to adapt to a price change is contingent upon the rate of change from flat rate tariff, and the variables mentioned above. There have been limited attempts at quantification of how much people will adapt their behaviors to TOU pricing, e.g., [39]. We define a simple model based on income and monthly bill (with no behavior change) to quantify the flexibility of a household to a price vector to move peak activities to non-peak hours in terms of an *ability to shift* factor s_i, given as

$$s_i = \frac{1}{1+e^{z_i}} \quad \text{where} \quad z_i = \frac{v_i + \overline{b}_i}{100} - 1. \tag{9}$$

[3] The peak demand reduction graphic in Figure 2 is taken from Omnes Energy blog:http://www.omnesenergy.com/blog/2016/8/18/peak-demand-reduction-with-energy-storage-1

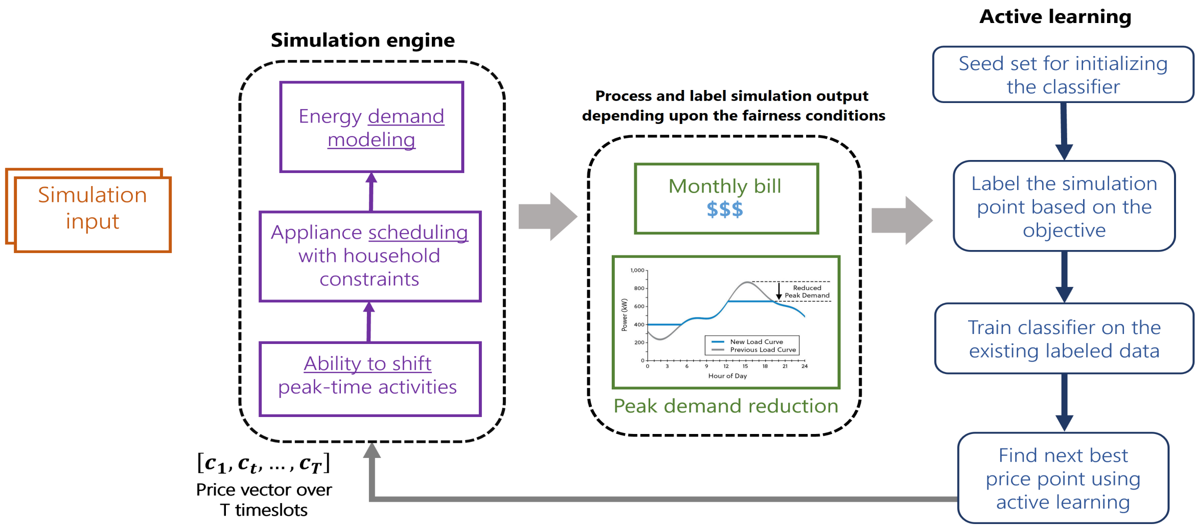

Figure 2: Framework for learning fairness in dynamic pricing (specifically TOU) using household-level behavior-induced agent-based modeling and active learning. Two objectives are considered for exploring the fairness of Time Of Use pricing in LMI and non-LMI communities: savings through the monthly bill and peak demand reduction.

v_i is the income percentile and $\bar{b}_i$ is the monthly bill of h_i for the TOU price vector $[c_{\text{peak}}, c_{\text{npeak}}]$ with no behavior change.

Let $\mathscr{A}$ be the set of activities/appliances observed during the peak period. The peak activities of interest are cooking, showering/bath, dishwasher, laundry, heating/cooling, vacuuming, lights, and device use such as TV. Of these peak activities, agents place the highest preference for shifting dishwasher and laundry activities [33, 39] during peak time. The *preference order* and *probability of shifting* an appliance outside of peak hours are adapted from Stelmach et al. [39]. These probabilities are recorded for a 30% increase in peak price. We adjust the probabilities to reflect changes in peak price when a new TOU price point is selected by active learning. This information is used in scheduling appliances. Let the probability of shifting an appliance/activity $a_j \in \mathscr{A}$ outside peak hours for the TOU price vector $[c_{\text{peak}}, c_{\text{npeak}}]$ be $\mathbb{P}(a_j)$.

$$\mathbb{P}(a_j) = \frac{c_{\text{peak}} \times \overline{\mathbb{P}(a_j)}}{\tilde{c}} \tag{10}$$

$\overline{\mathbb{P}(a_j)}$ is the probability of shifting appliance a_j at flat rate $\tilde{c}$.

4.4.2 Appliance scheduling. In this section, we focus on how agents respond to TOU pricing via *load-shifting* and *load reduction* strategies. This implies that appliances used in peak time are scheduled outside of peak hours for the same day based on the *probability of shift* for a TOU price vector. A data-driven behavior change algorithm is designed for scheduling appliances in a household for a TOU price vector chosen by active learning.

First, a 15-minute interval household occupancy sequence is constructed. The occupancy sequence records 3 states for each individual in the synthetic household for each 15-minute interval of the day. The recorded occupant states are *away*, *awake - at home*, and *asleep - at home*. This is easily constructed using the synthetic energy dataset. Next, we extract the existing schedule of appliances/activities in the households from the synthetic energy use data. Then, the probability of shifting each of the peak activities of interest is re-calibrated for a new price point based on the data found in [39] and Equation 10. Based on literature referred to in the Background section, some dynamic pricing trials have reported that households equipped with smart technology such as smart thermostats/appliances may be more responsive to dynamic pricing signals. Thus, we take into account the presence of smart technology (e.g., smart thermostats) in synthetic households. In this work, we focus on shifting/reducing energy demand from appliances such as dishwasher, laundry appliances (washer+dryer), heater/AC (HVAC), lighting, and cooking.

For every appliance, behavior change rules are defined based on the existing dynamic pricing trial literature [39]. The agent adopts behavior change only if the probability condition is satisfied. Dishwasher and laundry activities are scheduled outside peak hours only when the occupants are in the house and awake. However, if these appliances are smart technology enabled, the house can schedule dishwasher and laundry activities anytime outside peak hours. If HVAC is indicated as a peak activity, then, the occupants change the indoor thermostat setting by 2°F depending upon the season (e.g., the thermostat setting will increase by 2°F in summer) to reduce HVAC energy demand in peak hours. Similarly, if the lighting is indicated in peak activities, then, the household turns off any 2 bulbs during the peak period to reduce the peak time consumption. If a cooking activity needs to be shifted outside peak hours, then, it is shifted to either 1 hour before peak time or 2 hours after peak time as long as the occupants are at home and awake. If any electronic devices need to be shifted outside peak hours, then, these activities are randomly shifted to other timeslots when the occupant is at home. The energy demand modeling component computes the new energy demand schedules.

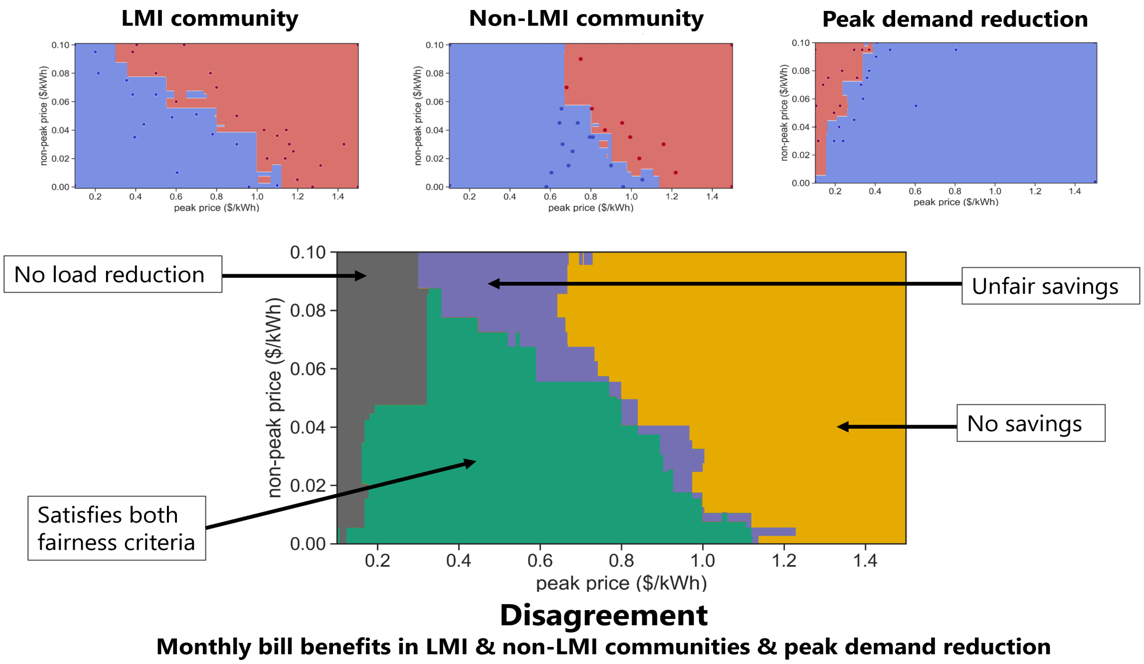

Figure 3: Summary of results from the three experiments. The top row represents three decision boundaries learned by active learning based on the two fairness criteria across income-segregated agents. The plot in the second row represents disagreement in the parameter space across the three experiments. The region that satisfies both the fairness criteria is denoted by green. The other (unfair) regions in the parameter space are denoted by different colors and a small caption beside them. Thus, if the power company were to design a fair TOU pricing for Rappahannock such that it achieves the utility's goal of peak demand reduction and fair distribution of monthly bills for the consumers, then, it would have to be a pricing point in the green region.

5 RESULTS AND DISCUSSION

Three experiments are conducted to explore fairness in designing a TOU scheme. The goal of the experiments is to learn three distinct decision boundaries in the peak and non-peak pricing 2-D parameter space that represents a TOU pricing scheme. The experiments are set up for the Rappahannock region in the state of Virginia, USA. It has approximately 3700 households. They are divided into two groups based on the area median income (AMI) of Virginia – LMI and non-LMI. LMI stands for Low-to-moderate income and is defined as 80% of AMI. For Virginia, the LMI threshold is a household annual income of $60,000. Approx. 600 out of 3700 households in Rappahannock are categorized as LMI. The remaining are categorized as non-LMI. The *Business As Usual* (BAU) scenario indicates that the agents do not change their appliance use behavior and the pricing is set to a flat-rate pricing scheme. In this setup, the flat rate is 0.11 $/kWh. Peak time is enforced between 5pm to 8pm.

Experiment 1. This experiment learns the decision boundary based on reference dependency and fair transition principles to discern where the monthly bill increases for agents in the LMI community when compared to the flat-rate pricing monthly bill without any behavior change. The decision boundary is based on fairness criterion 1.

Experiment 2. This experiment finds a decision boundary similar to the first experiment (i.e., based on fairness criterion 1), but for agents in the non-LMI community.

Experiment 3. This experiment examines the benefit for the power company in terms of peak demand reduction in the TOU pricing parameter space. A decision boundary is learned to separate the parameter space into 2 parts: the unfair region where the average peak demand reduction across agents is ≤ 1 kWh and the fair region corresponds to the average peak reduction ≥ 1 kWh (see Eq 8).

The top row in Figure 3 shows the fairness decision boundaries realized by active learning for experiments 1, 2, and 3. The blue-colored region indicates fair pricing and the red-colored region represents unfair pricing for each experiment. The red and blue dots in these plots indicate the price points chosen by the active learning method to be evaluated by the simulation. The first plot shows the decision boundary for fair pricing for Experiment 1. LMI agent responses are recorded and the fairness of a price point is determined based on fairness criterion 1. The second figure is similar to the first, except that the response of non-LMI agents is used. Note, the non-LMI agents benefit slightly more than LMI community agents (refer to the blue area). The last figure in the row shows the Experiment 3 decision boundary in the pricing parameter space where an average of 1kWh/household peak demand reduction is achieved. It is observed that as the peak price increases, households become more flexible to shift activities outside peak hours.

It is clear that the decision boundaries for the experiments across multiple fairness criteria are distinct. It will be useful if one can combine these separate results to analyze commonalities and differences in the decision space. We use a metric to combine these different results. A *disagreement metric* is defined for this purpose. If M is the total number of points generated in the parameter space, and M_0 is the number of points that are labeled as L_0. A similar definition applies to M_1. To calculate the disagreement metric, count the number of points M' that are labeled differently across multiple scenarios. Then, the disagreement is M'/M. The plot in the bottom row of Figure 3 highlights four distinct regions in the parameter space using the *disagreement metric*. Note, the region colored green (intersection of the blue region across three experiments) denotes the fair pricing region, i.e., it satisfies both the fairness criteria. The other partially fair and unfair regions in the parameter space are highlighted in different colors to indicate their meanings. Additional results and energy demand profiles before and after fair TOU pricing is implemented are attached in the Appendix.

5.1 Structural validation

All the datasets employed in the simulation framework are obtained from trustworthy sources and have been validated with ground truth data. The flat rate tariff, peak, and non-peak price ranges used in the simulation are adapted from power company websites. The digital twin of the energy dataset used in this work has been extensively validated at two different time scales and across representative regions of the U.S. using ground truth datasets [40, 41].

The data on the preference and probability of shifting different appliances outside peak hours is adapted from a survey conducted by Stelmach et al. [39] based on 337 households in California who were to be a part of the TOU pricing program. The survey questions were designed with concepts from social practice theory and focused on household practices and activities that understand the agent's appliance usage during peak times and how these affect overall energy use. This gives us an insight into household's preferences who would opt for a TOU pricing scheme. We use the survey findings to inform our calculations of agents' flexibility factors. It supports how agents will respond to price signals and if or how they change their behaviors.

The literature cited in the Background Section does not consider fairness in dynamic pricing. Thus, it is challenging to validate our results using ground truth data. There is also a general lack of data in the literature about user behavior toward nudges when responding to dynamic prices over a period of time. Our framework uses digital twins and surveys from the literature to show that modeling agent behaviors to dynamic price nudges is possible and reveals a reduction in both: monthly bills and peak reduction. These findings have been reported by dynamic pricing trials as well (mentioned in Background section). In addition, the simulation shows that it is possible for a ***fair TOU*** pricing scheme to achieve both goals.

5.2 Discussion

Designing agent-based simulations for reasonable population size (e.g., city, state) without consumer trial data is difficult. However, setting up large-scale trials is a long and arduous process with added complications such as consumer privacy. In such cases, the burden of conducting trials in the real world can be substantially reduced by introducing ML and simulations with domain knowledge. In this work, a fair dynamic pricing scheme is developed under an equal opportunity scenario by using active learning with agent-based simulations. Our results for a rural region show different TOU pricing thresholds for LMI and non-LMI communities to gain benefits in terms of monthly bills and peak demand reduction. This discussion has been extended in the Appendix.

Note that the simulation framework for analysis is flexible and generic, such that it can be extended to other social good problems that can be modeled by simulation functions. The active learning process and disagreement metrics are data-independent. The agent-based models can be switched with other appliance scheduling models from the literature subject to data availability. Our work showcases that comprehensive synthetic datasets can play crucial roles in building detailed agent models. This may be increasingly in favor of supporting economists and power companies to make important evidence-based policy decisions. Our work is developed using open-source data and published surveys.

Active learning provides an excellent foundation for characterizing simulation behaviors in unlabeled parameter space. The active learning algorithm can be extended to navigate high-dimensional spaces with appropriate visualizations. Thus, analyzing the fairness of real-time pricing is now possible using this framework. While techniques such as MILPs or stochastic optimization can also be used in place of active learning, they become slower as the number of data points as well as dimensions grow.

In future work, different fairness metrics can be employed by using household characteristics other than income to compute individual benefits (e.g., reduction in energy burden), and by including real-time pricing (RTP) and/or demand details at the utility level such as Electric Reliability Council of Texas (ERCOT). These types of metrics are used in incentive design trials where behavior change is induced by nudging consumers through communication channels such as emails or text messages just before the peak time pricing starts and advertising the possible savings that can be achieved. As summers become more extreme every year, one can study the effects of implementing dynamic pricing under such extreme weather scenarios to examine the vulnerability of households in terms of comfort violation (e.g., indoor thermostat setting).

6 CONCLUSION

In this work, fairness constructs are formulated for implementing residential dynamic pricing schemes after studying principles of rate design from economic and social theory. An agent-based framework is developed for simulating households' responses to dynamic pricing signals through load reduction and load shift strategies. Then, we integrate this agent-based simulation in an active learning framework to efficiently navigate the pricing parameter space. Our results show that fair pricing schemes can be designed that satisfy primary goals of residential consumers and utilities.

ACKNOWLEDGMENTS

This work was partially supported by University of Virginia Strategic Investment Fund award number SIF160 and NSF Grant No.: OAC-1916805 (CINES).

REFERENCES

[1] Thamer Alquthami, Ahmad H. Milyani, Muhammad Awais, and Muhammad B. Rasheed. 2021. An incentive based dynamic pricing in smart grid: A customer's perspective. *Sustainability (Switzerland)* 13 (6 2021). Issue 11. https://doi.org/10.3390/su13116066

[2] Khursheed Aurangzeb, Sheraz Aslam, Syed Muhammad Mohsin, and Musaed Alhussein. 2021. A Fair Pricing Mechanism in Smart Grids for Low Energy Consumption Users. *IEEE Access* 9 (2021), 22035–22044. https://doi.org/10.1109/ACCESS.2021.3056035

[3] Brendon Baatz. 2017. Rate Design Matters: The Intersection of Residential Rate Design and Energy Efficiency. *American Council for an Energy-Efficient Economy, Report U1703* (2017), 1–72.

[4] Arpita Biswas and Suvam Mukherjee. 2019. Fairness Through the Lens of Proportional Equality. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems* (Montreal QC, Canada) *(AAMAS '19)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1832–1834.

[5] James Bonbright. 1961. Principles of Public Utility Rates. *New York: Columbia University Press* (1961). https://doi.org/10.1016/j.energy.2022.124978

[6] Severin Borenstein. 2012. *NBER paper series: Effective and equitable adoption of opt-in residential dynamic electricity pricing.* Technical Report. http://www.nber.org/papers/w18037

[7] Peter Cappers, C Anna Spurlock, Annika Todd, and Ling Jin. 2016. *Experiences of Vulnerable Residential Customer Subpopulations with Critical Peak Pricing.* Technical Report.

[8] Peter Cappers, C Anna Spurlock, Annika Todd, and Ling Jin. 2016. Experiences of Vulnerable Residential Customer Subpopulations with Critical Peak Pricing. (2016).

[9] Emre Celebi and J. David Fuller. 2012. Time-of-Use Pricing in Electricity Markets Under Different Market Structures. *IEEE Transactions on Power Systems* 27, 3 (2012), 1170–1181. https://doi.org/10.1109/TPWRS.2011.2180935

[10] Zhi Chen, Lei Wu, and Yong Fu. 2012. Real-time price-based demand response management for residential appliances via stochastic optimization and robust optimization. *IEEE Transactions on Smart Grid* 3 (2012), 1822–1831. Issue 4. https://doi.org/10.1109/TSG.2012.2212729

[11] Weiwei Cui and Lin Li. 2018. A game-theoretic approach to optimize the Time-of-Use pricing considering customer behaviors. *International Journal of Production Economics* 201, C (4 2018), 75–88. https://doi.org/10.1016/j.ijpe.2018.04.02

[12] Celso M. de Melo, Stacy Marsella, and Jonathan Gratch. 2017. Increasing Fairness by Delegating Decisions to Autonomous Agents. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems* (São Paulo, Brazil) *(AAMAS '17)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 419–425.

[13] Jacqueline M Doremus, Irene Jacqz, Sarah Johnston, Carlos Flores, Teevrat Garg, Corbett Grainger, Steve Hamilton, Christopher Sullivan, Adam Theising, William Wheeler, Corey White, Justin Winikoff, and Eduardo Zambrano. 2021. Sweating the energy bill: Extreme weather, poor households, and the energy spending gap. (2021).

[14] Goutam Dutta and Krishnendranath Mitra. 2017. A literature review on dynamic pricing of electricity. *Journal of the Operational Research Society* 68 (10 2017), 1131–1145. Issue 10. https://doi.org/10.1057/s41274-016-0149-4

[15] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference* (Cambridge, Massachusetts) *(ITCS '12)*. Association for Computing Machinery, New York, NY, USA, 214–226. https://doi.org/10.1145/2090236.2090255

[16] Andrew Estornell, Sanmay Das, Yang Liu, and Yevgeniy Vorobeychik. 2023. Group-Fair Classification with Strategic Agents. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) *(FAccT '23)*. Association for Computing Machinery, New York, NY, USA, 389–399. https://doi.org/10.1145/3593013.3594006

[17] Ahmad Faruqui. 2015. *Time-Variant Pricing (TVP) in New York The REV agenda and residential time-variant pricing NYU School of Law, New York.* Technical Report.

[18] Shannon Gallagher, Lee F. Richardson, Samuel L. Ventura, and William F. Eddy. 2018. SPEW: Synthetic Populations and Ecosystems of the World. *Journal of Computational and Graphical Statistics* 27, 4 (2018), 773–784. https://doi.org/10.1080/10618600.2018.1442342

[19] Volker Grimm, Uta Berger, Donald L. DeAngelis, J. Gary Polhill, Jarl Giske, and Steven F. Railsback. 2010. The ODD protocol: A review and first update. *Ecological Modelling* 221, 23 (24 Nov. 2010), 2760–2768. https://doi.org/10.1016/j.ecolmodel.2010.08.019 Funding Information: We would like to thank two anonymous reviewers for comments on an earlier draft and all those who provided feedback on the ODD protocol, in particular: Thomas Banitz, Dan Brown, Jürgen Groeneveld, Jean Le Fur, Roger Jovani, Martin Köchy, Dawn Parker, Cyril Piou, Eva Roßmanith, Nadja Rüger, Gillian Salerno and Amelie Schmolke. DLD was supported by the Biological Resources Division of the U.S. Geological Survey, JGP was supported by funds from the Scottish Government Rural and Environment Research and Analysis Directorate.

[20] Bowei Guo and Melvyn Weeks. 2022. Dynamic tariffs, demand response, and regulation in retail electricity markets. *Energy Economics* 106 (2 2022). https://doi.org/10.1016/j.eneco.2021.105774

[21] Karen Herter. 2007. Residential implementation of critical-peak pricing of electricity. *Energy Policy* 35, 4 (2007), 2121–2130. https://doi.org/10.1016/j.enpol.2006.06.019

[22] William W Hogan. 2010. *Fairness and dynamic pricing: Comments.* Technical Report.

[23] Ijaz Hussain, Sajjad Mohsin, Abdul Basit, Zahoor Ali Khan, Umar Qasim, and Nadeem Javaid. 2015. A review on demand response: Pricing, optimization, and appliance scheduling. *Procedia Computer Science* 52, 843–850. Issue 1. https://doi.org/10.1016/j.procs.2015.05.141

[24] Mikko Jalas and S. Numminen. 2022. Prime-time access for whom? Rhythms fairness and the dynamic pricing of infrastructure services. *Local Environment* 27, 10-11 (2022), 1355–1371. https://doi.org/10.1080/13549839.2022.2040468 arXiv:https://doi.org/10.1080/13549839.2022.2040468

[25] E. A.M. Klaassen, C. B.A. Kobus, J. Frunt, and J. G. Slootweg. 2016. Responsiveness of residential electricity demand to dynamic tariffs: Experiences from a large field test in the Netherlands. *Applied Energy* 183 (12 2016), 1065–1074. https://doi.org/10.1016/j.apenergy.2016.09.051

[26] Amanda R. Kube, Sanmay Das, and Patrick J. Fowler. 2023. Fair and Efficient Allocation of Scarce Resources Based on Predicted Outcomes: Implications for Homeless Service Delivery. *J. Artif. Int. Res.* 76 (may 2023), 27. https://doi.org/10.1613/jair.1.12847

[27] Victor Li, Jacqueline Lam, Peiyang Guo, David Newbery, Kelvin Chan, Jacqueline CK Lam, and Kong OK Victor Li. 2022. A Data-driven Fairness-based Time of Use Tariff Design. (2022). https://doi.org/10.21203/rs.3.rs-1529952/v1

[28] Lydia T. Liu, Sarah Dean, Esther Rolf, Max Simchowitz, and Moritz Hardt. 2019. Delayed Impact of Fair Machine Learning. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization, 6196–6200. https://doi.org/10.24963/ijcai.2019/862

[29] Jan Maly, Simon Rey, Ulle Endriss, and Martin Lackner. 2023. Fairness in Participatory Budgeting via Equality of Resources. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems* (London, United Kingdom) *(AAMAS '23)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2031–2039.

[30] Tasfia Mashiat, Xavier Gitiaux, Huzefa Rangwala, and Sanmay Das. 2023. Counterfactually Fair Dynamic Assignment: A Case Study on Policing. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems* (London, United Kingdom) *(AAMAS '23)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2526–2528.

[31] Tasfia Mashiat, Xavier Gitiaux, Huzefa Rangwala, Patrick Fowler, and Sanmay Das. 2022. Trade-offs between Group Fairness Metrics in Societal Resource Allocation. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) *(FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 1095–1105. https://doi.org/10.1145/3531146.3533171

[32] Fan-Lin Meng and Xiao-Jun Zeng. 2013. A Stackelberg Game-Theoretic Approach to Optimal Real-Time Pricing for the Smart Grid. *Soft Comput.* 17, 12 (12 2013), 2365–2380. https://doi.org/10.1007/s00500-013-1092-9

[33] Stijn Neuteleers, Machiel Mulder, and Frank Hindriks. 2017. Assessing fairness of dynamic grid tariffs. *Energy Policy* 108 (2017), 111–120. https://doi.org/10.1016/j.enpol.2017.05.028

[34] Olabambo Ifeoluwa Oluwasuji, Obaid Malik, Jie Zhang, and Sarvapali Dyanand Ramchurn. 2020. Solving the fair electric load shedding problem in developing countries. *Auton. Agent. Multi. Agent. Syst.* 34, 1 (April 2020).

[35] Bonnie Wylie Pratt and Jon D. Erickson. 2020. Defeat the Peak: Behavioral insights for electricity demand response program design. *Energy Research and Social Science* 61 (3 2020). https://doi.org/10.1016/j.erss.2019.101352

[36] Bonnie Wylie Pratt and Jon D. Erickson. 2020. Defeat the Peak: Behavioral insights for electricity demand response program design. *Energy Research & Social Science* 61 (2020), 101352. https://doi.org/10.1016/j.erss.2019.101352

[37] Jose Luis Ramirez-Mendiola and Jacopo Torriti. 2022. The price is not right! Energy demand, Time of Use tariffs, values and social practices. *ECEEE Summer Study* (2022), 75–84.

[38] Tobias Scheffer, Christian Decomain, and Stefan Wrobel. 2001. Active Hidden Markov Models for Information Extraction. In *Proceedings of the 4th International Conference on Advances in Intelligent Data Analysis (IDA '01)*. Springer-Verlag, Berlin, Heidelberg, 309–318.

[39] Greg Stelmach, Chad Zanocco, June Flora, Ram Rajagopal, and Hilary S. Boudet. 2020. Exploring household energy rules and activities during peak demand to better determine potential responsiveness to time-of-use pricing. *Energy Policy* 144 (9 2020), 1–11. https://doi.org/10.1016/j.enpol.2020.111608

[40] Swapna Thorve, Young Yun Baek, Samarth Swarup, Henning Mortveit, Achla Marathe, Anil Vullikanti, and Madhav Marathe. 2023. High resolution synthetic residential energy use profiles for the United States. *Scientific Data* 10, 1 (06 Feb

2023), 76. https://doi.org/10.1038/s41597-022-01914-1
[41] Swapna Thorve, Anil Vullikanti, Henning S Mortveit, Samarth Swarup, and Madhav V. Marathe. 2022. Fidelity and diversity metrics for validating hierarchical synthetic data: Application to residential energy demand. In *2022 IEEE International Conference on Big Data (Big Data)*. 1377–1382. https://doi.org/10.1109/BigData55660.2022.10020837
[42] Lee V. White and Nicole D. Sintov. 2020. Health and financial impacts of demand-side response measures differ across sociodemographic groups. *Nature Energy* 5, 1 (01 1 2020), 50–60. https://doi.org/10.1038/s41560-019-0507-y
[43] Lee V. White and Nicole D. Sintov. 2020. Varied health and financial impacts of time-of-use energy rates across sociodemographic groups raise equity concerns. *Nature Energy* 5, 1 (01 1 2020), 16–17. https://doi.org/10.1038/s41560-019-0515-y
[44] Tim G Williams, Daniel G Brown, Seth D Guikema, Tom M Logan, Nicholas R Magliocca, Birgit Müller, and Cara E Steger. 2022. Integrating Equity Considerations into Agent-Based Modeling: A Conceptual Framework and Practical Guidance. *Journal of Artificial Societies and Social Simulation* 25, 3 (2022), 1. https://doi.org/10.18564/jasss.4816
[45] Peng Yang, Gongguo Tang, and Arye Nehorai. 2013. A game-theoretic approach for optimal time-of-use electricity pricing. *IEEE Transactions on Power Systems* 28 (2013), 884–892. Issue 2. https://doi.org/10.1109/TPWRS.2012.2207134
[46] Timur Yunusov and Jacopo Torriti. 2021. Distributional effects of Time of Use tariffs based on electricity demand and time use. *Energy Policy* 156 (2021), 112412. https://doi.org/10.1016/j.enpol.2021.112412