Modelling Weather Effects for Impact Analysis of Residential Time-of-Use Electricity Pricing

Reid Miller[a], Lukasz Golab[a,*], Catherine Rosenberg[a]

[a] *University of Waterloo, 200 University Avenue West, Waterloo, Ontario, Canada N2L 3G1*

Abstract

Analyzing the impact of pricing policies such as time-of-use (TOU) is challenging in the presence of confounding factors such as weather. Motivated by a lack of consensus and model selection details in prior work, we present a methodology for modelling the effect of weather on residential electricity demand. The best model is selected according to explanatory power, out-of-sample prediction accuracy, goodness of fit and interpretability. We then evaluate the effect of mandatory TOU pricing in a local distribution company in southwestern Ontario, Canada. We use a smart meter dataset of over 20,000 households which is particularly suited to our analysis: it contains data from the summer before and after the implementation of TOU pricing in November 2011, and all customers transitioned from tiered rates to TOU rates at the same time. We find that during the summer rate season, TOU pricing results in electricity conservation across all price periods. The average demand change during on-peak and mid-peak periods is -2.6% and -2.4% respectively. Changes during off-peak periods are not statistically significant. These TOU pricing effects are less pronounced compared to previous studies, underscoring the need for clear, reproducible impact analyses which include full details about the model selection process.

Keywords: time-of-use pricing, effect of weather on residential electricity demand, regression models

1. Introduction

Pricing schemes intended to reduce peak electricity consumption such as time-of-use (TOU) are becoming tractable as advanced metering proliferates. The Ontario Energy Board established a three-tier TOU pricing scheme with three objectives: *i*) to more accurately reflect the wholesale market cost of electricity in the price consumers pay; *ii*) to encourage electricity conservation across all hours of the day; and *iii*) to shift electricity use from high-demand periods to lower-demand periods (Ontario Energy Board, 2004). Properly evaluating the impact of such policies is critical for policy makers trying to reduce demand, reduce emissions

*Corresponding author: lgolab@uwaterloo.ca

and defer new generating capacity. However, isolating the moderate effects of TOU pricing is challenging in the presence of substantial confounding factors. For example, a mild or extreme summer may skew the estimated impact of TOU pricing if the effects of weather are not adequately modelled.

We observe that there is no consensus in prior work for modelling weather effects and discussion of variable selection criteria is limited. To ensure reliable results, policy makers should insist on clear, reproducible impact analyses which include details of the explanatory variable selection process and justification for any variable transformation used. To help produce such analyses, this paper presents a methodology for modelling the effects of weather on residential demand in the context of pricing policies.

The crux of our methodology is to compare a number of aggregate electricity demand models which have each modelled the effects of weather differently. We use statistical measures of their explanatory power, out-of-sample prediction accuracy, and goodness of fit to select a model that is both well-performing and readily interpretable. After careful analysis, we have chosen a multiple regression modelling structure for its interpretability, tractability, and modularity. To enumerate the possible models, we define three independent components: coincident weather (e.g., incorporating humidity and windchill in addition to temperature), delay or build-up of temperature that household thermal controls react to (e.g., moving average of temperature or cooling/heating degree-hours) and the non-linear relationship of temperature with demand (e.g., piecewise linear and natural spline transformations). We hypothesize that the effect of temperature on aggregate residential electricity demand is non-linear. Furthermore, we hypothesize that past temperature observations and coincident weather observations each provide additional explanatory value.

The second contribution of this paper is an application of the proposed methodology to evaluate the effects of Ontario's mandatory TOU implementation according to two of its stated objectives: energy conservation and shifting consumption out of peak demand periods. We use a smart meter dataset of over 20,000 households in southwestern Ontario, Canada that is particularly suited to our analysis. It has an adequate numbers of observations before and after the implementation of TOU pricing. Furthermore, the local distribution company transitioned all customers from tiered rates to TOU rates at a single point in time, meaning that there is no uncertainty introduced by a staggered TOU billing roll-out. Though the sample size and rate transition are positive assets of the dataset, the sample time period does not include adequate pre-TOU observations during the winter rate season to assess its effectiveness. Given this limitation, we present results only for the summer TOU rate season and make conclusions in that context.

2. Prior Work

A literature review performed by Newsham and Bowker (2010) discusses the impacts of three types of dynamic pricing pilots: critical peak pricing, time-of-use, and peak time rebates. Their review includes 13 TOU pilot studies conducted after 1997. They conclude that basic TOU pricing programs like Ontario's can expect to see residential on-peak demand change by -5%. An earlier TOU literature review by Faruqui and Sergici (2010) covering 12 TOU pilot studies concluded that TOU pricing induces a -3% to -6% change in residential on-peak demand. From 2010 onwards, there have been several impact studies of mandatory TOU pricing. We summarize these recent studies as well as several of the older ones in Table 1.

Our first observation is that results from opt-in experiments and pilot studies such as Hydro One (2008), Lifson and Miedema (1981), Ontario Energy Board et al. (2007) and Train and Mehrez (1994) are often more pronounced than mandatory studies such as Faruqui et al. (2013b), Navigant Research and Newmarket-Tay Power Distribution (2010) and Navigant Research and Ontario Energy Board (2013). Our second observation is that most studies in our review either have a pronounced demand shift from on-peak to off-peak hours or conservation across all hours. Only two subsets of one study by Jessoe et al. (2013) showed the opposite effect. Finally, we observe that the tiered roll-out of TOU to high-use customers first, analyzed by Jessoe et al. (2013), showed substantial flexibility to shift demand.

Across these TOU studies, we observed many different techniques being used to model weather. When deciding on which modelling techniques to consider in our methodology, we broadened our literature review to residential electricity demand analysis in general. Table 2 summarizes this broadened literature review, grouping prior work by the technique used to *transform* temperature observations. An explanatory variable transformation is a mathematical process that creates derived values from observed values. For example, a series of dry-bulb temperature observations may be transformed using humidity and wind chill to become a series of perceived temperatures. The derived variable would be used as input to the modelling procedure in place of the observed variable.

We define *coincident weather* to be measurable weather phenomena which coincide with temperature observations. For example, the humidity observed at time i is coincident with dry-bulb temperature observed at time i. Several studies transform temperature by taking humidity into account via the *temperature humidity index* (Faruqui et al., 2013a; Navigant Research and Ontario Energy Board, 2013), the Canadian *Humidex* (Faruqui et al., 2013b), or by incorporating humidity into some other transformation of temperature (Mountain and Lawsom, 1992). Humidity may have a direct effect on load via dehumidification equipment, or an indirect effect via human perception and comfort levels. Wind speed has also been incorporated into temperature transformations (Friedrich et al., 2014; Mountain and Lawsom, 1992). Wind may reduce electricity demand if customers choose to cool their home by leaving windows open during transition seasons. Wind chill may also affect perception of winter outdoor temperatures, inclining a customer to stay indoors.

Temporal transformations account for the delay between when an outdoor temperature occurs to when its effects are felt within a customer's home. *Heating degree-days* and *cooling degree-days* are derived values used to measure the prolonged heating and cooling requirements of a home over time. They have been extended to *heating degree-hours* and *cooling degree-hours*, derived by summing the difference between recent observations and a selected temperature break point. For modelling long-term and mid-term analysis horizons, heating and cooling degree-days are sufficient (Pardo et al., 2002; Cancelo et al., 2008). Heating and cooling degree-hours are better suited to the analysis of short-term and mid-term horizons (Navigant Research and Newmarket-Tay Power Distribution, 2010).

Harvey and Koopman (1993) considered *lagged* hours of temperature observations in early models of their study. Mountain and Lawsom (1992) used a four-hour *moving average* of recent temperatures as a component of the space heating index used in their model. Friedrich et al. (2014) refined work by Bruhns et al. (2005) to account for thermal transfer inertia. The authors define an exponentially weighted moving average filter to be the smoothed temperature.

Moral-Carcedo and Vicéns-Otero (2005) describe a single temperature break point as a *switching regression* to model temperature's non-linear relationship with electricity demand. The coefficient found for temperatures below the break point represents household heating effects. The coefficient for temperatures above the break point represents cooling effects. It is used by Faruqui et al. (2013b), Navigant Research and Newmarket-Tay Power Distribution (2010), and Navigant Research and Ontario Energy Board (2013). Lifson and Miedema (1981), and Train and Mehrez (1994) also use switching regression, but the lower region has a slope of zero because households in their regions of study have no heating effects.

Intuitively, the boundary between heating and cooling effects is not an abrupt break. When subjected to moderate temperatures, occupants may not heat or cool their home. Each household will heat or cool their home at different temperatures, resulting in a smoothed transition region when data is analyzed in aggregate (Bruhns et al., 2005; Friedrich et al., 2014; Moral-Carcedo and Vicéns-Otero, 2005). Cancelo et al. (2008) note that extreme low temperatures and extreme high temperatures exhibit saturation of heating and cooling effects. At these temperatures, all household thermal controls such as space heaters, electric baseboard heating, fans, or air conditioning are working constantly.

Regression splines, a widely-used explanatory variable transformation in econometric literature, are capable of modelling the smooth transitions between heating effects, mid-temperatures, cooling effects, and saturation plateaus at temperature extremes (Engle et al., 1986; Harvey and Koopman, 1993). The regression spline transformation first divides the range of temperatures into a number of regions. Within each region, a polynomial function is fit to the data and constraints may be placed on the polynomial functions to connect them at the region boundaries.

3. Data Description

The smart meter dataset used in this paper was provided by a local distribution company in southwestern Ontario. The observations occur over a period of 20 months, from March 1, 2011 through October 17, 2012. The switch from a seasonal, flat pricing scheme to TOU pricing occurred on November 1, 2011. The TOU rates, illustrated in Figure 1, are comprised of three price levels: off-peak, mid-peak and on-peak. Summer off-peak hours are 7:00pm through 6:59am (overnight) at 6.5¢/kWh. Mid-peak hours are 7:00am through 10:59am and 5:00pm through 6:59pm at 10¢/kWh. On-peak hours are 11:00am through 4:59pm at 11.7¢/kWh. All hours of weekends and holidays are off-peak rates.

List of Symbols	
N	the number of hours in the sample period
J	the number of residential smart meters (i.e., households) in the sample
$\boldsymbol{\tau}$	the $N \times 1$ vector of hourly, dry-bulb temperature observations
$\boldsymbol{\tau}'$	the intermediate $N \times 1$ vector resulting from the transformation of $\boldsymbol{\tau}$ incorporating coincident weather observations
$\boldsymbol{\tau}''$	the intermediate $N \times 1$ vector resulting from the transformation of $\boldsymbol{\tau}'$ incorporating past observations. It represents temperature's effects over time
$\boldsymbol{\Upsilon}$	the $N \times J$ matrix of hourly electricity demand per household
$\mathbf{Y}$	the $N \times 1$ vector of hourly, aggregate residential electricity demand
$\mathbf{X}$	the temporal explanatory variable transformation matrix
$\mathbf{V}$	the price explanatory variable transformation matrix
$\mathbf{T}$	the weather explanatory variable transformation matrix
$\hat{\mathbf{Y}}$	the $N \times 1$ vector representing the model's estimate of $\mathbf{Y}$
$\hat{\beta}_0$	the estimated intercept term from which all other coefficients are offset
$\hat{\boldsymbol{\beta}}$	the vector of coefficient estimates for temporal explanatory variables in $\mathbf{X}$
$\hat{\boldsymbol{\omega}}$	the vector of coefficient estimates for price explanatory variables in $\mathbf{V}$
$\hat{\boldsymbol{\theta}}$	the vector of coefficient estimates for weather explanatory variables in $\mathbf{T}$

The data contains hourly smart meter readings from 23,670 residential customers across a four-city service region. We removed 3,100 meters with customer account changes (e.g., tenant changes). Additionally, we performed a data cleaning step to remove extreme outliers. The maximum short-term overloading of a distribution transformer is 300% of its nameplate rating (IEEE Standards Association, 2012, Section 8.2.2). Using this as a guideline, 14 smart meters had an hourly reading that violated the maximum short-term overloading capacity of the transformer they were connected to and hence, were removed from the sample.

The remaining sample contains $J = 20,556$ smart meter time series for study, each with up to $N = 14,328$ data points (the number of hours in our sample period). Individual meter readings were stored with < 1Wh precision. Missing observations were stored as zero values, and nearly all meters have at least a few missing observations over the course of the sample period. 0.46% of the individual readings were missing. Often, a meter's missing values occur as irregularly positioned gaps lasting multiple hours, such that data interpolation is not suitable. We consider the data in aggregate by deriving the average household demand in each hour from all households. Let the variable $\mathbf{\Upsilon}$[1] represent the $N \times J$ matrix of household smart meter readings from March 1, 2011 through October 17, 2012. $I(\mathbf{\Upsilon}_{i,j} > 0)$ is an indicator function that returns 1 if there exists a reading during hour i for meter j. As Eq. (1) is evaluated from $i = 1, ..., N$, an $N \times 1$ vector $\mathbf{Y}$ representing the aggregate electricity demand for each hour of the sample period will be created.

$$\mathbf{y}_i = \frac{\sum_{j=1}^{J} \mathbf{\Upsilon}_{i,j}}{\sum_{j=1}^{J} I(\mathbf{\Upsilon}_{i,j} > 0)}, \quad i = 1, ..., N \tag{1}$$

The aggregate electricity demand observations fall in the range 0.49 kWh–3.54 kWh and are approximately lognormally distributed with mean 1.18 kWh and median 1.03 kWh. We use the vector $\mathbf{Y}$ as the response variable for the remainder of this study, plotted over time in Figure 2. Notice the summer air conditioning demands during the summer months and less noticeable heating effects during winter.

We also obtained the corresponding hourly weather data from two nearby Environment Canada (2015a) monitoring stations. Weather observations were paired with each meter by selecting the nearest monitoring station, all within 5-25 kilometres. Hourly observations recorded are dry-bulb temperature, relative humidity, dew point, wind direction, wind speed, visibility, atmospheric pressure, humidex, wind chill and a weather condition description. We define $\boldsymbol{\tau}$ to be an $N \times 1$ vector of hourly temperature observations averaged from the two weather stations, weighted by the number of meters reporting near that station each hour. The two summers are not drastically different from one another, as shown by key summary statistics in Table 3. Summer 2012 had a slightly higher median drybulb temperature of 20.2°C compared to 19.1°C in 2011.

Throughout Section 4.3, $\boldsymbol{\tau}$ will be used as input to temperature transformation functions which create a matrix $\mathbf{T}$ of dimension $N \times P_{weather}$, where $P_{weather}$ is the number of columns in $\mathbf{T}$, determined by the variable transformation applied. Left untransformed, $\mathbf{T} = \boldsymbol{\tau}$.

[1] We use regular font for scalar variables, and bold font for vector and matrix variables.

4. Methodology for Modelling the Effects of Weather

We use a multiple regression model shown in Eq. (2) to represent electricity consumption as a function of time, price and weather related variables. Let $\hat{\mathbf{Y}}$ be an $N \times 1$ vector representing the model's estimate of $\mathbf{Y}$. Let $\hat{\beta}_0$ be the estimated intercept term. We store the explanatory variables using three matrices $\mathbf{X}$, $\mathbf{V}$ and $\mathbf{T}$ which represent time, price and temperature transformations respectively. The effects of these explanatory variables are represented by the coefficient estimate vectors $\hat{\boldsymbol{\beta}}$, $\hat{\boldsymbol{\omega}}$ and $\hat{\boldsymbol{\theta}}$ fit using ordinary least squares.

$$\hat{\mathbf{Y}} = \hat{\beta}_0 + \mathbf{X}\hat{\boldsymbol{\beta}} + \mathbf{V}\hat{\boldsymbol{\omega}} + \mathbf{T}\hat{\boldsymbol{\theta}} \tag{2}$$

Our treatment of time and price explanatory variables, selecting categorical variables for inclusion using *forward selection* and analysis of variance (ANOVA), follows well-established statistical learning methods (James et al., 2013, Ch.6). ANOVA performs a hypothesis test comparing two models. The null hypothesis is that the less-complex model with fewer explanatory variables is sufficient to describe the response. The alternate hypothesis is that a more complex model is required. ANOVA tests whether the variance explained by an added explanatory variable or interaction is significantly different from the original model . We direct the reader to (Faraway, 2002, Ch. 10) for full details of the formulation and use of ANOVA. During forward selection, we begin with the null model, which contains the intercept β_0 but no explanatory variables in $\mathbf{X}$, $\mathbf{V}$, or $\mathbf{T}$. We then fit a number of alternate models, each with a single explanatory variable added to $\mathbf{X}$ or $\mathbf{V}$. The explanatory variable which results in an alternate model with the lowest residual sum of squares is added to the null model. The process of adding explanatory variables one at a time is continued until the ANOVA stopping condition is met. Some of the variables that were considered during forward selection but were ultimately ruled out are: weather description, visibility, day-of-week, and a schoolyear indicator.

The remainder of this section presents the details of our methodology. The first step is to select time-dependent variables (Section 4.1). The second step is to select price-related variables (Section 4.2). The third step is to select and justify a model for weather effects, which comprises the bulk of our effort and contribution (Sections 4.3 and 4.4).

4.1. Time-Related Variables

We model hour-of-day as a categorical variable with 24 terms, represented in $\mathbf{X}$ by 23 sparse columns with indicators for each hour. Figure 3 shows a box plot of electricity demand grouped by hour-of-day and exhibits the expected patterns of user activity within the home. People use less electricity in the middle of the night from 01:00–06:00 and are most active in the evening from 18:00–22:00.

Residential electricity demand also differs by day-of-week and on holidays. We are able to achieve a high level of explanatory power using only one degree of freedom by defining a *working day* indicator, similar to Møller Andersen et al. (2013); Moral-Carcedo and Vicéns-Otero (2005), such that weekends and holidays are non-working days.

The use of a working day indicator allows for meaningful variable *interactions* to be fit. The main effects of each explanatory variable represent deviation from the sample mean and the two-way interactions represent deviation from their main effects. An interaction between two categorical factors such as hour-of-day and working day is a sparse matrix with indicators for each unique combination of two variables not represented by their main effects. For example, the baseline for hour-of-day is 00:00 and the baseline for working day is `working_day=FALSE`. If observation i occurs at 00:00 on a non-working day, neither variable's main effects will be added to $\hat{\beta}_0$. If observation i is 07:00 on a non-working day, only the coefficient estimate for 07:00 will be added to $\hat{\beta}_0$ (i.e., main effects). If observation i is 07:00 on a working day, an interaction effect (denoted by 07:00 $\times$ `working_day=TRUE`) is added to $\hat{\beta}_0$ representing the deviation from the main effects of each variable. An example of working day $\times$ hour-of-day interaction coefficient estimates is given in Table 4. Aggregate electricity demand begins earlier on working days, indicated by a positive coefficient estimate that is of noticeable effect size and has a statistically significant p-value. This is likely caused by residential customers preparing for work around 07:00 or 08:00 on working days. 10:00 through 17:00 on working days has a negative coefficient estimate, likely because many residential customers are away at work.

As suggested by Figure 2, there are clear seasonal patterns during summer and winter months. Fitting a model with a categorical explanatory variable for month is statistically significant and increases $Adjusted\, R^2$. However, our goal is to evaluate temperature transformations used to generate $\mathbf{T}$. Any explanatory variable that is *collinear* with the temperature transformation matrix masks its effects, meaning that the estimated effects of two explanatory variables increase and decrease together. We check for collinearity using *variance inflation factor* (VIF) (Fox and Weisberg, 2011). Table 5 shows VIF values when a categorical variable for month is considered; a $VIF > 5$ indicates collinearity (James et al., 2013). Using this measure, we determine that addition of month masks the effects of temperature. For this reason we do not include month as a categorical variable.

As a result of forward selection and the justification process described above, we arrive at a desired set of temporal explanatory variables in $\mathbf{X}$ (we define the notation $\mathbf{x}_{\bullet,p}$ to represent the pth column and all rows of $\mathbf{X}$; this same notation will be used with other matrices going forward).

$\mathbf{x}_{\bullet,p=1}$ through $\mathbf{x}_{\bullet,p=23}$ are hour-of-day indicators representing 01:00 through 23:00.

$\mathbf{x}_{\bullet,p=24}$ is a working day indicator.

$\mathbf{x}_{\bullet,p=25}$ through $\mathbf{x}_{\bullet,p=48}$ are indicators representing the hour-of-day $\times$ working day interaction.

4.2. Price-Related Variables

We use two pricing categorical variables. The first is a TOU billing indicator. The second is a categorical variable representing the local distribution company's billing seasons: summer or winter. We were concerned that similar to the month categorical variable, the utility rate seasons might also be collinear with temperature. However, Table 6 shows that the addition of pricing variables are not collinear with a temperature transformation as we iterate through transformations of $\mathbf{T}$.

We will later saturate the price explanatory variable matrix $\mathbf{V}$ with interactions in our TOU case study in Section 6 after finding a suitable temperature transformation. To summarize, the explanatory variables included in $\mathbf{V}$ are:

$\mathbf{v}_{\bullet,p=1}$ is a utility rate season indicator representing summer and winter rates.

$\mathbf{v}_{\bullet,p=2}$ is a TOU active indicator representing whether customers are billed according to flat rates or TOU rates.

4.3. Weather-Related Variables and Transformations

To select a weather effects model, we define three steps of temperature transformations which are used in conjunction with one another to generate variations of $\mathbf{T}$.

1. **Coincident Weather Transformations**: dry-bulb temperature or *feels like* temperature

2. **Temporal Transformations**: current observation, lagged observations or moving average

3. **Non-Linear Transformations**: switching regression or natural cubic splines

Our methodology iterates over all combinations of temperature transformations listed above. Each iteration uses a different combination of transformation functions to generate the temperature transform matrix $\mathbf{T}$ while holding the matrices $\mathbf{X}$ and $\mathbf{V}$ fixed. We begin each temperature transformation with an $N \times 1$ vector of outdoor, dry-bulb temperature observations $\boldsymbol{\tau}$. Algorithm 1 shows the general process for transforming $\boldsymbol{\tau}$ into $\mathbf{T}$.

Algorithm 1 Overview of how temperature transformations are combined to generate the matrix $\mathbf{T}$.

1. Transform dry-bulb temperature observations $\boldsymbol{\tau}$ into the vector $\boldsymbol{\tau}'$ using coincident weather observations.
2. Transform the vector $\boldsymbol{\tau}'$ into the vector $\boldsymbol{\tau}''$ using a transformation which incorporates past observations. This transformation represents temperature's effects over time.
3. Finally, use the vector $\boldsymbol{\tau}''$ as input into a transformation which models the non-linear relationship between $\boldsymbol{\tau}''$ and aggregate electricity demand $\mathbf{Y}$. The result of this third step is the matrix $\mathbf{T}$ used in the multiple regression model.

In Section 4.3.4, we add two complex transformations to our comparison which violate Algorithm 1: the heating/cooling degree-hour transformation and the exposure-lag-response transformation. Both transformations combine algorithm steps two and three, transforming $\boldsymbol{\tau}'$ directly to the matrix $\mathbf{T}$.

4.3.1. Coincident Weather Transformations

During the first step of Algorithm 1, relative humidity and wind speed are used to transform dry-bulb temperature observations to a *feels like* temperature comprised of heat index and wind chill values where applicable. Algorithm 2 defines the *feels like* transformation.

Algorithm 2 The *feels like* temperature transformation. Formulation of heat index is described by Rothfusz (1990) and wind chill formulation is described by Environment Canada (2015b).

if $\boldsymbol{\tau}_i > 27$ **and** $Relative\, Humidity_i > 40\%$ **then**
  $\boldsymbol{\tau}'_i = Heat\, Index_i$
else if $\boldsymbol{\tau}_i \leq 10$ **and** $Wind\, Speed_i > 4.8kph$ **then**
  $\boldsymbol{\tau}'_i = Wind\, Chill_i$
else
  $\boldsymbol{\tau}'_i = \boldsymbol{\tau}_i$
end if

If $\boldsymbol{\tau}$ is left untransformed during this step, then $\boldsymbol{\tau}'$ would remain a vector of dry-bulb temperature observations such that $\boldsymbol{\tau}' = \boldsymbol{\tau}$.

4.3.2. Delayed Effects of Temperature

We also need to account for the delay between when an outdoor temperature occurs to when its effects are felt within a customer's home. To assess the importance of past temperature in predicting present electricity consumption, Table 7 shows the correlation coefficient of 0-12 lags of dry-bulb temperature $\boldsymbol{\tau}$ with $\mathbf{y}_i$. The correlation of $\mathbf{y}_i$ with past temperatures suggests that there may be an underlying temporal process interacting with temperature.

The *lagged observation* transformation shown in Eq. (3) considers the possibility that temperature's effects on electricity demand may be delayed by a number of hours ℓ, also known as lags. The cause for this delay may be the time it takes an outdoor temperature to pass through a building's insulation. After the time delay, the household's thermal controls react. This interim transformation vector has ℓ fewer rows than $\boldsymbol{\tau}'$ used as input, requiring that rows $i = 1, ..., \ell$ must also be removed from $\mathbf{Y}$, $\mathbf{X}$ and $\mathbf{V}$.

$$\boldsymbol{\tau}_i'' = \boldsymbol{\tau}_{i-\ell}', \quad i = (1+\ell), ..., N \tag{3}$$

A *moving average* of recent temperatures shown in Eq. (4) reflects the possibility that household thermal control systems are not reacting only to temperature at time i or some past time $i - \ell$. Instead, thermal control systems may be reacting to a number of recently experienced temperatures. The variable L represents the number of recent temperatures used in the moving average. The moving average transformation vector has $L - 1$ fewer rows than $\boldsymbol{\tau}'$ used as input, requiring that rows $i = 1, ..., \ell$ must also be removed from $\mathbf{Y}$, $\mathbf{X}$ and $\mathbf{V}$.

$$\boldsymbol{\tau}_i'' = \frac{\sum_{\ell=0}^{L-1} \boldsymbol{\tau}_{i-\ell}'}{L}, \quad i = L, ..., N \tag{4}$$

If $\boldsymbol{\tau}'$ is left untransformed during the second step of the temperature transformation Algorithm 1, then the output of the past weather observation transformation would be current observations $\boldsymbol{\tau}'$ such that $\boldsymbol{\tau}'' = \boldsymbol{\tau}'$.

4.3.3. Non-Linear Temperature Effects

The top graph of Figure 4 shows the coefficient estimate $\hat{\boldsymbol{\theta}}$ fit for an untransformed vector of dry-bulb temperature observations. It is clear that temperature's relationship with aggregate electricity demand in our dataset is *non-linear*. The non-linear relationship between temperature and aggregate electricity demand can be approximated by a number of linear regions. This approach is generally referred to as a piecewise linear transformation or linear splines (James et al., 2013). Moral-Carcedo and Vicéns-Otero (2005) describe piecewise linear models with two linear regions as *switching regression*, giving meaning for electricity demand analyses; the break point represents the switch from heating effects to cooling effects. Eq. (5) shows the transformation of $\boldsymbol{\tau}''$ into a column of $\mathbf{T}$ representing heating effects.

$$\mathbf{t}_{i,1} = (\xi_{break} - \boldsymbol{\tau}_i'')_+, \quad i = 1, ..., N \tag{5}$$

Similarly, Eq. (6) shows the transformation of $\boldsymbol{\tau}''$ into a column of $\mathbf{T}$ representing cooling effects.

$$\mathbf{t}_{i,2} = (\boldsymbol{\tau}_i'' - \xi_{break})_+, \quad i = 1, ..., N \tag{6}$$

Let $(x)_+ := max(0, x)$ and let ξ_{break} be a temperature break point estimated empirically (Muggeo, 2003, 2008). The fitted regression line for this switching regression transformation is shown in the middle graph of Figure 4.

Moving beyond piecewise transformations, *regression splines* have been used to first divide the range of temperatures into a number of regions. Within each region, a polynomial function is fit to the data and constraints are placed on the polynomial functions to connect them at the region boundaries, known as knots. Similar to switching regression, the goal of piecewise polynomial transformation is to break $\boldsymbol{\tau}''$ into regions using break points called knots, represented by the $K \times 1$ vector ξ. Let K be the number of knots, such that there are $K + 1$ regions. For each region, a polynomial function is used to transform observations in $\boldsymbol{\tau}''$. The bottom graph of Figure 4 illustrates $K = 3$ knots.

Additional restrictions about the continuity of the polynomial functions at each knot can be added, known as the order of the spline, denoted by M. An order $M = 1$ spline indicates that the polynomial function fit to each region can be discontinuous at the knots. Order $M = 2$ restricts piecewise polynomial functions of adjacent regions to be continuous at their shared knot. $M = 3$ places the additional restriction that the functions' first derivative must be continuous at the knots. $M = 4$ places yet another restriction that the functions' second derivative must be continuous at the knots. We have chosen order $M = 4$ splines, also known as *cubic splines*, which are widely used (Hastie et al., 2005). The first M columns of $\mathbf{T}$ represent the order of the spline (i.e., continuity restrictions), shown in Eq. (7).

$$\mathbf{t}_{i,m} = \boldsymbol{\tau}_i'' \tag{7}$$

The subsequent K columns of $\mathbf{T}$ represent the polynomial function applied to each temperature region, shown in Eq. (8).

$$\mathbf{t}_{i,M+k} = (\boldsymbol{\tau}_i'' - \boldsymbol{\xi}_k)_+^{M-1}, \quad k = 1, ..., K \tag{8}$$

One further refinement, used to address erratic behaviour of polynomials at the extremes where few observations exist, is to place additional constraints on the fit of the outer spline regions. *Natural cubic splines* restrict the polynomial functions of the outer regions to be linear beyond the sample boundaries. This added bias at the boundaries is often reasonable considering the sparse number of observations. The bottom graph of Figure 4 illustrates a natural cubic spline fit of aggregate electricity demand to dry-bulb temperature. There are three knots placed at 3 °C, 23 °C and 30 °C, selected empirically using the highest

Adjusted R^2 as the selection criterion. A smooth transition between heating and cooling effects is visible around 17 °C.

If $\boldsymbol{\tau}''$ is left untransformed during the third step of Algorithm 1, then the output of the non-linearity transformation would be a vector of observations generated by the first two transformation steps, such that $\mathbf{T} = \boldsymbol{\tau}''$.

4.3.4. Complex Temperature Transformations

Heating degree-hours (HDH) and cooling degree-hours (CDH) are derived values which represent the build-up of temperature beyond a given threshold during a recent window of time. Similar to switching regression, a temperature break point ξ_{break} is chosen. HDH is determined by summing the number of degrees below ξ_{break} during a window of L recent hours, shown in Eq. (9).

$$\boldsymbol{t}_{i,1} = \sum_{\ell=0}^{L} (\xi_{break} - \boldsymbol{\tau}'_{i-\ell})_+, \quad i = L, ..., N \tag{9}$$

Similarly, CDH is determined by summing the number of degrees above ξ_{break} during a window of L recent hours, shown in Eq. (10).

$$\boldsymbol{t}_{i,2} = \sum_{\ell=0}^{L} (\boldsymbol{\tau}'_{i-\ell} - \xi_{break})_+, \quad i = L, ..., N \tag{10}$$

The resulting $(N - L) \times 2$ transformation matrix $\mathbf{T}$ is a piecewise linear regression, similar to switching regression. Rows $i = 1, ..., L$ from $\mathbf{Y}$, $\mathbf{X}$ and $\mathbf{V}$ must also be removed from the sample. Since CDH and HDH values are approximately linear, we do not fit a model using these values as input to a natural cubic splines transformation.

A finite distributed lag model was initially proposed by Almon (1965) to compute a weighted sum of past explanatory variable effects on a response variable. A more recent implementation of this concept by Gasparrini et al. (2010); Gasparrini (2011) has come to be known as distributed lag non-linear models (DLNM). In the DLNM framework, the effects of weather and its relation with time are represented by the concept of *basis*. It assumes that the effect at time i is a basis that can be expressed as a linear combination of exposure and lag transformations of $\boldsymbol{\tau}'$. These transformations are known as basis functions. For example, the basis function of temperature's effects may be modelled with natural cubic splines and is known as the exposure-response association. The weight of the effect may change with time. The basis function describing effect weights over time is known as the lag-response association. Together, they comprise the basis known as *exposure-lag-response association*.

4.4. Metrics for Model Selection

Having described the space of possible models, we now explain how to choose the best model for the task at hand. For each model, we compute a measure of variance explained, $Adjusted\, R^2$. We also check the value of the Bayesian Information Criterion (BIC). As $Adjusted\, R^2$ increases, BIC's value should decrease. If $Adjusted\, R^2$ decreases and BIC increases or if both $Adjusted\, R^2$ and BIC increase, then added explained variance is not justified by added model complexity (James et al., 2013; Ramsey and Schafer, 2012). Additionally, we compute the Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE) to indicate out-of-sample predictive power. We seek a model that balances explanatory power with out-of-sample predictive power, while being parsimonious and interpretable.

Aside from examining the relationship of each explanatory variable with aggregate electricity demand individually, the residuals remaining after fitting a model to data can provide an indication of underlying issues with the estimated model. The $N \times 1$ vector of residuals should be normally distributed, mean zero and independent of each explanatory variable.

5. Results for Modelling the Effects of Weather

Table 8 shows the results for each weather model. We also include several trivial models for comparison: a null model (i.e., intercept-only) in which $\mathbf{X}$, $\mathbf{V}$ and $\mathbf{T}$ have been omitted; non-temperature explanatory variables only in which $\mathbf{T}$ has been omitted; and dry-bulb temperature without any transformation in which $\mathbf{T} = \boldsymbol{\tau}$. By comparing all combinations of temperature variable transformations and selecting a well-performing model, a substantial amount of variance can be explained by weather. Though combinations of temperature transformation steps each produce incremental improvements, the proportion of variance explained by any temperature transformations is notable.

Our clearest descriptive results pertain to the time delay between observed temperature and its effects within residential households. If an analyst is to use a single temperature observation to explain electricity demand at time i, the temperature observation at time $i-2$ should be used. Of single temperature variables, it also has the highest $Adjusted\, R^2$ and out-of-sample predictive power. We interpret this to mean that residential customer's household thermal controls are reacting to temperatures experienced in the past, not the current hour.

All three temporal transformations which include a window of past observations have high $Adjusted\, R^2$ values and improved out-of-sample prediction accuracy. This suggests that a window of recently-observed temperatures is important to properly describe its relationship with electricity demand. Both CDH/HDH and the moving average transformations showed that a six-hour window of temperature observations yielded the highest $Adjusted\,R^2$. We feel this validates part of our hypothesis, that past hours' temperature observations have an effect on the current hour's electricity demand. Notably, *despite the prevalence of the CDH/HDH metric in literature, the moving average transformation has greater explanatory power and predictive power in our dataset*. This may be caused by the smoothing effect that moving average has on the temperature explanatory variable.

The use of heat index and wind chill as components of *feels like* temperature has greater $Adjusted\,R^2$ than the use of dry-bulb temperature in all cases but two. Our analysis cannot provide additional insight about the underlying process, whether human perception or mechanical. Conversely, *feels like* temperature has less out-of-sample predictive power than dry-bulb temperature. Due to this mixed result, we reject part of our hypothesis. Namely, the part that stated coincident weather observations have an effect on electricity demand. We conclude that the *feels like* temperature transformation has little added value over simply using dry-bulb temperature observations.

Despite the strong assumption of linearity made by the switching regression transformation, it explains untransformed aggregate electricity demand reasonably well using either dry-bulb temperature or *feels like* temperature. When estimating unlagged temperature observations, its $Adjusted\,R^2 \approx 0.85$ is comparable to $Adjusted\,R^2 \approx 0.86$ using natural splines. The temperature breakpoint has a straightforward interpretation in relation to electricity demand. The empirical switching point for dry-bulb temperature in our data is 17.9 °C. Natural cubic splines do provide more flexibility in modelling the temperature's non-linear relationship with aggregate electricity demand and has higher $Adjusted\,R^2$ than switching regression. Both results support part of our hypothesis, that temperature's effects on electricity demand are non-linear, having greater impact at low and high temperature extremes.

Finally, we comment on the exposure-lag-response transformation, which we have not found in electricity demand analysis literature. Its intended purpose, to model the weight of an exposure effect over time, is not easily interpretable when applied to our data sample. Combined with its minimal improvements to explanatory power and prediction accuracy, we do not feel its use is justified.

Based on these results, we select the model that uses dry-bulb temperature, combined with the six-hour moving average and the natural spline transform. This model obtained an $Adjusted\,R^2$ of 0.902. A similar model that uses feels-like temperature instead has a slightly higher $Adjusted\,R^2$ of 0.904 but lower out-of-sample predictive power. Furthermore, a similar model with exposure-lang-response transformation rather than the six-hour moving average has an even higher $Adjusted\,R^2$ of 0.910, but as mentioned above, is not easily interpretable.

Finally, we comment on the residuals of the selected model, which fulfill the first assumption of linear regression analysis, that errors be normally distributed with mean zero. The next step of residual analysis is to assess heteroscedasticity of residuals. The plots in Figure 5 indicate that heteroscedastic and autocorrelation consistent (HAC) standard errors must be used when performing hypothesis tests in the TOU pricing case study (Zeileis, 2004). The first plot shows increased variability of residuals at warm temperatures. The middle plot shows greater variance associated with summer and winter seasons. This is likely a result of temperature and season's collinearity with dry-bulb temperature. This result is supported quantitatively by the Durbin-Watson statistic in Table 8. A Durbin-Watson value < 1 indicates positive serial correlation of residuals Bhargava et al. (1982). The bottom plot illustrates that variance of residuals increases with larger values of the response variable. This too is likely related to temperature. Because warmer temperatures are associated with higher electricity demand, it follows that greater residual variance is associated with higher electricity demand.

6. Methodology for TOU Impact Analysis

In Section 4, we described the methodology for modelling time, price and weather explanatory variables as the matrices $\mathbf{X}$, $\mathbf{V}$ and $\mathbf{T}$. We show Eq. (2) below again for clarity, since the same form will be used during the TOU impact analysis.

$$\hat{\mathbf{Y}} = \hat{\beta}_0 + \mathbf{X}\hat{\boldsymbol{\beta}} + \mathbf{V}\hat{\boldsymbol{\omega}} + \mathbf{T}\hat{\boldsymbol{\theta}}$$

We set the temperature transformation matrix $\mathbf{T}$ to be a dry-bulb, six-hour moving average, natural cubic splines transformation. Because collinearity of temperature and seasonal explanatory variables is not a concern when analyzing the effects of TOU, we add a categorical explanatory variable for month to $\mathbf{X}$, supported by ANOVA, such that the categorical variables in $\mathbf{X}$ are:

$\mathbf{x}_{\bullet,p=1}$ through $\mathbf{x}_{\bullet,p=11}$ are month indicators representing January through December.

$\mathbf{x}_{\bullet,p=12}$ through $\mathbf{x}_{\bullet,p=34}$ are hour-of-day indicators representing 00:00 through 23:00.

$\mathbf{x}_{\bullet,p=35}$ is a working day indicator.

$\mathbf{x}_{\bullet,p=36}$ through $\mathbf{x}_{\bullet,p=58}$ are indicators representing the hour-of-day $\times$ working day interaction.

To model effects associated with TOU pricing with hourly fidelity, the time and temperature matrices $\mathbf{X}$ and $\mathbf{T}$ are then held constant. We will use *backward selection*, ANOVA and HAC standard errors to remove insignificant variables from a saturated matrix $\mathbf{V}$ of explanatory variables related to price. Backward selection starts with all possible explanatory variables in $\mathbf{X}$, $\mathbf{V}$ and $\mathbf{T}$. All two-way and three-way interactions combining a TOU billing indicator, working day, hour-of-day and utility rate season are included. This initial model is also called the *saturated model*. The variable interactions provide the necessary degrees of freedom to explain the effects of TOU billing for each hour of day. We remove variables with the largest p-value (i.e., the least statistically significant variable) one at a time until the analysis of variance stopping condition is met (James et al., 2013). The remaining, significant explanatory variables in $\mathbf{V}$ are used as components of the multiple regression model in a "what if" analysis. We use the results from the "what if" analysis to quantify the change in demand associated with TOU pricing. The categorical variables in $\mathbf{V}$ before backward selection are:

$\mathbf{v}_{\bullet,p=1}$ is a utility rate season indicator representing summer and winter rates.

$\mathbf{v}_{\bullet,p=2}$ is a TOU active indicator representing whether customers are billed according to flat rates or TOU rates.

$\mathbf{v}_{\bullet,p=3}$ through $\mathbf{v}_{\bullet,p=25}$ are indicators representing the hour-of-day $\times$ rate season interaction.

$\mathbf{v}_{\bullet,p=26}$ through $\mathbf{v}_{\bullet,p=48}$ are indicators representing the hour-of-day $\times$ TOU active interaction.

$\mathbf{v}_{\bullet,p=49}$ is an indicator representing the working day $\times$ rate season interaction.

$\mathbf{v}_{\bullet,p=50}$ is an indicator representing the working day $\times$ TOU active interaction.

$\mathbf{v}_{\bullet,p=51}$ is an indicator representing the rate season $\times$ TOU active interaction.

$\mathbf{v}_{\bullet,p=52}$ through $\mathbf{v}_{\bullet,p=74}$ are indicators representing the hour-of-day $\times$ working day $\times$ rate season interaction.

$\mathbf{v}_{\bullet,p=75}$ through $\mathbf{v}_{\bullet,p=97}$ are indicators representing the hour-of-day $\times$ working day $\times$ TOU active interaction.

$\mathbf{v}_{\bullet,p=98}$ through $\mathbf{v}_{\bullet,p=120}$ are indicators representing the hour-of-day $\times$ rate season $\times$ TOU active interaction.

$\mathbf{v}_{\bullet,p=121}$ is an indicator representing the working day $\times$ rate season $\times$ TOU active interaction.

Using backward selection, we remove variables from the saturated matrix $\mathbf{V}$ to create a more parsimonious model. The following interactions cannot be justified by ANOVA and are dropped from $\mathbf{V}$: working day $\times$

rate season × TOU active, hour-of-day × working day × TOU active and working day × TOU active. The model used in this case study yields $Adjusted\, R^2 = 0.935$.

Algorithm 3 "What if" analysis used to quantify the effects of mandatory TOU electricity pricing.

1. Fit a model to the entire sample of data. Our sample runs from March 1, 2011 – October 17, 2012.
2. For the summer utility rate season, select subset where TOU active = `FALSE` (i.e., May 2011 – October 2011).
3. Group the selected observations by working day indicator. For each working day type find the mean electricity demand observations for each hour. These hourly averages for each working day type represent the *observed summer*.
4. Copy the selected sample data from step 2 into a new hypothetical sample of data called the *counterfactual summer*.
5. In the counterfactual summer, change the TOU active indicator from `FALSE` to `TRUE`.
6. Estimate a response vector using the adjusted counterfactual summer from step 5. Because TOU active has been changed to `TRUE`, the coefficients estimated in step 1 will create a response vector as if TOU billing had been active during summer 2011. This estimated response vector is the "what if" analysis.

Algorithm 3 describes our "what if" methodology, similar to that used by Navigant Research and Ontario Energy Board (2013). However, because our sample of data does not have a complete winter utility rate season of TOU active = `FALSE` from November 2010 through May 2011, we are only able to carry out the analysis for summer utility rate season.

7. Results for TOU Impact Analysis

The average demand change during summer on-peak and mid-peak periods is −2.6% and −2.4% respectively. This translates to −0.035kWh (±0.024kWh) per household each hour during on-peak periods and −0.030kWh (±0.024kWh) change during mid-peak periods. Changes during working day and non-working day off-peak periods are −0.9% and −0.6% but are not statistically significant. Table 9 summarizes the hourly effects averaged by TOU price period. Figure 6 shows this same information graphically. The estimated effects of TOU pricing for each hour are plotted over coloured regions representing the three price periods of a summer working day.

The results from Table 9 can be extrapolated to all 20,556 residential customers in the local distribution company's service region. Demand during each on-peak hour would change by −0.72 MWh (±0.49 MWh), mid-peak hours would change by −0.62 MWh (±0.49MWh), off-peak would change by −0.23 MWh (±0.49 MWh), and each hour of non-working days would change by −0.17 MWh (±0.62 MWh).

We study the daily peak-to-average ratio since it is a metric often used by utilities to measure how extreme demand fluctuations are. Each day's peak-to-average ratio is defined as the peak demand for the day divided by the average demand during that day. The average observed peak-to-average ratio for summer 2011 under flat pricing was 1.441. The estimated summer peak-to-average ratio of the counterfactual sample is 1.429. This represents an estimated change of -0.844% to the peak-to-average ratio, with a 95% confidence interval of $\pm 0.6\%$.

The local distribution company's peak hour observed during the pre-TOU summer occurred on Thursday, July 21, 2011 at 18:00 EDT, averaging 3.54 kWh per household. Using estimated demand from the "what if" analysis, on-peak TOU pricing would have reduced the average household consumption during that hour to 3.42 kWh (± 0.03 kWh), a reduction of 3.4%. The most extreme peak-to-average ratio was observed on Tuesday, June 21, 2011 with a value of 1.65. Had TOU pricing been in place that summer, the estimated peak-to-average ratio on that date would have been 1.57, a reduction of 4.8%.

8. Conclusions and Policy Implications

The main policy implication of this paper is the introduction of a methodology that energy researchers and practitioners may use to model residential demand, including the effects of weather, when analyzing the impact of pricing strategies such as TOU. Our methodology evaluates a wide variety of approaches used to cope with the effects of time, weather and price, and selects the best model based on explanatory power, out-of-sample prediction accuracy, interpretability and goodness of fit. These effects can vary greatly by region, so no single residential electricity demand model is universally applicable. For this reason, policy makers should insist on clear, reproducible methodologies such as ours, which include details about variable selection and model assessment measures. The results of an analysis should only be considered reliable if adequate supporting metrics are provided.

Furthermore, policy makers should strive to make more electricity demand data available in order to validate new pricing schemes and conservation programs; our analysis would not have been possible without access to a large smart meter dataset.

The second policy implication stems from our TOU impact analysis. We conclude that TOU helped mitigate peak electricity demand, reducing summer on-peak demand by 2.6%. Our findings are consistent with those of Newsham and Bowker (2010), which estimates that TOU implementations typically see on-peak demand reductions $< 5\%$. However, we observe that the estimated effects in our dataset are less pronounced than initial results elsewhere in Ontario. Faruqui et al. (2013b) estimated first-year results from four Ontario TOU programs with summer on-peak reductions in the range 2.6% to 5.7%. Our result falls at the bottom of that range. Our result is also lower than that of Navigant Research and Ontario Energy Board (2013), which analyzed a sample of 10,000 residential consumers in various locations within Ontario, finding a summer on-peak reduction of 3.3%. It is worth noting that Navigant's 3.3% demand reduction estimate falls within our 95% confidence interval.

Both the slight decrease in peak-to-average ratio and the hourly demand reduction across all TOU price periods indicate that mandatory TOU pricing can achieve electricity conservation. However, analysis of electricity demand shifting is more complex. Table 9 shows that the majority of estimated summer demand reduction occurs in on-peak and mid-peak periods. Change during off-peak periods for both working days and non-working days is minimal. We interpret this to mean that electricity demand is not being shifted to off-peak periods, but is only being conserved. Conservation is focused during on- and mid-peak periods. Given this finding, the local distribution company should adjust its long term forecasts. If conservation is the trend in many other local distribution companies, the province might be able to defer construction of new generation facilities. If demand shifting were more substantial (e.g., an increase during off-peak periods) it would result in a flattened demand curve. If such a trend were to exist in many local distribution companies, then the make-up of generation facilities throughout the province could shift from those with fast ramp rates, such as natural gas or reservoir hydroelectric, to those that are more constant, such as nuclear.

The demand reduction during the off-peak hours of 19:00 through 21:00 during working days is counter-intuitive. When the hours 17:00 through 18:00 during the second mid-peak period are also considered, demand reduction seems focused during the evening after typical work hours. Residential customers may be attempting to conserve electricity, but they may only have flexibility in their after-work household activity. Because Ontario's TOU pricing also applies to commercial customers, it may not be optimally structured around residential demand flexibility. This misalignment was also noted by Adepetu et al. (2013) in their results when studying aggregate provincial data. The province of Ontario could study the impact of placing commercial and residential customers on separate TOU schedules and adjusting rates accordingly. If Ontario's residential TOU rate schedule remains unchanged, there is an opportunity for technology companies in the realm of connected devices. This result suggests that residential customers, unaided by automated devices, have difficulty reacting to TOU rates when outside the home. Devices and software which can incorporate the user's TOU rate schedule could reduce the household electricity bill and associated on-peak emissions to a greater extent.

Because our sample of data is from one local distribution company in south west Ontario, we acknowledge that our results are only directly applicable to that region. Additionally, because we only have data for one summer of before and after the switch to TOU pricing, we cannot assess the effects of TOU pricing during winter rates. We restate our original question in this context: Is Ontario's mandatory TOU policy associated with energy conservation or load shifting during the winter rate season in this local distribution company's service region?

9. Acknowledgements

This research was funded by the Natural Sciences and Engineering Research Council of Canada (Rosenberg) and the Waterloo Institute for Sustainable Energy/Cisco Smart Grid Research Fund (Golab).

References

References

Adepetu, A., Rezaei, E., Lizotte, D., Keshav, S., 2013. Critiquing time-of-use pricing in Ontario, in: 2013 IEEE International Conference on Smart Grid Communications, SmartGridComm 2013, pp. 223–228. URL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6687961, doi:10.1109/SmartGridComm.2013.6687961.

Almon, S., 1965. The distributed lag between capital appropriations and expenditures. Econometrica: Journal of the Econometric Society 33, 178–196. URL: http://www.jstor.org/stable/1911894.

Bhargava, A., Franzini, L., Narendranathan, W., 1982. Serial Fixed Correlation and Effects the Model. Review of Economic Studies 49, 533–549. URL: http://www.jstor.org/stable/2297285.

Breheny, P., Burchett, W., 2015. visreg: Visualization of Regression Models. URL: http://CRAN.R-project.org/package=visreg. r package version 2.1-1.

Bruhns, A., Deurveilher, G., Roy, J.S., 2005. A non-linear regression model for mid-term load forecasting and improvements in seasonality, in: Power Systems Computation Conference, Curran Associates, Inc., Liege. pp. 1–8. URL: http://pscc.ee.ethz.ch/uploads/tx_ethpublications/fp168.pdf, doi:10.1.1.124.891.

Cancelo, J.R., Espasa, A., Grafe, R., 2008. Forecasting the electricity load from one day to one week ahead for the Spanish system operator. International Journal of Forecasting 24, 588–602. URL: http://linkinghub.elsevier.com/retrieve/pii/S016920700800071X, doi:10.1016/j.ijforecast.2008.07.005.

Engle, R.F., Granger, C.W.J., Rice, J., Weiss, A., 1986. Semiparametric Estimates of the Relation Between Weather and Electricity. Journal of the American Statistical Association 81, 310–320. URL: http://www.jstor.org/stable/2289218, doi:10.1080/01621459.1986.10478274.

Environment Canada, 2015a. Environment Canada — Climate: Hourly Data Report. URL: http://climate.weather.gc.ca.

Environment Canada, 2015b. Environment Canada — Wind Chill Calculator. URL: http://www.ec.gc.ca/meteo-weather/default.asp?lang=EnI&n=0F42F92D-1.

Faraway, J.J., 2002. Practical Regression and Anova using R. URL: http://cran.utstat.utoronto.ca/doc/contrib/Faraway-PRA.pdf, doi:10.1016/0360-1315(91)90006-D.

Faruqui, A., Sergici, S., 2010. Household response to dynamic pricing of electricity: a survey of 15 experiments. Journal of Regulatory Economics 38, 193–225. URL: http://link.springer.com/10.1007/s11149-010-9127-y, doi:10.1007/s11149-010-9127-y.

Faruqui, A., Sergici, S., Akaba, L., 2013a. Dynamic pricing of electricity for residential customers: the evidence from Michigan. Energy Efficiency 6, 571–584. URL: http://link.springer.com/10.1007/s12053-013-9192-z, doi:10.1007/s12053-013-9192-z.

Faruqui, A., Sergici, S., Lessem, N., Mountain, D., Denton, F., Spencer, B., King, C., 2013b. Impact Evaluation of Ontario's Time-of-Use Rates: First Year Analysis. Technical Report. Ontario Power Authority. Toronto. URL: http://powerauthority.on.ca/sites/default/files/conservation/Preliminary-Report-First-Year-Impact-Evaluation-of-Ontario-TOU

Fox, J., Weisberg, S., 2011. An R Companion to Applied Regression. Second ed., Sage, Thousand Oaks CA. URL: http://socserv.socsci.mcmaster.ca/jfox/Books/Companion.

Friedrich, L., Armstrong, P., Afshari, A., 2014. Mid-term forecasting of urban electricity load to isolate air-conditioning impact. Energy and Buildings 80, 72–80. URL: http://dx.doi.org/10.1016/j.enbuild.2014.05.011, doi:10.1016/j.enbuild.2014.05.011.

Gasparrini, A., 2011. Distributed lag linear and non-linear models in R: the package dlnm. Journal of Statistical Software 43, 1–20. URL: http://www.jstatsoft.org/v43/i08/.

Gasparrini, A., Armstrong, B., Kenward, M.G., 2010. Distributed lag non-linear models. Statistics in Medicine 29, 2224–2234. URL: http://onlinelibrary.wiley.com/doi/10.1002/sim.3940/pdf, doi:10.1002/sim.3940.

Harvey, A., Koopman, S.J., 1993. Forecasting Hourly Electricity Demand Using Time-Varying Splines. Journal of the American Statistical Association 88, 1228–1236. URL: http://www.jstor.org/stable/2291261, doi:10.2307/2291261.

Hastie, T., Tibshirani, R., Friedman, J., 2005. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. 2nd ed., Springer. URL: http://www.springerlink.com/index/D7X7KX6772HQ2135.pdf.

Hydro One, 2008. Hydro One Networks Inc. Time-of-Use Pricing Pilot Project Results. Technical Report. Hydro One. Toronto. URL: http://www.ontarioenergyboard.ca/documents/cases/EB-2004-0205/smartpricepilot/TOU_Pilot_Report_HydroOne_20080513.pdf.

Hyndman, R.J., Fan, S., 2010. Density forecasting for long-term peak electricity demand. IEEE Transactions on Power Systems 25, 1142–1153. doi:10.1109/TPWRS.2009.2036017.

IEEE Standards Association, 2012. IEEE Guide for Loading Mineral- Oil-Immersed Transformers and Step-Voltage Regulators. URL: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6166928.

James, G., Witten, D., Hastie, T., Tibshirani, R., 2013. An Introduction to Statistical Learning with Applications in R. 4th ed., Springer, New York. URL: http://www-bcf.usc.edu/ gareth/ISL/ISLR Fourth Printing.pdf, doi:10.1007/978-1-4614-7138-7.

Jessoe, K., Rapson, D., Smith, J.B., 2013. The Effect of a Mandatory Time-of-Use Pricing Reform on Residential Electricity Use, in: American Economic Association Annual Meeting, American Economic Association, San Diego. pp. 1–54. URL: http://www.aeaweb.org/aea/2013conference/program/retrieve.php?pdfid=318.

Lifson, D.P., Miedema, A.K., 1981. A comparative analysis of time-of-use electricity rate effects: the Arizona experiment. Energy 6, 403–408. URL: http://linkinghub.elsevier.com/retrieve/pii/0360544281900037 http://www.sciencedirect.com/science/article/pii/0360544281900037, doi:10.1016/0360-5442(81)90003-7.

Maggiore, S., Gallanti, M., Grattieri, W., Benini, M., 2013. Impact of the enforcement of a time-of-use tariff to residential customers in Italy, in: 22nd International Conference and Exhibition on Electricity Distribution (CIRED 2013), Institution of Engineering and Technology, Stockholm. pp. 1–4. URL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6683276, doi:10.1109/EEEIC.2013.6549543.

Mei, L., Qiulan, W., 2011. Study on TOU Price Implantation Effect on Huainan Resident, in: 2011 International Conference on Computer Distributed Control and Intelligent Environmental Monitoring, IEEE. pp. 2221–2224. URL: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5748266, doi:10.1109/CDCIEM.2011.397.

Møller Andersen, F., Larsen, H.V., Boomsma, T.K., 2013. Long-term forecasting of hourly electricity load: Identification of consumption profiles and segmentation of customers. Energy Conversion and Management 68, 244–252. URL: http://dx.doi.org/10.1016/j.enconman.2013.01.018, doi:10.1016/j.enconman.2013.01.018.

Moral-Carcedo, J., Vicéns-Otero, J., 2005. Modelling the non-linear response of Spanish electricity demand to temperature variations. Energy Economics 27, 477–494. URL: http://linkinghub.elsevier.com/retrieve/pii/S0140988305000174, doi:10.1016/j.eneco.2005.01.003.

Mountain, D.C., Lawsom, E.L., 1992. A Disaggregated Nonhomothetic Modeling of Responsiveness to Residential Time-of-Use Electricity Rates. International Economic Review 33, 181–207. URL: http://www.jstor.org/stable/2526990, doi:10.2307/2526990.

Muggeo, V.M., 2003. Estimating regression models with unknown break-points. Statistics in Medicine 22, 3055–3071.

Muggeo, V.M., 2008. segmented: an R Package to Fit Regression Models with Broken-Line Relationships. R News 8, 20–25. URL: http://cran.r-project.org/doc/Rnews/.

Navigant Research, Newmarket-Tay Power Distribution, 2010. The Effects of Time-of-Use Rates on Residential Electricity Consumption. Technical Report. Ontario Energy Board. Toronto. URL: http://www.nmhydro.ca/pdf/NMH_TOU_FINAL.PDF.

Navigant Research, Ontario Energy Board, 2013. Time of Use Rates in Ontario, Part 1: Impact Analysis. Technical Report. Ontario Energy Board. Toronto. URL: http://www.ontarioenergyboard.ca/oeb/_Documents/EB-2004-0205/Navigant_report_TOU_Rates_in_Ontario_Part_1_201312.pdf.

Newsham, G.R., Bowker, B.G., 2010. The effect of utility time-varying pricing and load control strategies on residential summer peak electricity use: A review. Energy Policy 38, 3289–3296. URL: http://linkinghub.elsevier.com/retrieve/pii/S0301421510000510, doi:10.1016/j.enpol.2010.01.027.

Ontario Energy Board, 2004. Board Proposal: Regulated Price Plan for Electricity Consumers. URL: http://www.ontarioenergyboard.ca/documents/cases/EB-2004-0205/development/rpp_proposal_071204.pdf.

Ontario Energy Board, 2012. TOU Prices. URL: http://www.ontarioenergyboard.ca/OEB/Consumers/Electricity/Electricity%20Prices.

Ontario Energy Board, IBM, eMeter Strategic Consulting, 2007. Ontario Energy Board Smart Price Pilot Final Report. Technical Report. Ontario Energy Board. Toronto. URL: http://www.centerpointenergy.com/staticfiles/CNP/Common/SiteAssets/doc/OSPP Final Report - Final070726.pdf.

Pardo, A., Meneu, V., Valor, E., 2002. Temperature and seasonality influences on Spanish electricity load. Energy Economics 24, 55–70. doi:10.1016/S0140-9883(01)00082-2.

Ramsey, F., Schafer, D., 2012. The Statistical Sleuth: A Course in Methods of Data Analysis. 3rd ed., Cengage Learning.

Rothfusz, L.P., 1990. The Heat Index Equation. URL: http://www.srh.noaa.gov/images/ffc/pdf/ta_htindx.PDF.

Train, K., Mehrez, G., 1994. Optional time-of-use prices for electricity: econometric analysis of surplus and pareto impacts. The RAND Journal of Economics 25, 263–283. URL: http://www.jstor.org/stable/2555830.

Zeileis, A., 2004. Econometric computing with hc and hac covariance matrix estimators. Journal of Statistical Software 11, 1–17. URL: http://www.jstatsoft.org/v11/i10.

Table 1: Results from prior TOU electricity pricing studies.

Study	**Pilot**	**Mand.**	**Season**	**Total Change (%)**	**On-Peak (%)**	**Mid-Peak (%)**	**Off-Peak (%)**	**Weekend**
Hydro One (2008)	Yes	No	summer	−3.30	−3.70	*NR*	*NR*	*NR*
Lifson and Miedema (1981)	Yes	No	summer	−3.17	−8.84	−3.95	+2.86	*NA*
Ontario Energy Board et al. (2007)	Yes	No	summer	−6.00	−2.40 (*NS*)	*NR*	*NR*	*NR*
Train and Mehrez (1994)	Yes	No	full year	*NR*	−9.02	*NA*	+6.51	*NA*
Jessoe et al. (2013)	No	Yes	summer	−3.14[a]	−6.09[a]	*NA*	−2.00[a]	*NA*
			summer	+0.39[b]	+1.16[b]	*NA*	+0.06[b]	*NA*
			summer	+2.64[c]	+3.11[c]	*NA*	+2.4[c]	*NA*
Faruqui et al. (2013b)	No	Yes	summer	0 to −0.45[d]	−2.60 to −5.70	decrease	increase	*NR*
			winter	0 to −0.45[d]	−1.60 to −3.20	decrease	increase	
Navigant Research and Newmarket-Tay Power Distribution (2010)	No	Yes	full year	−0.66 (*NS*)	−2.80	−1.39	+0.16 (*NS*)	+2.21
Navigant Research and Ontario Energy Board (2013)	No	Yes	summer	0 to −0.10	−3.30	−2.20	+1.20	+1.90
			summer shoulder	*NR*	−2.20	−1.50	+1.50	+1.40
			winter	*NR*	−3.40	−3.90	−2.50	−1.20
			winter shoulder	*NR*	−2.10	−2.30	−1.10	+0.50 (*NS*)
Maggiore et al. (2013)	No	Yes	Jan–Jun	*NR*	−0.83	*NA*	*NR*	*NA*
Mei and Qiulan (2011)	No	Yes	Feb–Dec	increase	increase	*NA*	increase	*NA*

NR – not reported, *NA* – not applicable, *NS* – not statistically significant

[a] high-use customers only, [b] medium-use customers only, [c] low-use customers only, [d] annual

Table 2: Categories of temperature transformations found in prior work, used when modelling residential electricity demand.

Coincident Weather Transformations	
Humidity	Mountain and Lawsom (1992)
Humidex	Faruqui et al. (2013b)
Temperature Humidity Index	Faruqui et al. (2013a); Navigant Research and Ontario Energy Board (2013)
Wind Speed	Friedrich et al. (2014); Mountain and Lawsom (1992)
Temporal Transformations	
Lagged Observations	Harvey and Koopman (1993)
Heating and Cooling Degree-Days	Pardo et al. (2002); Cancelo et al. (2008)
Heating and Cooling Degree-Hours	Navigant Research and Newmarket-Tay Power Distribution (2010)
Moving Average	Mountain and Lawsom (1992)
Weighted Moving Average	Friedrich et al. (2014); Bruhns et al. (2005)
Non-Linear Transformations	
Switching Regression	Moral-Carcedo and Vicéns-Otero (2005); Faruqui et al. (2013b); Navigant Research and Newmarket-Tay Power Distribution (2010); Navigant Research and Ontario Energy Board (2013); Lifson and Miedema (1981); Train and Mehrez (1994)
Linear Regions with Smoothed Transitions	Bruhns et al. (2005); Friedrich et al. (2014); Moral-Carcedo and Vicéns-Otero (2005)
Regression Splines	Engle et al. (1986); Harvey and Koopman (1993)

Table 3: Summary statistics for $\boldsymbol{\tau}$, the weighted average of drybulb temperatures (°C) within the service region during summer 2011 and summer 2012 rate seasons.

Year	Minimum	1st Quartile	Median	Mean	3rd Quartile	Maximum
Summer 2011	0.4	14.0	19.1	18.8	23.8	37.2
Summer 2012	1.3	15.6	20.2	19.8	24.3	37.7

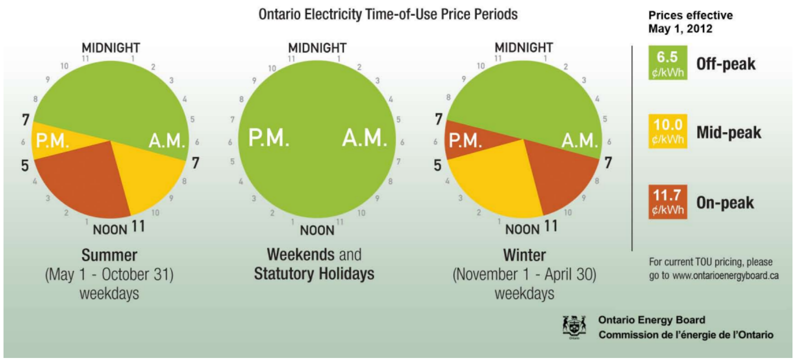

Figure 1: This chart from the Ontario Energy Board (2012) shows the 24-hour schedule in a clock-like format. Off-peak prices are shown in green, mid-peak in yellow, and on-peak in orange. The summer schedule is on the left, weekend schedule in the middle (both seasons), and winter schedule on the right.

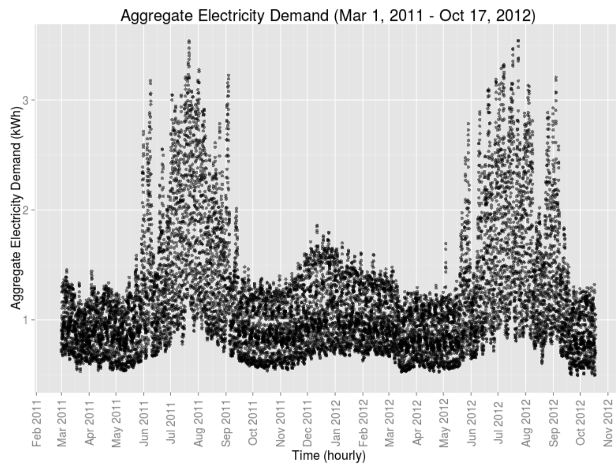

Figure 2: Aggregate residential electricity demand plotted as a function of time. Transparency has been used to give a sense of density.

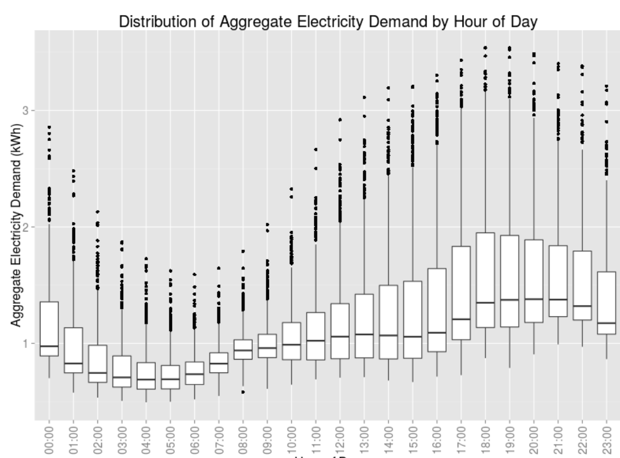

Figure 3: Plot of aggregate electricity demand grouped by hour. Note that this plot contains data from both working and non-working days.

Table 4: Coefficient estimates and p-values illustrating the intuition behind the hour-of-day × working day interaction. Starred p-values are statistically significant. The coefficient estimates will change slightly with each temperature transformation compared, but the sign, intuition and statistical significance remain applicable.

Interaction Term	**Coefficient Estimate**	**p-value**	
01:00×`working_day=TRUE`	0.001	0.9772	
02:00×`working_day=TRUE`	-0.009	0.8642	
03:00×`working_day=TRUE`	0.007	0.8814	
04:00×`working_day=TRUE`	0.017	0.7278	
05:00×`working_day=TRUE`	0.031	0.5377	
06:00×`working_day=TRUE`	0.066	0.1850	
07:00×`working_day=TRUE`	0.135	0.0066	**
08:00×`working_day=TRUE`	0.152	0.0022	**
09:00×`working_day=TRUE`	-0.010	0.8336	
10:00×`working_day=TRUE`	-0.140	0.0048	**
11:00×`working_day=TRUE`	-0.208	0.0000	***
12:00×`working_day=TRUE`	-0.246	0.0000	***
13:00×`working_day=TRUE`	-0.261	0.0000	***
14:00×`working_day=TRUE`	-0.268	0.0000	***
15:00×`working_day=TRUE`	-0.252	0.0000	***
16:00×`working_day=TRUE`	-0.214	0.0000	***
17:00×`working_day=TRUE`	-0.153	0.0020	**
18:00×`working_day=TRUE`	-0.086	0.0817	.
19:00×`working_day=TRUE`	-0.061	0.2205	
20:00×`working_day=TRUE`	-0.025	0.6088	
21:00×`working_day=TRUE`	0.016	0.7501	
22:00×`working_day=TRUE`	0.043	0.3919	
23:00×`working_day=TRUE`	0.038	0.4386	

Table 5: VIF of explanatory variable main effects with a categorical variable for month. *Note: A natural cubic spline transformation of temperature has been used in this example, though similar results are achieved with nearly all temperature transforms discussed in Section 4.3.*

Explanatory Variable(s)	**VIF**	**Degrees of Freedom**
Natural Cubic Splines **T**	8.52	4
Month	7.60	11
Hour-of-Day	1.29	23
Working Day	1.01	1

Table 6: VIF of explanatory variable main effects with addition of utility rate season and a TOU billing indicator. *Note: A natural cubic spline transformation of temperature has been used in this example, though similar results are achieved with nearly all temperature transforms discussed in Section 4.3.*

Explanatory Variable(s)	**VIF**	**Degrees of Freedom**
Natural Cubic Splines $\mathbf{T}$	3.08	4
Rate Season	2.85	1
TOU Active	1.08	1
Hour-of-Day	1.17	23
Working Day	1.00	1

Table 7: Up to 5 lags of dry-bulb temperature are correlated with aggregate electricity demand at levels comparable to dry-bulb temperature at time i.

Lagged Dry-Bulb Temperature	**Correlation with $\mathbf{y}_i$**
$\boldsymbol{\tau}_i$	0.539
$\boldsymbol{\tau}_{i-1}$	0.551
$\boldsymbol{\tau}_{i-2}$	0.558
$\boldsymbol{\tau}_{i-3}$	0.558
$\boldsymbol{\tau}_{i-4}$	0.550
$\boldsymbol{\tau}_{i-5}$	0.533
$\boldsymbol{\tau}_{i-6}$	0.509
$\boldsymbol{\tau}_{i-7}$	0.477
$\boldsymbol{\tau}_{i-8}$	0.440
$\boldsymbol{\tau}_{i-9}$	0.400
$\boldsymbol{\tau}_{i-10}$	0.361
$\boldsymbol{\tau}_{i-11}$	0.328
$\boldsymbol{\tau}_{i-12}$	0.302

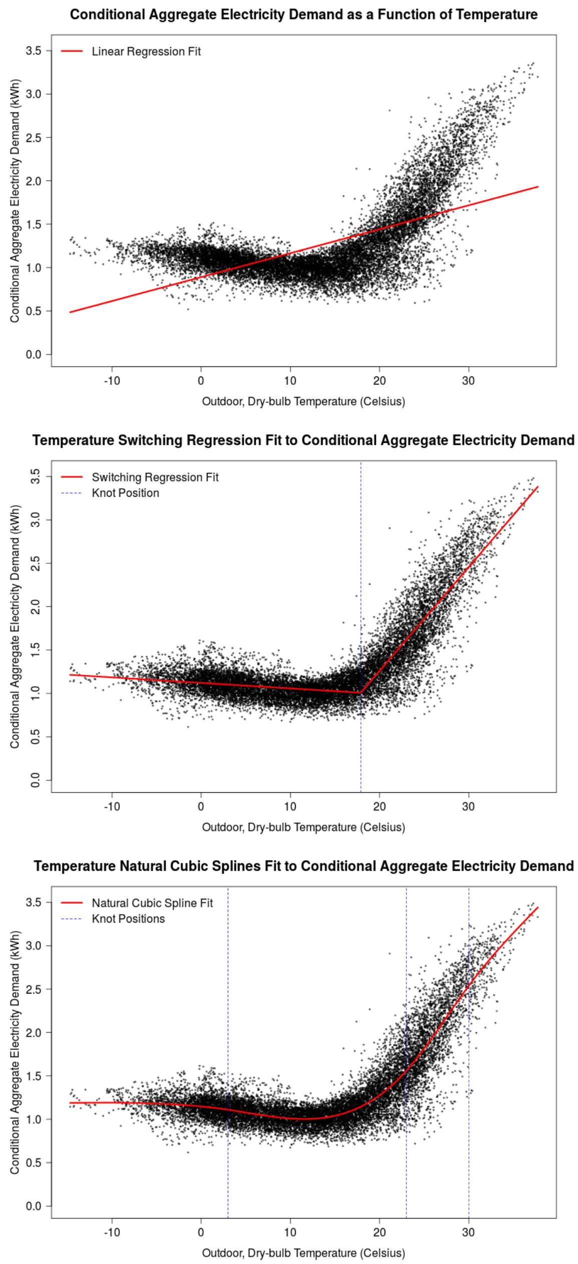

Figure 4: *Top*: Linear regression line fit to untransformed, outdoor, dry-bulb temperature observations. *Middle*: Fitted regression line for switching regression transformation of outdoor, dry-bulb temperature. Temperature break point at 17.9 °C. *Bottom*: Natural cubic splines fit of outdoor, dry-bulb temperature fit to aggregate electricity demand. Knots are placed at 3 °C, 23 °C and 30 °C. All three plots are conditioned for visualization purposes using a process described by Breheny and Burchett (2015).

Table 8: Results of temperature transformation comparison. The first three columns show how temperature transformations from the three categories are combined. *Adjusted* R^2 column is our primary evaluation criterion. *BIC* and Durbin-Watson columns provide secondary measures of model complexity and serially correlated errors. The final two columns report predictive accuracy using average MAE and average MAPE measured using time series cross-validation (Hyndman and Fan, 2010). The model identified in Section 5 as having the greatest explanatory power, out-of-sample prediction accuracy, and interpretability has been bolded.

Temporal Transform	**Weather Transform**	**Non-Linearity Transform**	**Adj.** R^2	**BIC**	**Durbin-Watson**	**Avg. MAE (kWh)**	**Avg. MAPE (%)**
			ideal=1	ideal=low	ideal=2	ideal=0	ideal=0, max=100
Null Model (i.e., intercept only)			0.000	22339.7	0.056	0.432	37.57
Non-Temperature Explanatory Variables Only			0.438	14497.9	0.042	0.397	37.36
None (i-0)	None (Drybulb)	None (Linear)	0.580	10324.2	0.060	0.266	23.51
None (i-0)	None (Drybulb)	Switching Regression	0.854	-4822.3	0.201	0.166	14.20
None (i-0)	None (Drybulb)	Natural Splines	0.862	-5551.9	0.198	0.157	13.19
None (i-0)	Feels Like	Switching Regression	0.857	-5097.0	0.208	0.164	14.00
None (i-0)	Feels Like	Natural Splines	0.862	-5550.1	0.210	0.158	13.35
i-1	None (Drybulb)	Switching Regression	0.875	-7068.5	0.229	0.149	13.05
i-1	None (Drybulb)	Natural Splines	0.884	-8028.3	0.228	0.141	12.02
i-1	Feels Like	Switching Regression	0.878	-7326.9	0.236	0.149	12.82
i-1	Feels Like	Natural Splines	0.884	-8036.4	0.241	0.143	12.19
i-2	None (Drybulb)	Switching Regression	0.881	-7741.3	0.258	0.144	12.71
i-2	None (Drybulb)	Natural Splines	0.889	-8722.6	0.262	0.135	11.64
i-2	Feels Like	Switching Regression	0.883	-7974.6	0.266	0.144	12.46
i-2	Feels Like	Natural Splines	0.890	-8810.8	0.276	0.137	11.76
i-3	None (Drybulb)	Switching Regression	0.873	-6803.5	0.248	0.149	12.96
i-3	None (Drybulb)	Natural Splines	0.880	-7627.2	0.250	0.138	11.72
i-3	Feels Like	Switching Regression	0.875	-7026.7	0.257	0.149	12.73
i-3	Feels Like	Natural Splines	0.882	-7807.8	0.266	0.139	11.81
i-4	None (Drybulb)	Switching Regression	0.853	-4683.5	0.232	0.159	13.50
i-4	None (Drybulb)	Natural Splines	0.859	-5294.1	0.231	0.147	12.05
i-4	Feels Like	Switching Regression	0.855	-4925.3	0.240	0.158	13.28
i-4	Feels Like	Natural Splines	0.862	-5563.3	0.244	0.147	12.13
i-5	None (Drybulb)	Switching Regression	0.825	-2171.8	0.192	0.173	14.29
i-5	None (Drybulb)	Natural Splines	0.830	-2628.2	0.191	0.160	12.84
i-5	Feels Like	Switching Regression	0.828	-2437.6	0.199	0.172	14.11
i-5	Feels Like	Natural Splines	0.834	-2940.9	0.201	0.161	12.94
i-6	None (Drybulb)	Switching Regression	0.790	387.4	0.166	0.189	15.39
i-6	None (Drybulb)	Natural Splines	0.796	-9.7	0.165	0.179	14.34
i-6	Feels Like	Switching Regression	0.794	114.4	0.171	0.188	15.24
i-6	Feels Like	Natural Splines	0.801	-321.1	0.172	0.179	14.41
CDH/HDH (L=6)	None (Drybulb)	Switching Regression	0.895	-9493.4	0.183	0.133	11.53
CDH/HDH (L=6)	None (Drybulb)	Natural Splines	*N/A*	*N/A*	*N/A*	*N/A*	*N/A*
CDH/HDH (L=6)	Feels Like	Switching Regression	0.896	-9629.3	0.184	0.134	11.36
CDH/HDH (L=6)	Feels Like	Natural Splines	*N/A*	*N/A*	*N/A*	*N/A*	*N/A*
Moving Avg. (L=6)	None (Drybulb)	Switching Regression	0.895	-9492.1	0.195	0.139	12.33
Moving Avg. (L=6)	**None (Drybulb)**	**Natural Splines**	**0.902**	**-10537.5**	**0.196**	**0.128**	**10.94**
Moving Avg. (L=6)	Feels Like	Switching Regression	0.897	-9771.0	0.193	0.138	11.99
Moving Avg. (L=6)	Feels Like	Natural Splines	0.904	-10718.9	0.199	0.130	11.16
Lag-Response: Cubic Polynomial (L=6)	None (Drybulb)	Exposure-Response: Switching Regression	0.902	-10388.5	0.197	0.127	10.93
Lag-Response: Cubic Polynomial (L=6)	None (Drybulb)	Exposure-Response: Natural Splines	0.910	-11559.9	0.209	0.118	9.86
Lag-Response: Cubic Polynomial (L=6)	Feels Like	Exposure-Response: Switching Regression	0.901	-10257.2	0.192	0.129	10.88
Lag-Response: Cubic Polynomial (L=6)	Feels Like	Exposure-Response: Natural Splines	0.911	-11732.2	0.213	0.123	10.43

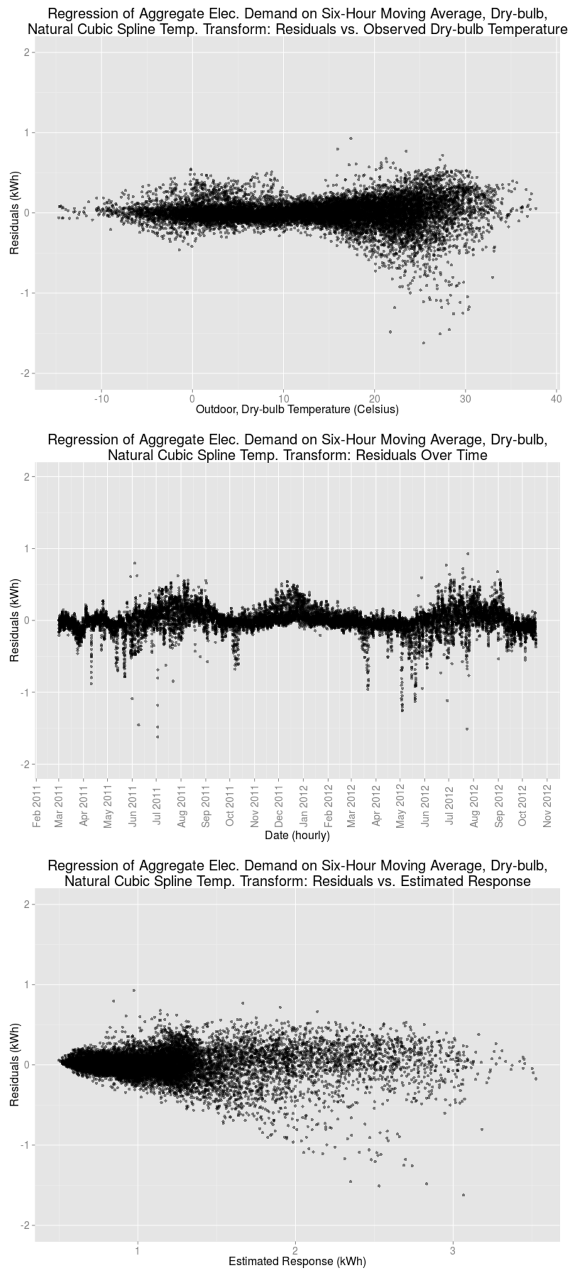

Figure 5: *Top*: Residuals as a function of dry-bulb temperature observations, resulting from a comparison model using dry-bulb, six-hour moving average and natural cubic splines to generate the temperature transformation matrix **T**. *Middle*: Residuals as a function of time, resulting from a comparison model using dry-bulb, six-hour moving average and natural cubic splines to generate the temperature transformation matrix **T**. *Bottom*: Residuals as a function estimated response, resulting from a comparison model using dry-bulb, six-hour moving average and natural cubic splines to generate the temperature transformation matrix **T**.

Table 9: Estimated change in average household electricity demand for each TOU price period.

Summer Price Period	**Hourly Impact (kWh)**	**95% Conf. Interval (kWh)**	**Hourly Impact (%)**	**95% Conf. Interval (%)**
On-Peak	-0.035	± 0.024	-2.641	± 1.819
Mid-Peak	-0.030	± 0.024	-2.403	± 1.933
Off-Peak	-0.011	± 0.024	-0.888	± 1.901
Non-Working Day	-0.009	± 0.030	-0.617	± 2.212

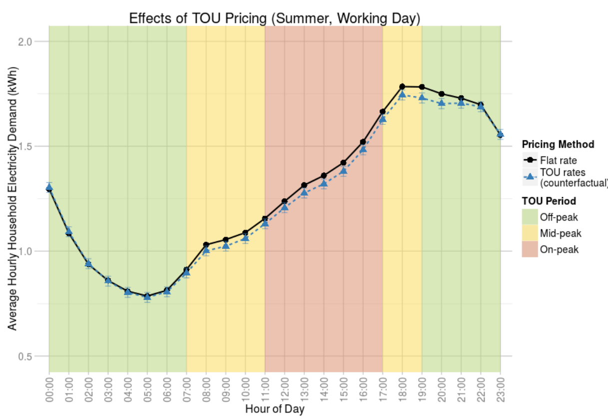

Figure 6: The hourly effects of a "what if" analysis estimated using summer 2011 data from our sample. The observed data is the solid, black line, indicating the mean of observed demand for each hour of working days. The dotted blue line indicates the mean of estimated demand for each hour of working days, had TOU billing been in place. A 95% confidence interval is also plotted for each hour.