Evaluating Probabilistic Load Forecasting in Predicting Long-Term Household Savings from Electricity Rate Plan Recommendations

Xavier Farrell
Electrical and Computer Engineering
University of Michigan
Ann Arbor, MI, USA
xfarrell@umich.edu

Johanna L. Mathieu
Electrical and Computer Engineering
University of Michigan
Ann Arbor, MI, USA
jlmath@umich.edu

2024 56th North American Power Symposium (NAPS) | 979-8-3315-2103-5/24/$31.00 ©2024 IEEE | DOI: 10.1109/NAPS61145.2024.10741851

***Abstract*—Electricity rate plan recommender systems (EPRS) are tools that aim to predict and recommend which rate plans will be the least expensive for a household. In addition to its rate plan recommendation, EPRS could be greatly improved if they also generated probabilistic forecasts for the savings the household might incur if enrolled. Probabilistic load forecasting could be used to forecast cumulative usage costs and potential savings. However, load forecasts are generated from imperfect models, and so cumulative savings forecast distributions will also be imperfect. In this paper, we apply probabilistic load forecasting to generate cumulative savings forecast distributions for EPRS recommendations and show how load forecasting error affects savings forecasts. We show that these effects differ in severity over forecast horizons of a day, between months in the year, and over the full year. We show across 865 real households and 12 months that inaccuracies in forecasted load distributions can form bias in short-term savings forecasts but also compound, worsening the forecasted cumulative savings distributions. We conclude that forecasting cumulative savings from EPRS recommendations using probabilistic load forecasting has potential but requires careful consideration or mitigation of forecast bias induced by the nonstationarity of load.**

***Index Terms*—probabilistic load forecasting, electricity rate plan recommender systems, forecast bias, residential load**

I. Introduction

With the increasing prevalence of advanced metering infrastructure, electric utilities are able to offer their residential customers a suite of electricity rate plan options [1], [2]. Algorithms and models called electricity rate plan recommender systems (EPRS) have emerged to help households save on their electricity bills [3]. EPRS recommend the rate plans predicted to be lowest-cost based on a household's features (i.e., monthly usage, appliance ownership, household income, etc.) [4]. However, current EPRS do not generate predictions for how much money and at what times over the coming months or year the household can expect to save money after enrolling. Such predictions would allow the household to better plan or adapt their energy usage and other expenses, ultimately enabling them to have greater control over both, which is important to low income homes [2].

This work was supported by U.S. NSF Award ECCS-1952038.

One approach to predicting future (i.e., "forecasting") savings from an EPRS recommendation is to forecast home load and compute the cost on each rate plan of interest. The difference between the costs on each rate plan estimates the expected savings. However, due to load uncertainty stemming from weather uncertainty and demand variation and change, it is common to approach forecasting probabilistically [5]. Furthermore, probabilistic load forecasting offers more information regarding long-term possibilities. This is relevant to forecasting over horizons comparable to the duration of enrollment contracts, which can require enrollment for a year [6]. From these probabilistic load forecasts, probabilistic cost and savings forecasts can be generated. However, load forecast distributions are generated from imperfect load models, and so the savings distribution forecasts will also be imperfect.

Inaccuracies in the forecasted load distributions can give rise to bias in the savings forecasts. This can lead to increasingly erroneous long-term cumulative savings forecasts as the effect of bias accumulates over multi-day forecast horizons. If such savings forecasts were used by a household in decision making, it could expose them to unanticipated financial risk and hardship. Therefore, it is important to understand how probabilistic load forecasting inaccuracies affect long-term probabilistic savings forecasts.

In this paper, we demonstrate the potential downstream effects of probabilistic load forecast error on savings forecasts corresponding to EPRS recommendations. We employ a probabilistic load forecasting method in which we first combine and post process the outputs of two temperature dependent load models to generate home load scenarios. Then, we use the load scenarios to forecast the range of savings that the home can incur on each day, and forecast a probability distribution describing the likeliness of the savings outcome. We present a case study utilizing multi-year home load data from 865 real households in Detroit, MI and show how periods of bias in short-term savings forecasts compound in the long-term to produce inaccurate cumulative savings forecasts. Furthermore, we explore these effects considering various forecast horizon durations and periods of the year.

Past work has developed EPRS to identify and recommend to a household its least-expensive rate plan options, often leveraging indirect relationships between household features and rate plan suitability [3], [4], [7], [8]. However, many of these methods often cannot forecast usage costs and savings if a household were to enroll in the recommended rate plan. Some works that develop methods to probabilistically forecast load have considered its application to usage cost predictions and rate plan recommendations [9], [10]. Methods presented for short- and long-term probabilistic load forecasting often generate distributions for load at each step in time in a forecast horizon [9], [11]–[13]. A method of this kind is extended to predict the potential for positive savings from a rate plan recommendation over forecast horizons of one week [10]. Performance metrics applied to these methods often measure how well each individual distribution predicts the load at its corresponding step in time. However, the interaction between each distribution becomes critical when forecasting costs and savings that accumulate over time steps in the forecast horizon, but has mostly gone unconsidered. This motivates our work which applies and assesses probabilistic load forecasting in cumulative savings predictions for EPRS recommendations.

The contributions of this paper are three-fold. First, we expand the electricity rate plan recommendation problem to include a forecast for long-term cumulative savings and we develop an initial method to generate probabilistic savings forecasts leveraging load models. Second, we study how load forecast error affects daily savings and long-term cumulative savings, highlighting variation as the forecast horizon duration and period in the year changes. Third, we highlight the compounding effects of bias in short-term savings forecasts on cumulative savings forecasts, motivating the use of methods to mitigate forecast bias induced by the nonstationarity of load.

This paper is organized as follows. Section II introduces the methods used in this work, including the electricity load model and how the model output and residuals are used to generate probabilistic load scenarios. These load scenarios are then used to create savings scenarios, from which we form daily and cumulative savings distribution forecasts. Section III presents a case study illustrating the effects of load forecast bias on savings distribution forecasts. Section IV concludes.

II. Methods

Fig. 1 summarizes our method for forecasting a probability distribution that describes the likeliness of cumulative savings outcomes for a household after enrolling in a recommended rate plan. First, we train and implement regression-based load models to forecast home load from temperature forecasts and calendar information. We use these load models to generate daily point forecasts of home load as well as a discrete probability distribution of forecast error (the green block). For each day, we form all combinations of the point forecast and forecast errors and generate a subset of feasible load scenarios (the blue block), approximating the domain of home load outcomes. Next, the difference in tariff charge (i.e., "savings") incurred by each load scenario is computed and used to construct a set of savings scenarios for the household (the purple block). Then, leveraging the probability of each forecast error, we create a probability mass function (PMF) of short-term (daily) savings (the pink block). Finally, we convolve the short-term savings distributions to obtain the long-term cumulative savings distribution forecast (the yellow block). Note that our method does not model behavioral changes in response to enrollment in the recommended rate plan. The following subsections describe each step in detail.

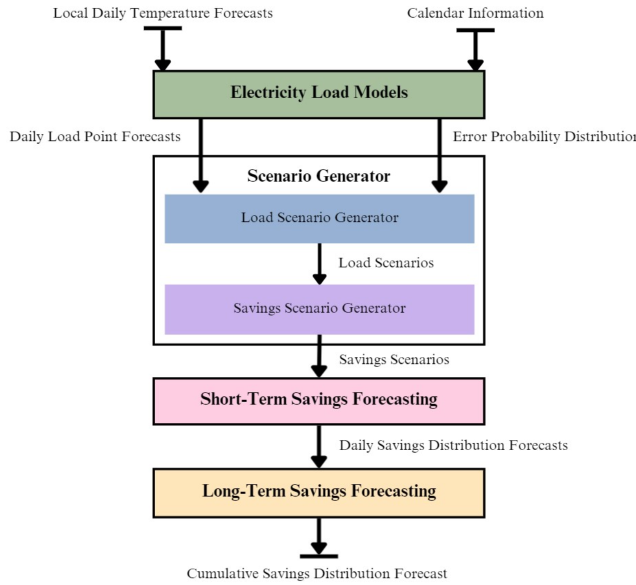

Fig. 1. A block diagram for probabilistically forecasting home load and cumulative savings from rate plan enrollment.

A. Electricity Load Models

In this section, we describe the "Electricity Load Models" block in Fig. 1. In our case study, we consider two electric rates – a time-of-use (TOU) rate and an increasing block tariff (IBT) rate – and so we use models that allow us to estimate relevant loads and costs on both rates. Specifically, we use two home load models to predict total load during two different periods of a day – on-peak and off-peak – based on forecasted temperature and the part of week. Note that this method could be easily extended to consider more pricing periods. Here, we describe the historical data used in constructing the model, the model formulation and training process, and our approach to probabilistically model load forecast error, wherein we first approximate the error domain with forecast error scenarios and then ascribe probabilities to the scenarios.

1) Historical Data: We first describe the historical load, temperature and calendar information used in constructing the two home load models. First, let z_j^{On} be an indicator whose value is 1 if the j^{th} hour is within the "on-peak period" and is 0 for all hours j outside of the on-peak period, referred to as the "off-peak period". Vectors $\boldsymbol{u}^{\text{On}}$ and $\boldsymbol{u}^{\text{Off}}$ contain each day i's total on-peak period load u_i^{On} and off-peak period load u_i^{Off} in kWh, respectively, i.e., $\boldsymbol{u}^{\text{On}} = [u_1^{\text{On}}, \dots u_N^{\text{On}}]$ and $\boldsymbol{u}^{\text{Off}} =$

$[u_1^{\text{Off}}, \dots u_N^{\text{Off}}]$, where N is the number of days of historical data. Vectors $\boldsymbol{t}^{\text{On}}$ and $\boldsymbol{t}^{\text{Off}}$ contain each day i's mean on-peak period outdoor air temperature t_i^{On} and off-peak period outdoor air temperature t_i^{Off}, respectively, i.e., $\boldsymbol{t}^{\text{On}} = [t_1^{\text{On}}, \dots, t_N^{\text{On}}]$ and $\boldsymbol{t}^{\text{Off}} = [t_1^{\text{Off}}, \dots, t_N^{\text{Off}}]$. We construct the load models such that the load's sensitivity to temperature can differ between low and high temperature conditions using a method similar to that presented in [14]. The key difference is that rather than using fixed temperature ranges to define the load's sensitivity, we allow the low and high temperature ranges for load sensitivity to change based on the home's historical usage patterns. For the on-peak period load model, we split the temperature into two temperature components,

$$\tau_{i,1}^{\text{On}} = \min(t_i^{\text{On}} - \underline{t}^{\text{On}}, t^{\text{On,c}} - \underline{t}^{\text{On}}), \tag{1}$$
$$\tau_{i,2}^{\text{On}} = \max(0, t_i^{\text{On}} - t^{\text{On,c}}), \tag{2}$$

where $\underline{t}^{\text{On}}$ is the historical minimum mean on-peak temperature and $t^{\text{On,c}}$ is the change-point temperature, which is the point of transition between low and high temperatures. Then, $\boldsymbol{\tau}_i^{\text{On}} = [\tau_{i,1}^{\text{On}}, \tau_{i,2}^{\text{On}}]$, and $\boldsymbol{\tau}^{\text{On}} \in \mathbb{R}^{N \times 2}$ stacks all $\boldsymbol{\tau}_i^{\text{On}}$. The off-peak period load model is defined similarly. The values for $t^{\text{On,c}}$ and $t^{\text{Off,c}}$ are chosen to maximize the R-squared value between each model's predictions and historical observation.

We also process historical calendar information to form a vector $\boldsymbol{w}$ with elements w_i equal to 1 if day i corresponds to Monday-Friday or equal to 0 if day i corresponds to Saturday or Sunday, i.e., $\boldsymbol{w} = [w_1, w_2, \dots, w_N]$.

2) Model Formulation and Training: We next describe the model formulation and training process. The load models used in this paper are also inspired by those presented in [14] but adapts the model outputs to forecast total load over several hours. Both the on-peak period and off-peak period load models are linear regression models each with its own set of parameters $\boldsymbol{\theta} = \{\alpha^{\text{WD}}, \alpha^{\text{WE}}, \boldsymbol{\beta}\}$, where α^{WD} captures the weekday influence, α^{WE} captures the weekend influence, and $\boldsymbol{\beta} = [\beta_1, \beta_2]$ captures the temperature influence on the load. The on-peak period load model uses future calendar information w_i, temperature components $\boldsymbol{\tau}_i^{\text{On}}$ computed from a temperature forecast t_i^{On}, and model parameters $\boldsymbol{\theta}^{\text{On}}$ to forecast the total on-peak load on day i according to

$$u_i^{\text{On}} = \alpha^{\text{On,WD}} w_i + \alpha^{\text{On,WE}}(1 - w_i) + \sum_{k=1}^{2} \beta_k^{\text{On}} \tau_{i,k}^{\text{On}}, \tag{3}$$

and the off-peak period load model is defined similarly. Using ordinary least squares, we compute the on-peak period load model parameters $\boldsymbol{\theta}^{\text{On}}$ using on-peak period historical data for $\{\boldsymbol{u}, \boldsymbol{w}, \boldsymbol{\tau}\}$, and we compute the off-peak period load model parameters similarly. Note that future work may incorporate household price elasticity and load flexibility into the load models to capture a household's response to rate plan enrollment (e.g., reducing/shifting load to minimize costs).

3) Modeling the Domain of Load Forecast Error: Here, we describe how we approximate the domain of the load forecast error with a set of error scenarios. We use historical data to evaluate the load models to generate an on-peak period load estimate $\hat{u}_i^{\text{On}}$ and off-peak period load estimate $\hat{u}_i^{\text{Off}}$ and compute the model residuals for each day i as the difference between the actual load and the estimated load,

$$e_i^{\text{On}} = u_i^{\text{On}} - \hat{u}_i^{\text{On}}, \tag{4}$$
$$e_i^{\text{Off}} = u_i^{\text{Off}} - \hat{u}_i^{\text{Off}}. \tag{5}$$

We include all residuals in the sets $\mathcal{E}^{\text{On}} = \{e_1^{\text{On}}, e_2^{\text{On}}, \dots, e_N^{\text{On}}\}$ and $\mathcal{E}^{\text{Off}} = \{e_1^{\text{Off}}, e_2^{\text{Off}}, \dots, e_N^{\text{Off}}\}$. Next, we form a histogram with K^{On} uniform bins from the set $\mathcal{E}^{\text{On}}$, and a histogram with K^{Off} uniform bins from the set $\mathcal{E}^{\text{Off}}$, and define the approximate joint residual for each day i as the pair $\bar{e}_i = (\bar{e}_i^{\text{On}}, \bar{e}_i^{\text{Off}})$, where $\bar{e}_i^{\text{On}}$ corresponds to the midpoint of the bin within which e_i^{On} falls, and $\bar{e}_i^{\text{Off}}$ corresponds to the midpoint of the bin within which e_i^{Off} falls. Let $\bar{\mathcal{E}}$ include all $\bar{e}_i$, and $\bar{\mathcal{E}}^{\text{r}} \subset \bar{\mathcal{E}}$ include all unique $\bar{e}_i$, with elements $\bar{e}_r$ for $r = 1, \dots, N^{\text{r}}$. Then, $\bar{\mathcal{E}}^{\text{r}}$ is the domain of the load forecast error.

4) Forecast Error Probability Distributions: Here, we describe how probabilities are ascribed to each scenario within $\bar{\mathcal{E}}^{\text{r}}$. The use of residuals to form parametric and nonparametric distributions for load forecasts is well established [15] and we implement a nonparametric approach. This allows us to take advantage of our regression-based (point forecast) load model and historical observations to inform error modeling while limiting the assumptions imposed on the shape of the error distributions. We model load forecast error with separate distributions for high-temperature days and low-temperature days, because we have found that errors on these days often differ significantly. We define high-temperature days as days in which $t_i^{\text{On}} \geq t^{\text{On,c}}$ or $t_i^{\text{Off}} \geq t^{\text{Off,c}}$, and all other days are low-temperature days. Subset $\bar{\mathcal{E}}^{\text{H}} \subset \bar{\mathcal{E}}$ includes all $\bar{e}_i$ occurring on all high-temperature days (and includes N^{H} total elements) and subset $\bar{\mathcal{E}}^{\text{L}} \subset \bar{\mathcal{E}}$ includes all $\bar{e}_i$ occurring on all low-temperature days (and includes N^{L} total elements). Elements of $\bar{\mathcal{E}}^{\text{r}}$ can appear in both $\bar{\mathcal{E}}^{\text{H}}$ and $\bar{\mathcal{E}}^{\text{L}}$.

Assuming the approximate joint residuals $\bar{e}_i$ are i.i.d. (albeit a somewhat unrealistic assumption), we can ascribe probabilities to both the high- and low-temperature day load forecast error scenarios. For high-temperature days, the approximate probability of the pair $\bar{e}_r$ is

$$\Pr(x = \bar{e}_r) \approx \frac{n_r^{\text{H}}}{N^{\text{H}}} = p_r^{\text{H}}, \tag{6}$$

where n_r^{H} is the number of instances of $\bar{e}_r$ in $\bar{\mathcal{E}}^{\text{H}}$. Similarly, for low-temperature days, the approximate probability of the pair $\bar{e}_r$ is

$$\Pr(x = \bar{e}_r) \approx \frac{n_r^{\text{L}}}{N^{\text{L}}} = p_r^{\text{L}}, \tag{7}$$

where n_r^{L} is the number of instances of $\bar{e}_r$ in $\bar{\mathcal{E}}^{\text{L}}$. Then, the empirical high-temperature day joint PMF is

$$\boldsymbol{F}^{\text{H}}(x = \bar{e}_r) = \begin{cases} p_r^{\text{H}} & \text{if } \bar{e}_r \in \bar{\mathcal{E}}^{\text{H}} \\ 0 & \text{otherwise} \end{cases}, \text{ for } r = 1, \dots, N^{\text{r}}, \tag{8}$$

and the empirical low-temperature day joint PMF is

$$\boldsymbol{F}^{\mathrm{L}}(x=\bar{e}_r)=\begin{cases} p_r^{\mathrm{L}} & \text{if } \bar{e}_r\in\bar{\mathcal{E}}^{\mathrm{L}} \\ 0 & \text{otherwise}\end{cases}, \text{ for } r=1,\ldots,N^{\mathrm{r}}, \tag{9}$$

and the full joint PMF for each day i is

$$\boldsymbol{F}_i(x=\bar{e}_r)=\begin{cases} \boldsymbol{F}^{\mathrm{H}}(x=\bar{e}_r) & \text{if high-temperature day} \\ \boldsymbol{F}^{\mathrm{L}}(x=\bar{e}_r) & \text{if low-temperature day}\end{cases}. \tag{10}$$

B. Scenario Generator: Load and Savings

This section describes the "Scenario Generator" block in Fig. 1. We first generate home load scenarios using the load forecast error scenarios. We then compute the savings scenarios, assuming two specific electricity rate options.

1) Load Scenario Generator: We present the process for constructing the load scenarios corresponding to a high-temperature day; however, the same process is followed for low-temperature days. Let $\bar{e}_r=(\bar{e}_r^{\mathrm{On}},\bar{e}_r^{\mathrm{Off}})$. We assume the true load might differ from the forecast $(u_i^{\mathrm{On}},u_i^{\mathrm{Off}})$ by any $\bar{e}_r\in\bar{\mathcal{E}}^{\mathrm{H}}$. We compute the corrected on-peak period and off-peak period load scenario for day i and unique residual r as

$$u_{i,r}^{\mathrm{On}}=u_i^{\mathrm{On}}+\bar{e}_r^{\mathrm{On}}, \tag{11}$$
$$u_{i,r}^{\mathrm{Off}}=u_i^{\mathrm{Off}}+\bar{e}_r^{\mathrm{Off}}. \tag{12}$$

We use the indicator function $f(u_{i,r}^{\mathrm{On}},u_{i,r}^{\mathrm{Off}})=1$ to indicate whether a scenario is feasible, i.e., $u_{i,r}^{\mathrm{On}}\geq 0$ and $u_{i,r}^{\mathrm{Off}}\geq 0$. Otherwise, $f(u_{i,r}^{\mathrm{On}},u_{i,r}^{\mathrm{Off}})=0$.

2) Savings Scenario Generator: Next, we describe how we approximate a domain for the savings the household could incur on each day. Let the TOU rate plan be recommended to the household and let the alternative rate plan be an IBT. We compute the savings on day i with unique residual r as

$$s_{i,r}=C_{i,r}^{\mathrm{B}}-C_{i,r}^{\mathrm{T}}, \tag{13}$$

where $C_{i,r}^{\mathrm{B}}$ is the cost of load under the IBT rate plan and $C_{i,r}^{\mathrm{T}}$ is the cost of load under the TOU rate plan on day i with unique residual r. Then, the set of feasible savings values for day i is

$$\mathcal{S}_i=\{s_{i,r},\, r=1,\ldots,N^{\mathrm{r}}\mid f(u_{i,r}^{\mathrm{On}},u_{i,r}^{\mathrm{Off}})=1\}, \tag{14}$$

and the set of all savings across all days i within a forecast horizon is $\mathcal{S}^+=\{\mathcal{S}_i,\, i=1,\ldots,N^{\mathrm{F}}\}$, where N^{F} is the number of days in the forecast horizon. We form a histogram with K^{s} uniform bins from the set $\mathcal{S}^+$, and define the approximate savings scenarios $x_k^{\mathrm{s}}\in\mathcal{X}^{\mathrm{s}}$, for $k=1,\ldots,K^{\mathrm{s}}$, as the midpoints of the bins, where $\mathcal{X}^{\mathrm{s}}$ approximates the savings domain across the forecast horizon.

C. Short-Term Savings Forecasts

We next describe the "Short-Term Savings Forecast" block in Fig. 1, which forecasts a probability distribution for each day's savings. We use the indicator function $v_k(s_{i,r})=1$ to indicate whether $s_{i,r}$ is in bin k. Otherwise, $v_k(s_{i,r})=0$. Then, the approximate probability of x_k^{s} for each day i in the forecast horizon is

$$\Pr(x=x_k^{\mathrm{s}})=\frac{\sum_{r=1}^{N^{\mathrm{r}}}\boldsymbol{F}_i(x=\bar{e}_r)v_k(s_{i,r})f(u_{i,r}^{\mathrm{On}},u_{i,r}^{\mathrm{Off}})}{\sum_{r=1}^{N^{\mathrm{r}}}\boldsymbol{F}_i(x=\bar{e}_r)f(u_{i,r}^{\mathrm{On}},u_{i,r}^{\mathrm{Off}})}=p_{i,k}, \tag{15}$$

and the empirical savings PMF for each day i in the forecast horizon is

$$\boldsymbol{F}_i^{\mathrm{s}}(x=x_k^{\mathrm{s}})=\{p_{i,k},k=1,2,\ldots,K^{\mathrm{s}}\}. \tag{16}$$

D. Long-Term Savings Forecast

Finally, we describe the "Long-Term Savings Forecast" block in Fig. 1, which computes the probability distribution for the cumulative savings incurred over a multi-day forecast horizon, resulting in a long-term probabilistic savings forecast for a household's enrollment into the recommended rate plan. If we were using point forecasts, the long-term savings would be the sum of the short-term (daily) savings. Since we are using forecast distributions, we approach the computation as determining the distribution corresponding to a sum of independent random variables [16]. We make the assumption that the distributions are independent (again, a somewhat unrealistic assumption) and convolve the short-term (daily) savings distributions to obtain the long-term savings distribution. Therefore, the cumulative savings distribution is computed as the convolution of the PMFs of each day in the forecast horizon, i.e.,

$$\boldsymbol{F}(x)=\boldsymbol{F}_i^{\mathrm{s}}(x)*\boldsymbol{F}_{i+1}^{\mathrm{s}}(x)*\ldots*\boldsymbol{F}_{N^{\mathrm{F}}}^{\mathrm{s}}(x). \tag{17}$$

III. Case Study

We next present a case study in which we demonstrate the potential downstream effects of inaccuracies in probabilistic load forecasts in daily and cumulative savings forecasts. We first describe our case study set-up and then describe the results.

A. Set Up

Here, we describe the specifics of our implementation of the aforementioned methodology including the preprocessing of data, the rate plans considered, and other modeling choices.

1) Input Data: We use hourly smart meter data from approximately 1000 real households in Detroit, MI. We use one year of home load data for model training (1 March 2019 - 29 February 2020) and one year of home load data for testing (1 March 2020 - 28 February 2021). We required all homes to have no missing data over the two year period. We also required that the training year home load have a total usage above 500 kWh as a proxy for the home being inhabited most of the year. All toll, we filtered out over 100 homes, resulting in 865 homes.

We use NOAA temperature data from Detroit Metropolitan Wayne County Airport Detroit, MI [17]. We use two years of hourly temperature measurements, corresponding to the same period as the load data, one year for model training and one year for forecasting. Both years of data are interpolated using spline interpolation to have readings on the hour. In forecasting load, we use a perfect temperature forecast because, though assessing the sensitivity of this method to temperature forecast uncertainty is valuable, this paper primarily aims to assess the effects of load forecast bias in savings forecasts with emphasis on the nonstationarity of load.

2) Rate Plans: The rate plans are modeled after rates historically offered to Detroit electricity customers [6]. The TOU rate plan's on-peak price period is 11AM - 7PM on weekdays and the off-peak price period are all hours outside of the on-peak period. The hot season months are from June to September and all other months are considered the cold season. The prices differ between seasons for the same price period. Let the total forecasted load on day i with unique residual r be

$$u_{i,r}^{\text{Tot}} = u_{i,r}^{\text{On}} + u_{i,r}^{\text{Off}}. \tag{18}$$

The charge associated with this load incurred on the TOU rate plan is

$$C_{i,r}^{\text{T}} = \pi_i^{\text{On}} u_{i,r}^{\text{On}} + \pi_i^{\text{Off}} u_{i,r}^{\text{Off}}, \tag{19}$$

where π_i^{On} and π_i^{Off} are the on-peak period and off-peak period prices for day i, determined by whether the day is a weekday or weekend and whether the day occurs during the hot season or cold season months.

The IBT rate plan has two blocks, with electricity consumption up to 17 kWh/day charged at a different rate than electricity consumption beyond 17 kWh. The charge associated with $u_{i,r}^{\text{Tot}}$ incurred on the IBT rate plan is

$$C_{i,r}^{\text{B}} = \pi_2^{\text{B}} \max(u_{i,r}^{\text{Tot}} - 17, 0) + \pi_1^{\text{B}} \min(u_{i,r}^{\text{Tot}}, 17), \tag{20}$$

where π_1^{B} and π_2^{B} are the prices for the first usage block and second usage block, respectively.

3) Other Modelling Choices: We apply a uniformly weighted 7x7 smoothing kernel to the high-temperature day and low-temperature day forecast error distributions. This was motivated by the sometimes sparse or jagged (i.e., many peaks and valleys) error distributions we observed and the assumption that errors within proximity of one another have similar probabilities. Similarly, we smooth each daily savings distribution using uniformly weighted moving averaging with a window size automatically chosen by the MATLAB function "smoothData".

B. Results

Other work has focused on improving probabilistic load forecasts and, as such, their results often include measures of precision and accuracy. In contrast, our work aims to highlight how errors in forecasted load distributions propagate to savings. Therefore, we use the Prediction Interval Coverage Probability Score (PICP) [13], which measures how frequently the true value of a forecasted quantity falls within forecasted prediction intervals.

We first generate daily load distribution forecasts for a single home over a year-long forecast horizon and explore the occurrence of load forecast bias presumably due to the nonstationarity of load. In Fig. 2, we plot the off-peak marginal distribution of each day's joint load distribution forecast. In

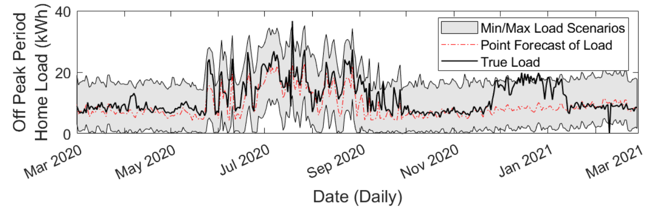

Fig. 2. Forecasted load distribution vs. true load, for a single home.

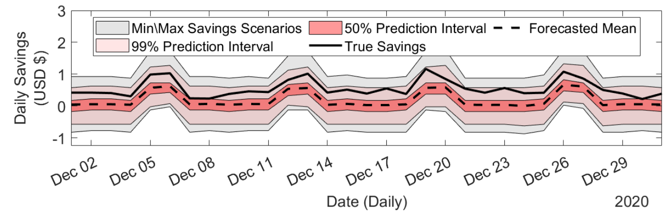

Fig. 3. Bias in forecasted daily savings distributions, for a single home.

December 2020 and the start of January 2021, the true off-peak period load (the solid black line) is much larger than the load during the same period in the historical training data set. This suggests nonstationarity in the home load, causing bias in the mean load forecast (the dotted red line) over that period. From this we observe that nonstationarity creates challenges for precisely forecasting future load with unobserved features; however, such scenarios can be accounted for when forecasting is approached probabilistically.

Next, we explore the effects of load forecast bias on daily savings forecasts. As more load is consumed in a day, especially as the day's cumulative load exceeds 17 kWh, the IBT charge grows more rapidly than the TOU charge, leading to greater savings generally. If the total forecasted load is underestimated, then it is possible that the savings will be underestimated as well. We observe this outcome in Fig. 3, where we plot the savings incurred for the same home using the true load compared to the savings forecasts, and their corresponding prediction intervals, for each day. We learn that though the true daily savings always remains within the 99% prediction interval (the light red area), the load forecast bias causes the mean daily savings forecast (the dotted black line) to be biased to underestimate the true daily savings.

Here, we shift to consider the effects of load forecast bias on cumulative savings forecasts over forecast horizons of a month, again for the same home. We convolve the savings forecasts of each day sequentially, from the first day of the month to the last day of the month, and plot each day's cumulative savings distribution in Fig. 4(a). In December, the mean cumulative savings forecasts (the red dotted line) increasingly underestimate the true savings (the solid black line). In December and January, the true cumulative savings exceeds the upper bound of the forecasted 99% prediction interval (the dark grey area), unlike the case for daily savings. From this, we learn that bias in load distribution forecasts propagate to daily savings forecasts but compound in cumulative savings.

Next, we consider a forecast horizon of a year. Note that the true daily savings associated with this home always remain within the forecasted 99% prediction interval for all 365 days

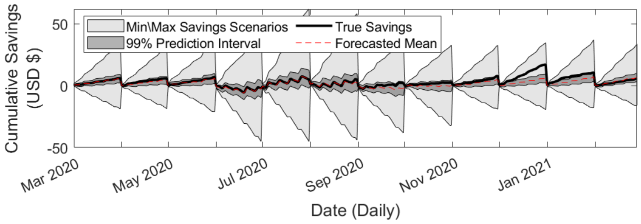

(a) Monthly prediction horizon.

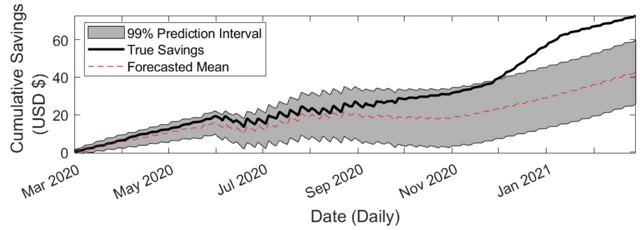

(b) Yearly prediction horizon.

Fig. 4. True savings vs. forecasted prediction intervals, for a single home.

in the forecast horizon. We once again convolve the savings forecast of each day, now from the first day of the year to the last day of the year, and plot each day's cumulative savings distribution in Fig. 4(b). We observe that the difference between the true cumulative savings and mean cumulative savings forecast grows, increasing noticeably post August, as a consequence of daily savings forecasts being biased to underestimate the true daily savings. This figure also reveals the residual effects of the forecast bias when compared to Fig. 4(a) in that while February overestimates the true savings, it does not offset the effect of the cumulative bias in earlier parts of the year, resulting in the last three months of the true savings being outside of the 99% prediction interval for cumulative savings.

Lastly, we explore how load forecast bias propagates to and compounds in cumulative savings for all 865 real homes. For a given month, we consider the set of prediction intervals for daily savings generated across all homes and compute the PICP. Repeating this for each month, we plot the PICP values (the blue line) in Fig. 5. For each month, we consider the prediction intervals for month-end savings across all homes and plot the corresponding PICP (the orange line). The high daily savings PICP shows that instances of extremely poor calibration in daily savings forecasts are generally rare. Consequently, the disparity in their PICP scores is attributable to the compounding effects of daily savings forecasts bias, originating from load distribution forecast inaccuracies. The results demonstrate how probabilistic load forecasting inaccuracies appear downstream in probabilistic savings forecasts.

IV. Conclusion

In this paper, we presented a regression-based home load model and a residual-based nonparametric probabilistic forecast error model that generates probabilistic load and savings forecasts from electricity rate switching. We show that bias in forecasted load distributions propagate to short-term savings forecasts and compound in cumulative savings forecasts. Future work should assess the application of bias correction to EPRS savings forecasting and explore methods to model the nonstationarity of load due to changes in appliance ownership and behavioral variation including changes in behavior in response to enrollment in the recommended rate plan.

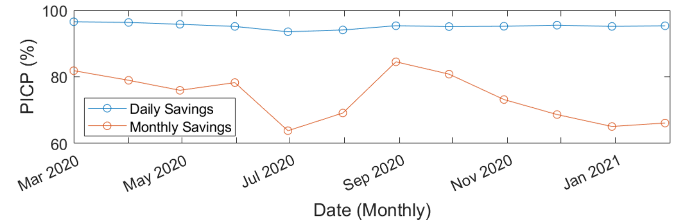

Fig. 5. The rate at which the 99% prediction interval captures daily savings vs. cumulative monthly savings across real 865 homes.

References

[1] F. vom Scheidt, P. Staudt, and C. Weinhardt, "Assessing the economics of residential electricity tariff selection," in *IEEE International Conference on Smart Energy Systems and Technologies*, 2019.

[2] S. J. Darby, "Metering: EU policy and implications for fuel poor households," *Energy Policy*, vol. 49, pp. 98–106, 2012.

[3] P. Zhao, Z. Y. Dong, K. Meng, W. Kong, and J. Yang, "Household power usage pattern filtering-based residential electricity plan recommender system," *Applied Energy*, vol. 298, p. 117191, 2021.

[4] Y. Zhang, K. Meng, W. Kong, and Z. Y. Dong, "Collaborative filtering-based electricity plan recommender system," *IEEE Transactions on Industrial informatics*, vol. 15, no. 3, pp. 1393–1404, 2018.

[5] J. Xie and T. Hong, "Temperature scenario generation for probabilistic load forecasting," *IEEE Transactions on Smart Grid*, vol. 9, no. 3, pp. 1680–1687, 2018.

[6] DTE Energy, "DTE electric company rate book for electric service," pp. 113–115, 2020. [Online]. Available: https://www.michigan.gov/-/media/Project/Websites/mpsc/consumer/rate-books/electric/dte/dtee1cur.pdf?rev=2e38fd3df4f44dc1a8a501aa48975df5

[7] J. Zheng, C. S. Lai, H. Yuan, Z. Y. Dong, K. Meng, and L. L. Lai, "Electricity plan recommender system with electrical instruction-based recovery," *Energy*, vol. 203, p. 117775, 2020.

[8] Y. M. Chung, S. Kang, J. Jung, B. J. Chung, and D. S. Kim, "Residential electricity rate plans and their selections based on statistical learning," *IEEE Access*, vol. 10, pp. 74 012–74 022, 2022.

[9] Z. Xia, H. Ma, T. K. Saha, and R. Zhang, "Consumption scenario-based probabilistic load forecasting of single household," *IEEE Transactions on Smart Grid*, vol. 13, no. 2, pp. 1075–1087, 2021.

[10] S. Arora and J. W. Taylor, "Forecasting electricity smart meter data using conditional kernel density estimation," *Omega*, vol. 59, pp. 47–59, 2016.

[11] T. Hong, J. Wilson, and J. Xie, "Long term probabilistic load forecasting and normalization with hourly information," *IEEE Transactions on Smart Grid*, vol. 5, no. 1, pp. 456–462, 2013.

[12] J. Zhang, Y. Wang, M. Sun, N. Zhang, and C. Kang, "Constructing probabilistic load forecast from multiple point forecasts: A bootstrap based approach," in *IEEE Innovative Smart Grid Technologies-Asia*, 2018, pp. 184–189.

[13] J. Lemos-Vinasco, P. Bacher, and J. K. Møller, "Probabilistic load forecasting considering temporal correlation: Online models for the prediction of households' electrical load," *Applied Energy*, vol. 303, p. 117594, 2021.

[14] J. Mathieu, P. Price, S. Kiliccote, and M. Piette, "Quantifying changes in building electricity use, with application to demand response," *IEEE Transactions on Smart Grid*, vol. 2, no. 3, pp. 507–518, 2011.

[15] T. Hong and S. Fan, "Probabilistic electric load forecasting: A tutorial review," *International Journal of Forecasting*, vol. 32, no. 3, pp. 914–938, 2016.

[16] S. M. Ross, *Introduction to Probability Models (Thirteenth Edition)*. Academic Press, 2024.

[17] NOAA, "Local climatological data," Station WBAN:94847, 2023. [Online]. Available: https://www.ncdc.noaa.gov/cdo-web/datasets/LCD/stations/WBAN:94847/detail