Analysis of Weather, Time, and Mandatory Time-of-Use Pricing Effects on Aggregate Residential Electricity Demand

by

Reid Miller

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Electrical and Computer Engineering

Waterloo, Ontario, Canada, 2015

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Residential electricity demand is affected by three types of external factors: weather, time, and the price of electricity. Ontario has mandated time-of-use pricing for all residential customers in the province. We implement a data-driven study using multiple linear regression to quantify the effects of mandatory time-of-use pricing on residential electricity demand in south west Ontario. In order to isolate the effects of this pricing policy on electricity demand, we first account for the combined effects of weather and time.

Our treatment of temporal variables such as month, working days, and hour-of-day is consistent with prior work. However, there is no consensus in prior work for modelling the effects of temperature and weather over time. In temperate regions like Ontario, the relationship between residential electricity demand and temperature is notably non-linear across winter and summer seasons. A mild or extreme summer may skew the estimated impacts of time-of-use pricing if the effects of temperature are not properly accounted for. To address this challenge, we formulate a detailed comparison of existing methods used to transform dry-bulb temperature observations. We consider piecewise linear and natural spline transformations for modelling non-linearity. We also consider coincident weather observations such as humidity, wind chill, and qualitative weather conditions. Finally, we consider variable transformations that take into account the time delay or build-up of temperature that household thermal controls react to. We consider lagged observations, cooling degree-hour, heating degree-hour, moving average, and exposure-lag-response transformations. For all combinations of temperature variable transformations we report the explanatory power, out-of-sample prediction accuracy, and discuss impacts on model interpretability.

Using the results from our temperature transformation comparison, we select a well-performing, descriptive model for use in a time-of-use case study. Ontario's time-of-use pricing policy is evaluated according to two of its stated objectives: energy conservation and shifting consumption out of peak demand periods. We show that during the summer rate season, time-of-use pricing is associated with electricity conservation across all price periods. The average demand change during on-peak and mid-peak periods is -2.6% and -2.4% respectively. Change during working day and non-working day off-peak periods is -0.9% and -0.6% but is not statistically significant. The peak-to-average ratio, a separate metric to measure shifted electricity demand, changed -0.8% under time-of-use pricing from 1.441 to 1.429. These results are consistent with prior time-of-use evaluations carried out within the province, though less pronounced compared to pilot studies.

Acknowledgements

I would like to thank my two supervisors Professor Catherine Rosenberg and Professor Lukasz Golab for their guidance and consultation.

I would also like to thank my two readers Professor Srinivasan Keshav and Professor Mark Crowley for their insightful feedback.

Table of Contents

List of Tables

List of Figures

List of Symbols

$\mathbf{Y}$	An $N \times 1$ vector of hourly aggregate electricity demand, see equation (3.1), page 24
$\hat{\mathbf{Y}}$	An $N \times 1$ vector of aggregate residential electricity demand estimates produced by our multiple regression model, see equation (2.1), page 8
$\mathbf{X}$	An $N \times P_{time}$ matrix of temporal explanatory variables, see equation (2.1), page 8
$\mathbf{V}$	An $N \times P_{price}$ matrix of explanatory variables reflecting electricity price, see equation (2.1), page 8
$\mathbf{T}$	An $N \times P_{weather}$ explanatory variable matrix produced by transforming temperature observations, see equation (2.1), page 8
$\hat{\beta}_0$	The estimated intercept term which other coefficient estimates in the multiple regression model are offset from, see equation (2.1), page 8
$\mathbf{e}$	An $N \times 1$ vector of residuals, see equation (2.2), page 8
$\hat{\beta}$	A $P_{time} \times 1$ vector of coefficient estimates for $\mathbf{X}$, see equation (2.1), page 8
$\hat{\omega}$	A $P_{price} \times 1$ vector of coefficient estimates for $\mathbf{V}$, see equation (2.1), page 8
$\hat{\theta}$	A $P_{weather} \times 1$ vector of coefficient estimates for $\mathbf{T}$, see equation (2.1), page 8
N	The number of observations in our sample, see equation (3.1), page 24
i	An index variable frequently used to iterate over observations in our sample, see equation (3.1), page 24
P	The number of explanatory variables used in the multiple linear regression model, see equation (2.1), page 8

τ An $N \times 1$ vector of hourly dry-bulb outdoor temperature observations, page 21

τ' An interim vector of explanatory variables generated by the coincident weather transformation, page 36

τ'' An interim vector of explanatory variables generated by the past weather observation transformation, page 36

ξ_{break} A knot used in the piecewise linear transformation or degree-hour transformation, see equation (3.5), page 44

ξ A $K \times 1$ vector of knots used in the regression splines transformation, see equation (3.8), page 47

K The number of knots used in the regression splines transformation, see equation (3.8), page 47

M The order of desired continuity in regression splines transformation, see equation (3.7), page 47

ℓ A lag index frequently used when iterating over past observations, see equation (3.3), page 42

L The maximum number of lags used in a temperature transformation such as moving average, degree-hours, or exposure-lag-response, see equation (3.4), page 42

$\pounds$ An $L \times 1$ vector of lag indices, see equation (3.11), page 51

$\mathbf{Z}$ The $N \times U$ exposure-response matrix created by the exposure basis function in the exposure-lag-response transformation, page 50

U The number of explanatory variables in the exposure basis matrix, page 50

$\mathbf{C}$ The $(L+1) \times D$ lag-response matrix created by the lag basis function in the exposure-lag-response transformation, see equation (3.12), page 51

D The degree of a polynomial transformation, see equation (3.12), page 51

$\dot{\mathbf{R}}$ The $N \times U \times (L+1)$ array resulting from adding a lag dimension to $\mathbf{Z}$, see equation (3.13), page 51

$\dot{\mathbf{H}}$ The $N \times (U \cdot D) \times (L+1)$ cross-basis array, see equation (3.14), page 51

$G_{i,j}$ The permutation operator used to generate the cross-basis array $\dot{\mathbf{H}}$, see equation (3.14), page 51

SS_{total} Total sum of squares, see equation (2.3), page 8

$SS_{explained}$ Explained sum of squares, see equation (2.5), page 10

$SS_{residuals}$ Residual sum of squares, see equation (2.6), page 10

$Adjusted\, R^2$ The R^2 value adjusted for the number of model parameters, see equation (2.8), page 11

ANOVA Analysis of Variance, page 12

BIC The Bayesian Information Criterion, see equation (2.9), page 11

DW The Durbin-Watson test for serial correlation, see equation (2.10), page 12

HAC Heteroscedasticity and autocorrelation consistent estimators, page 13

MAE Mean absolute error, see equation (2.12), page 13

$MAPE$ Mean absolute percentage error, see equation (2.12), page 13

HDH Heating degree-hours, see equation (3.9), page 49

CDH Cooling degree-hours, see equation (3.10), page 49

Chapter 1

Introduction

1.1 Motivation

As a ubiquitous component of daily life, the generation and use of electricity represents a substantial portion of greenhouse gas emissions. In April 2005, the Ontario Energy Board created a time-of-use (TOU) electricity pricing framework pursuant to efficiency and environmental conservation goals [52]. In August 2010, the Ontario Energy Board mandated that all electricity distributors establish schedules to bill their customers according to the TOU pricing framework [53]. By August 2012, TOU pricing was in place for 4.4 million residential customers, comprising 91% of residential customers in the province [54].

TOU is a time-based pricing framework which divides the day into rate periods that correspond to well known, aggregate electricity demand patterns. In Ontario there are three rate tiers for residential customers: off-peak, mid-peak, and on-peak. These rate tiers correspond to typical levels of electricity demand throughout the day. During off-peak hours the demand on the electricity grid is the lowest, so consumers are charged at the lowest price rate. Conversely, during on-peak hours consumers pay a higher price for electricity. The Ontario Energy Board introduced this pricing framework with three objectives: *i*) to more accurately reflect the 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 [51]. Lower, more constant demand allows electricity generation facilities to operate in a manner that has less impact on the environment. A flattened demand pattern lowers nuclear generation facilities' risk of difficult and costly curtailment. It also reduces the amount that fossil fuel facilities must ramp generation up and down to match demand.

TOU pricing has been studied regularly since the 1970s, when the U.S. Department of Energy ran several experiments and pilot programs to explore the concept as a conservation-inducing measure during the energy crisis [2]. Most studies have been experimental, pilot projects, implementations that require customer opt-in, or implementations that are mandatory only for customers with average demand above a certain threshold. Ontario is one of only a few large jurisdictions in the world that has mandated TOU pricing for all residential customers.

In this work, enabled by the availability of smart meter data, we explore the effects of weather, time, and price on aggregate residential electricity demand. We are interested in evaluating the effectiveness of Ontario's mandatory TOU pricing policy according to its stated efficiency and environmental objectives. Has increased consumer mindfulness about the amount and timing of electricity usage resulted in conservation across all hours of the day? Has electricity use been shifted from on-peak periods to off- and mid-peak periods? What is the magnitude of demand change at the household and local distribution company levels?

1.2 Problem Statement

In order to quantify the demand change associated with TOU pricing, its effects must be isolated from other external factors relating to time and weather. Treatment of temporal variables is consistent in prior work. Explanatory variables typically reflect hour-of-day, weekdays, weekends, and holidays. However, the manner in which weather has been modelled in prior work varies significantly, depending primarily on the geographic region in which the study was performed and the frequency of observations. Our data from a local distribution company in south west Ontario is hourly, exhibits significant variability associated with weather, and has been subject to two different pricing frameworks.

Therefore, we propose a multiple regression analysis to quantify change in aggregate residential electricity demand associated with TOU pricing. A multiple regression model is chosen for its interpretability and its modularity. Time, weather, and price can each be incorporated and discussed as separate components of the regression.

1.3 Challenges

Real world data is often incomplete and noisy. Because this is a data-driven study, our sample of hourly smart meter readings from a local distribution company is no different.

We first manually label residential meters and remove erroneous data. Using the cleaned data set, we create a normalized sample of hourly residential electricity demand from the aggregate set of meters reporting each hour. This reduces the noise of individual customer time series samples. It also resolves gaps in individual time series data without interpolating values.

A second challenge, the core component of our study, is to properly model the predominant effects of temperature so that the moderate effects of TOU electricity pricing may be isolated and quantified. A mild or extreme summer may skew the estimated impacts of TOU pricing if the effects of weather and time are not adequately modelled.

The final challenge, common to statistical modelling, is variable selection. With a large number of observations, it is possible for a researcher to explain small amounts of variance in the data. Similarly, with a large number of potential explanatory variables, it easy to overfit a model in pursuit of explaining small amounts of variance. We define a parsimonious model by using statistical tests to analyze components of variance, account for heteroscedasticity in model residuals, quantify model complexity, and guard against overfitting.

1.4 Contributions

Our primary focus is to infer the effects of TOU pricing associated with aggregate residential electricity demand.

- There is no clear consensus in existing literature on how to model the effects of weather on electricity demand. We conduct a detailed comparison of existing temperature variable transformations which incorporate coincident weather observations (e.g. humidity and wind chill), past temperature observations, and varying degrees of non-linearity. We conclude that a six-hour moving average of dry-bulb temperature observations transformed using natural cubic splines has significant explanatory power and is easily interpretable.

- In particular, we apply a variable transformation known as *exposure-lag-response association* as a component of our temperature transformation comparison. To the best of our knowledge, exposure-lag-response association has not been used in electricity demand modelling literature. We hypothesize that the variable transformation might provide nuanced insight into how electricity demand is affected by prolonged exposure to cold or warm temperatures and how those effects are weighted over time. Our result, how temperature effects are weighted over time, is difficult to interpret and provides no added explanatory value. Though this is a somewhat negative result, we share our findings because we have not seen exposure-lag-response applied in prior work.

- We conclude that during the summer rate season, TOU pricing is associated with electricity conservation across all price periods. The average demand change during on-peak and mid-peak periods is -2.6% and -2.4% respectively. Change during working day and non-working day off-peak periods is -0.9% and -0.6% but is not statistically significant. The peak-to-average ratio of electricity demand changed -0.8% from 1.441 to 1.429.

Effort has been made to describe the process behind explanatory variable selection by visualizing decisions and reporting supporting statistics. Source code for our analysis is also provided [39].

In addition to our analytical contributions, our data set has several characteristics that make it a novel item of study. First, large samples of hourly smart meter readings are difficult to obtain due to privacy issues. We have a large data set with adequate numbers of observations before and after the implementation of TOU pricing. Second, the local distribution company transitioned all customers from flat rates to TOU rates at a single point in time, meaning that there is no uncertainty introduced by a staggered customer billing roll-out. Third, TOU pricing is mandatory for all residential customers, high-use and low-use. Studying the effects of mandatory TOU pricing within Ontario may provide insight for other regions considering similar mandatory implementations.

1.5 Outline

Chapter 2 first gives background on multiple linear regression, residual analysis, measures of explanatory power, and out-of-sample predictive power. After providing background information on relevant statistical methods, we then review related work. Literature is grouped by techniques used to include coincident weather information, the structure given to temperature over time, and temperature's non-linear relationship with electricity demand.

Chapter 3 first describes our data set provided by a local distribution company in south west Ontario. We also describe a set of weather observations and how they are paired with smart meter readings. Because both data sets are real-world observations, data cleaning is required to make them suitable for study. Each data cleaning step is described and supported.

To properly model the effects of weather on electricity demand, we propose a detailed comparison of existing temperature variable transformations. Descriptions, visualizations, and equations for the fixed and varying components of our multiple regression model are provided. Our comparison generates models using all combinations of temperature transformations falling into three categories: coincident weather, past observations, and non-linearity.

Finally, we provide empirical results from our temperature transformation comparison. The fitted models are primarily evaluated according to their explanatory power and interpretability. Though the primary purpose of our study is not forecasting, we also report the out-of-sample prediction error for each model. Out-of-sample predictive power supports the general applicability of our conclusions and shows that we are not overfitting the data.

Chapter 4 carries out a case study to quantify the effects associated with TOU pricing in our data sample. First, noteworthy TOU pilot studies and mandatory TOU deployments are reviewed. We then use a well-performing model from chapter 3 to carry out a "what-if" analysis to quantify the effects of Ontario's TOU policy on aggregate residential electricity demand.

Chapter 5 reviews conclusions drawn from our temperature transformation comparison and TOU case study. We then list avenues of future work possible by building on this study and our data set.

Chapter 2

Background and Prior Work

2.1 Background

The focus of our study is to describe the functional relationship between hourly aggregate residential electricity demand $\mathbf{Y}$ and a set of explanatory variables. A single-variable, linear relationship between a response variable and an explanatory variable is referred to as *simple linear regression*. *Multiple linear regression* is an extension of simple linear regression which adds multiple explanatory variables, each with its own slope coefficient [31]. This is also simply referred to as *multiple regression* and will be referred to as such going forward. Appendix A provides a review of multiple regression.

We assume that the underlying process driving residential electricity demand can be estimated using a multiple regression model (2.1). The $N \times 1$ vector $\hat{\mathbf{Y}}$ is an estimate of aggregate electricity demand $\mathbf{Y}$ produced by adding the estimated effects of three types of explanatory variables: time, price, and weather. Let $\mathbf{X}$ be an $N \times P_{time}$ matrix of temporal explanatory variables, $\mathbf{V}$ be an $N \times P_{price}$ matrix of explanatory variables reflecting electricity price, and $\mathbf{T}$ be an $N \times P_{weather}$ matrix produced by transforming temperature observations. N is the number of observations in our sample of hourly time series data and each subscript of P is the number of explanatory variables of a given type.

We break the explanatory variables up into separate matrices so that each type may be discussed separately. The temperature transformation comparison described in chapter 3 will hold $\mathbf{X}$ and $\mathbf{V}$ fixed and compare combinations of temperature transformations which produce $\mathbf{T}$. Similarly, the TOU case study described in chapter 4 will hold $\mathbf{X}$ and $\mathbf{T}$ fixed while adding many TOU explanatory variables to $\mathbf{V}$.

$$\hat{\mathbf{Y}} = \hat{\beta}_0 + \mathbf{X}\hat{\beta} + \mathbf{V}\hat{\omega} + \mathbf{T}\hat{\theta} \tag{2.1}$$

Let $\hat{\beta}_0$ be the estimated intercept term which other coefficient estimates are offset from. $\hat{\beta}$ is a $P_{time} \times 1$ vector of coefficient estimates for $\mathbf{X}$. $\hat{\omega}$ is a $P_{price} \times 1$ vector of coefficient estimates for $\mathbf{V}$. $\hat{\theta}$ is a $P_{weather} \times 1$ vector of coefficient estimates for $\mathbf{T}$.

2.1.1 Least Squares

The process of choosing values for $\hat{\beta}$, $\hat{\omega}$, and $\hat{\theta}$ is known as *model fitting* or *training* the model. The best set of values chosen are those that minimize the difference between observations of $\mathbf{Y}$ and the estimate $\hat{\mathbf{Y}}$. The differences between each observed response and estimated response are known as *residuals*, shown in (2.2).

$$\mathbf{e} = \mathbf{Y} - \hat{\mathbf{Y}} \tag{2.2}$$

Let $\mathbf{e}$ be an $N \times 1$ vector of residuals. Figure 2.1 illustrates a generic set of observations, estimates, and residuals. The goal of estimating model coefficients is to find the best set of values that produce an estimated response $\hat{\mathbf{Y}}$ which is as close to all values of the observed response $\mathbf{Y}$ as possible. There are many ways to measure the best set of coefficient estimates. The most common approach is to minimize the *least squares* criterion [31].

First, let us define the *total sum of squares* to describe the amount that observed data varies from the sample mean $\bar{y}$, shown in (2.3). The notation $\mathbf{y}_i$ represents the ith element of the vector $\mathbf{Y}$. This notation for referencing specific elements in vectors will be used going forward.

$$SS_{total} = \sum_{i=1}^{N} (\mathbf{y}_i - \bar{y})^2 \tag{2.3}$$

SS_{total} may be partitioned into two parts shown in (2.4): *explained sum of squares* and *residual sum of squares*.

$$SS_{total} = SS_{explained} + SS_{residuals} \tag{2.4}$$

Explained sum of squares (2.5) represents the amount of deviation from the sample mean explained by the fitted model.

$$SS_{explained} = \sum_{i=1}^{N} (\hat{\mathbf{y}}_i - \bar{y})^2 \tag{2.5}$$

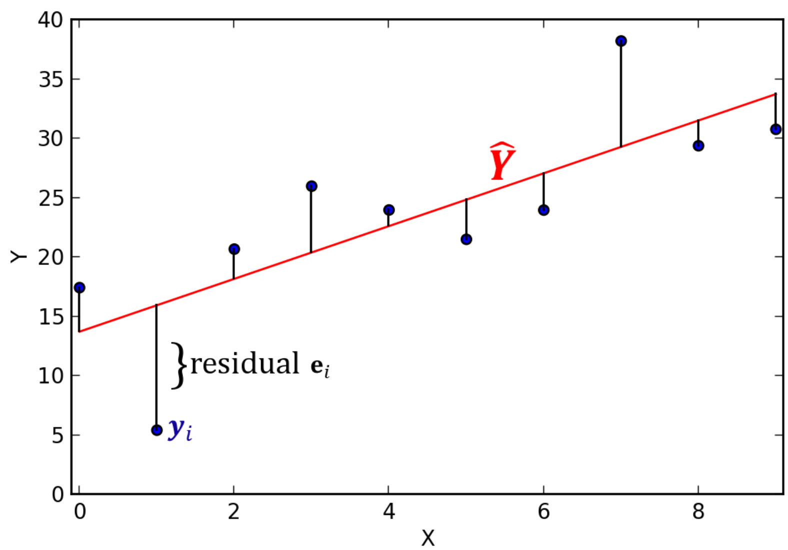

Figure 2.1: Blue points represent a generic set of observations $\mathbf{Y}$. Each point on the plot represents an individual observation $\mathbf{y}_i$ in $\mathbf{Y}$. The red line $\hat{\mathbf{Y}}$ is produced by an estimated model of the true, underlying function. The vertical black lines represent residuals $\mathbf{e}$. (Creative-Commons 3.0 image [25], edited by Reid Miller)

When the value of $\hat{\beta}_0$ is set to $\bar{y}$, the estimated coefficient vectors $\hat{\beta}$, $\hat{\omega}$, and $\hat{\theta}$ explain observations' deviation from the mean. The variance left unexplained is the residual sum of squares, shown in (2.6).

$$\begin{aligned} SS_{residuals} &= \sum_{i=1}^{N} (\mathbf{y}_i - \hat{\mathbf{y}}_i)^2 \\ &= \sum_{i=1}^{N} \mathbf{e}_i^2 \end{aligned} \tag{2.6}$$

$SS_{residuals}$ provides an absolute measure of the lack of fit of the model to observations. To report the explanatory power of a multiple regression model in more familiar terms, the R^2 statistic (2.7) reports the proportion of variance explained on a scale from 0 to 1 [31].

$$R^2 = 1 - \frac{SS_{residuals}}{SS_{total}} \tag{2.7}$$

For details on how the coefficient estimates that minimize (2.6) are computed (i.e. ordinary least squares), the reader is directed to chapter 3 of [26]. Throughout this study the core `stats` package in the statistical computing environment, R, is used to fit multiple regression models [58].

2.1.2 Model Parsimony

The number of explanatory variables that may be used to describe electricity demand is potentially infinite. Many additional explanatory variables could be used such as solar irradiance, wind speed, visibility, indicators for hours surrounding daylight savings, and many more. Each may potentially explain some part of the variance in electricity demand. Adding additional variables, even those weakly associated with the response, will always fit the training data more accurately. R^2 will always increase as more parameters are added to the model [31].

Including too many parameters can have a number of negative effects. First, coefficient estimates associated with too many explanatory variables can become difficult to interpret, especially in the presence of *multicollinearity*. Multicollinearity occurs when several explanatory variables describe the same underlying phenomenon. Second, as more explanatory variables are added, they increase the *degrees of freedom* of the model. If a model has too many degrees of freedom the coefficient estimates may be overfit to irreducible error in the observed data (i.e. noise), rather than the true underlying function. An overfit model will exhibit poor predictive power when it encounters out-of-sample observations.

An overfit model also has less general applicability, due to the fact that it describes the nuances of one particular data sample.

To guard against overfitting, several measures have been defined to punish overly-complex models. They strive to find a balance between model complexity and explanatory power, referred to as *model parsimony*. $Adjusted\, R^2$, shown in (2.8), is the R^2 value adjusted for sample size N and the total number of explanatory variables P. Let the total number of explanatory variables be the combined number of explanatory variables in each matrix such that $P = P_{time} + P_{price} + P_{weather}$.

$$Adjusted\, R^2 = 1 - \frac{SS_{residuals}/(N - P - 1)}{SS_{total}/(N - 1)} \tag{2.8}$$

If there are a large number of observations, then it is possible to detect small effect sizes with large N. When considering the addition of an explanatory variable to the model, if the model has a large number of explanatory variables but only a small decrease of $SS_{residuals}$, then the added model complexity is not justified. $Adjusted\, R^2$ has a value in the range $[0, 1]$ similar to R^2. However, $Adjusted\, R^2$ will always be less than $\leq R^2$.

A second measure of model parsimony is the Bayesian Information Criterion (BIC) shown in (2.9) [59]. It is a unitless value that decreases for models with low test error, but adds a penalty as the number of explanatory variables increase. When comparing two models estimating the same response, the model with a lower BIC value is considered more parsimonious.

$$BIC = N \cdot ln(\frac{SS_{residuals}}{N}) + (P + 1) \cdot ln(N)) \tag{2.9}$$

2.1.3 Explanatory Variable Selection

We use two conceptually-similar methods of selecting explanatory variables for our multiple regression models. The first is called *forward selection*. In forward selection we begin with the *null model*, a model that contains the intercept β_0 but no explanatory variables in $\mathbf{X}$, $\mathbf{V}$, and $\mathbf{T}$. We then fit a number of alternate models, each with a single explanatory variable added. The explanatory variable which resulted in an alternate model with the lowest $SS_{residuals}$ is added to the null model. This method of adding explanatory variables one at a time is continued until the analysis of variance stopping condition is met.

Similarly, *backward selection* starts with all possible explanatory variables in $\mathbf{X}$, $\mathbf{V}$, and $\mathbf{T}$. This initial model is also called the *saturated model*. 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.

Analysis of variance (ANOVA) performs a hypothesis test comparing two models $\mathcal{M}_1$ and $\mathcal{M}_2$. The null hypothesis is that the less-complex model $\mathcal{M}_1$, with fewer explanatory variables, is sufficient to describe the response. The alternate hypothesis that a more complex model $\mathcal{M}_2$ is required. ANOVA tests whether the variance explained by an added explanatory variable or interaction are significantly different than the original model [31]. The formulation of ANOVA is beyond the scope of our study. We direct the reader to [15] for full details of the procedure.

2.1.4 Heteroscedastic Residuals

Fitted models are said to be *heteroscedastic* if the variance of residuals increases or decreases systematically with the value of the estimated response vector $\hat{\mathbf{Y}}$, with the value of an explanatory variable, or with previous residuals. In our results, we will visualize heteroscedasticity by plotting $\mathbf{e}$ as a function of $\hat{\mathbf{Y}}$ and as a function of several explanatory variables.

When modelling time series data, it is common that the residual at a given index $\mathbf{e}_i$ will be correlated with previous residuals, a form of heteroscedasticity known as *serially correlated residuals*. We will report the Durbin-Watson test for serial correlation [5], shown in (2.10).

$$DW = \frac{\sum_{i=2}^{N}(\mathbf{e}_i - \mathbf{e}_{i-1})^2}{\sum_{i=2}^{N} \mathbf{e}_i^2} \tag{2.10}$$

The value of the Durbin-Watson statistic DW will always be in the range $(0, 4)$. $DW = 2$ indicates that there is no serial correlation present within the residuals. $DW < 1$ indicates positive correlation, meaning sequential residuals are often similar. $DW > 3$ indicates negative correlation, meaning sequential residuals are often far apart.

In the presence of heteroscedastic residuals, normal standard error estimates for each explanatory variable are too small, implying that we have too much confidence in the coefficient estimates. This has implications when using standard error estimates in confidence intervals, ANOVA, and other hypothesis tests. Because normal standard error estimates are erroneously small in the presence of heteroscedasticity, hypothesis tests would report explanatory variables as significant when they are not. To correct for this problem, standard errors should be estimated using heteroscedasticity and autocorrelation consistent (HAC) standard errors. The HAC procedure increases the standard error of coefficient estimates, indicating that we have less confidence in our model. These wider HAC standard errors are appropriate for use in hypothesis tests. The formulation of HAC standard errors is beyond the scope of our study. We direct the reader to [63] for full details of the procedure.

2.1.5 Out-of-Sample Validation

Though our focus is on the descriptive power of our model, we also report models' out-of-sample predictive power using mean absolute error (MAE) and mean absolute percentage error (MAPE) shown in (2.11) and (2.12) respectively [28].

$$MAE = \frac{\sum_{i=1}^{N} \mid \mathbf{e}_i \mid}{N} \tag{2.11}$$

$$MAPE = 100 \cdot \frac{\sum_{i=1}^{N} \mid \mathbf{e}_i / \mathbf{y}_i \mid}{N} \tag{2.12}$$

To measure out-of-sample prediction accuracy, we use a form of time series cross-validation described in [29] by training our model on the first 12,312 hours of data and test using the next 168 hours (i.e. one week). We then slide the windows of training and test data forward 168 hours yielding a new set of training and testing data. This process is repeated 12 times, yielding 12 MAE values and 12 MAPE values. We report the mean MAE and mean MAPE averaged over the 12 out-of-sample validation folds.

2.2 Prior Work

Both aggregate residential electricity demand and the effects of time-based electricity pricing are well-studied problems. However, there are several recent changes enabled by smart meters that make our study relevant. Household, hourly smart meter readings are still a fairly recent source of data. They may be used individually to identify customer groupings and common behaviours. They may also be used in aggregate to provide short-term and mid-term insights about electricity demand. Our study focuses on the latter, providing insight about aggregate electricity demand relevant at the mid-term planning horizon of several months or several years. Table 2.1 provides an overview of temperature observation transformations used in prior work. These techniques use temperature observations as an input to a number of basis functions, transforming it into into a basis matrix $\mathbf{T}$. The prior work described throughout the remainder of this section each describe different transformations to create $\mathbf{T}$.

Table 2.1: Overview of three temperature transformation categories which we consider.

Coincident Weather Transformations	
Temperature Humidity Index	[17, 48]
Humidex	[18]
Other Humidity Transform	[42]
Wind Speed	[20, 42]
Solar Irradiance and Cloud Cover	[20, 8]
Temporal Transformations	
Lagged Observations	[24]
Heating Degree-Days / Cooling Degree-Days	[57, 9]
Heating Degree-Hours / Cooling Degree-Hours	[47]
Moving Average	[42]
Weighted Moving Average	[20, 8]
Non-Linear Transformations	
Switching Regression	[41, 18, 47, 48, 35, 62]
Threshold Regression	[41, 6, 4]
Threshold Regression with Saturation at Extremes	[9]
Regression Splines	[12, 24]

2.2.1 Coincident Weather Variables and Transformations

We define *coincident weather* to be other 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* [17, 48], the Canadian *Humidex* [18], or by incorporating humidity into some other transformation of temperature [42]. Humidity may have a direct effect on load via dehumidification equipment, or it may have an indirect effect on load via human perception and comfort levels. Wind speed has also been incorporated into temperature transformations by [20, 42]. Wind may reduce electricity demand if customers choose to cool their home by leaving windows open during transition seasons. It may also affect human perception of cold outdoor temperatures via wind chill, inclining them to stay indoors.

2.2.2 Temperature's Effects Over Time

Effort has also been taken to account for the delay between when an outdoor temperature occurs to when its effects are felt within a customer's home. Depending on the quality of housing insulation prevalent in the area, this heat transfer can take a number of hours. *Heating degree-days* and *cooling degree-days* are common derived values used to measure the prolonged heating and cooling requirements of a home over time. As hourly electricity demand readings have become commonplace, these metrics have been extended to *heating degree-hours* and *cooling degree-hours*. The traditional degree-day or degree-hour procedure is based on the idea that for residential buildings, electricity demand will be proportional to the difference between the mean daily temperature and a temperature break point [56]. It is found by summing the number of recent observations that have been below or above a given break point. For analysis of long-term and mid-term horizons, heating and cooling degree-days have been used and are sufficient [57, 9]. Cooling and heating degree-hours are better suited to analysis of short-term and mid-term horizons [47]. This transformation is described with greater detail in section 3.7.4.

In [24], the authors considered lagged hours of temperature in early models of their study, though ultimately they did not use lagged temperature in the final model. The authors of [42] used a four-hour *moving average* of recent temperatures as a component of the space heating index used in their model. The most sophisticated method of accounting for thermal transfer inertia is found in [20], which is a refinement of [8]. In these papers, the authors define an exponentially weighted moving average filter to be the smoothed temperature. This smoothed temperature represents the heat transfer inertia of buildings. It is combined with an instantaneous outdoor temperature offset by cloud cover and solar irradiance. Together the smoothed temperature and instantaneous temperature are combined into a *sensible temperature* which is ultimately used as a component in their multiple regression model.

2.2.3 Non-linear Effects of Temperature

Recall (2.1), in which the underlying functional form of explanatory variables is assumed to be linear and additive. Temperature has a non-linear relationship with electricity demand, discussed in greater detail throughout chapter 3. To model temperature's non-linear relationship with electricity demand, the temperature observations must be transformed. There are a number of common techniques used to model explanatory variables which have a non-linear association with the response variable.

The authors of [41] describe two types of piecewise linear models and create a terminology that adds meaning for electricity demand analyses. They describe a single temperature break point as a *switching regression*. The line fit to temperatures below the break point represents household heating effects. The line fitted to temperatures above the break point are cooling effects. When two break points are selected, the authors describe it to be a *threshold regression*. Temperatures below the cooler temperature break point are heating effects. Temperatures above the warmer temperature break point are cooling effects. The range of temperatures between these two points, the threshold, are assumed to be temperatures where no heating or air conditioning is used within the household. The threshold is typically a region of moderate temperatures experienced during spring and autumn seasons. Using the terminology established in [41], a switching regression is used in [18, 47, 48]. The authors of [35, 62] also use switching regression, but the lower region has a slope of zero because residential users in Arizona and California do not have heating requirements. A threshold regression is used in [6] and [4] to approximate customer-specific thermal responses.

The authors of [9] note that extreme low temperatures and extreme high temperatures exhibit saturation of heating and cooling effects. At these extreme temperatures, all household thermal controls such as space heaters, electric baseboard heating, fans, or air conditioning available to residential are working constantly. Electricity demand plateaus at the maximum amount associated with heating or cooling.

Extending the multiple regression model beyond piecewise linear regions, temperature observations may be transformed into derived values that better support non-linearity. Building on piecewise linear transformations, smoothed transitions between threshold regions and heating/cooling effects are described in [8, 20, 41].

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 [12, 24]. 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, known as *knots*. This transformation is described with greater detail in section 3.7.3.

Chapter 3

Comparison of Temperature Transformations

Statistical models may fall into one of three categories: explanatory models, descriptive models, or predictive models [61]. Explanatory models must demonstrate causality, formally testing a causal hypothesis. They are often parametric models chosen for their modularity and the interpretability of their coefficient estimates. Models with high explanatory power may not necessarily have high out-of-sample prediction accuracy. Conversely, models that have high predictive power may not have much explanatory value. The primary focus of predictive modelling is the ability to forecast future observations accurately.

Our study falls under the third category: descriptive modelling. We define a parametric multiple regression model with explanatory variables relating to weather, time, and residential electricity price. Because our data sample is not from a controlled experiment and has no classic control group, we cannot test causal relationships between our explanatory variables and electricity demand. However, we can describe changes in aggregate electricity demand associated with each explanatory variable. Our primary objective is inference about the effects of weather and TOU pricing on aggregate electricity demand. We also use time series cross-validation to evaluate our model's out-of-sample predictive accuracy and to maintain model parsimony. We focus on a mid-term planning horizon, evaluating how the local distribution company's service region responds to time-of-day, changes in weather, and Ontario's TOU pricing policy.

3.1 Data Description

3.1.1 Smart Meter Data

The smart meter data used in our study was provided by a local distribution company in south west Ontario. It contains hourly smart meter readings from 28,890 customers across a four-city service region. The time range of observations is from October 29, 2010 through October 17, 2012. Attributes available for each customer include a unique meter number, the longitude/latitude of the meter, kWh of electricity consumed per hour, connection/disconnection dates, and the distribution transformer that the meter is connected to.

Prior to November 1, 2011, residential customers were billed according to a seasonal, flat rate. The long-standing flat rate pricing structure was comprised of a summer price of 6.8¢/kWh over all hours of the day from May through October. It then changed to a winter price of 7.1¢/kWh over all hours of the day from November through April.

On November 1, 2011 the electricity pricing structure changed to TOU rates 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. Winter off-peak hours are 7:00pm through 6:59am (overnight) at 6.2¢/kWh. Mid-peak hours are 11:00am through 4:59pm at 9.2¢/kWh. On-peak hours are 7:00am through 10:59am and 5:00pm through 6:59pm at 10.8¢/kWh. In both summer and winters, all hours of weekends and holidays are off-peak rates. The weekday TOU pricing seasons are illustrated in Figure 3.1. The distribution company's concise transition from flat rates to TOU pricing provides a clear opportunity to measure the effectiveness of TOU pricing in accomplishing two of its stated objectives: overall conservation and demand-shift between price periods.

3.1.2 Weather Data

The weather data set used in this study is hourly, historical observations from two nearby Environment Canada monitoring stations [13, 38]. Weather observations are paired with each meter by selecting those from the nearest monitoring station, all within 5-25 kilometres. Section 3.2 describes the process of aggregating the smart meter readings to create a single electricity demand time series. Similarly, a weighted average of weather observations is created based on the number of meters reporting each hour. Additional weather

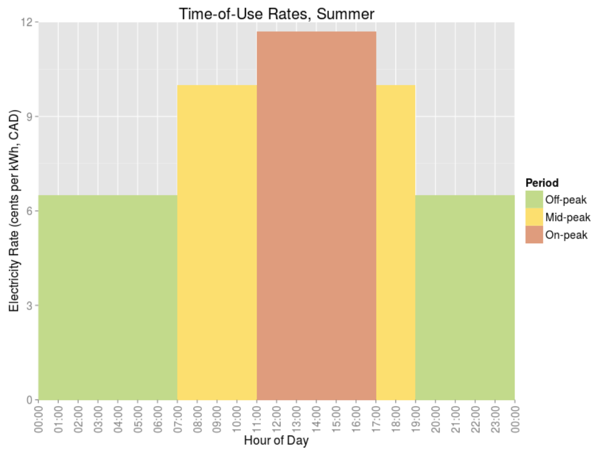

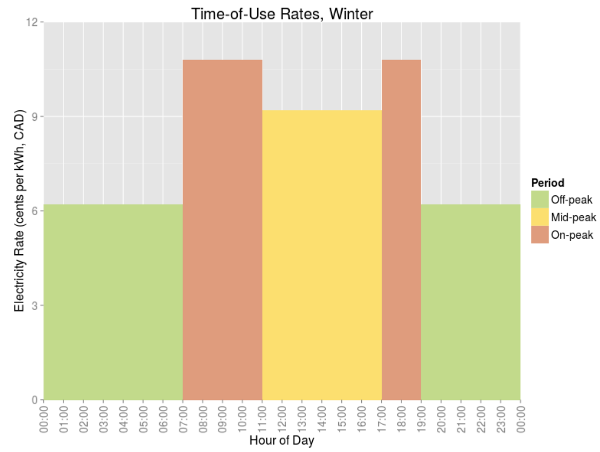

Figure 3.1: Plot of the TOU rate structure for weekdays. All hours of weekends and holidays are billed at the off-peak rate. *Left*: Winter weekday rates, November through April. *Right*: Summer weekday rates, May through October.

attributes recorded each hour are dry-bulb temperature, relative humidity, dew point, wind direction, wind speed, visibility, atmospheric pressure, Humidex, wind chill, and a weather condition description. The term *dry-bulb* refers to air temperature measured by thermometer, shielded from moisture and radiation.

Let τ be an $N \times 1$ vector of hourly dry-bulb outdoor temperature observations. This is not to be confused with Kendall's τ rank correlation coefficient, common in statistics literature; we do not discuss Kendall's τ in our study. The vector of outdoor dry-bulb temperature observations τ will be used as input to temperature transformation functions with 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, $\mathbf{T} = \tau$ and $\hat{\theta}$ is a scalar.

3.2 Data Cleaning Process

Our first data cleaning step ensures that only residential meters are sampled. The local distribution company's data is a mixture of unlabelled commercial and residential meters. Using the distribution transformer attribute for each meter, we are able to infer labels for meters. First, only single-phase transformers and their connected meters are selected. Second, because the longitude and latitude are known for each transformer and meter, it is possible to plot them in geographic information system software. The locations of meters and transformers are exported to a Keyhole Markup Language (KML) file and viewed in Google Earth [23]. In the KML file, `folder` elements are created for each distribution transformer and child `placemark` elements are created for each meter connected to that transformer, shown in Figure 3.2. In this way, the visibility of all meters connected to a transformer can be toggled on and off easily within Google Earth's interface. Using this method, transformers are manually labelled as either residential, commercial, or mixed-use based on visual information. If all meters connected to the transformer are clearly residential or commercial, the transformer is labelled accordingly. Transformers that are connected to both residential meters and commercial meters are labelled as mixed-use to avoid type I errors (i.e. false-positives). Using this process, it is possible to infer labels for 28,890 meters by manually labelling 2,657 distribution transformers. Selecting only meters connected to residential labelled transformers results in 23,670 residential meters to be considered in subsequent data cleaning steps.

Second, the connection date and disconnection date of meters are known. For example, meters in newly constructed or demolished buildings would have to be connected or disconnected from the distribution grid. We remove meters that were connected or disconnected during the sample period.

Third, customer changes associated with each meter are known. For example, if the tenant changes within a rental unit, then the customer identifier connected to the meter changes. We remove meters from our sample that had a customer change during the sample period.

Finally, high level diagnostics are performed on this interim data sample to uncover data quality issues. The only concerning issue identified are that some residential meters have extremely high readings. The worst of which is a meter reading of 214,902.7 kWh for a given hour. A 215 MWh reading is clearly an equipment failure or data transmission error. Rather than introducing bias into selecting which meters have reported erroneous values, a reasonable filter criteria is used to throw out meters reporting erroneous values. IEEE C57.91-2011 section 8.2.2 describes the maximum short-term overloading of a distribution transformer to be 300% of its nameplate rating [30]. Using this guideline, if a smart meter has an hourly reading that single-handedly violates the maximum short-term overloading capacity of the transformer it is connected to, then that meter is removed from the sample.

The remaining residential meter sample after all data cleaning steps contains 20,556 meters. A possible bias our data cleaning technique may have introduced is that meters from mixed-zone housing in city-centres or rentals with changing tenants may be under-represented due to removal of mixed-use transformers and removal of meters with a

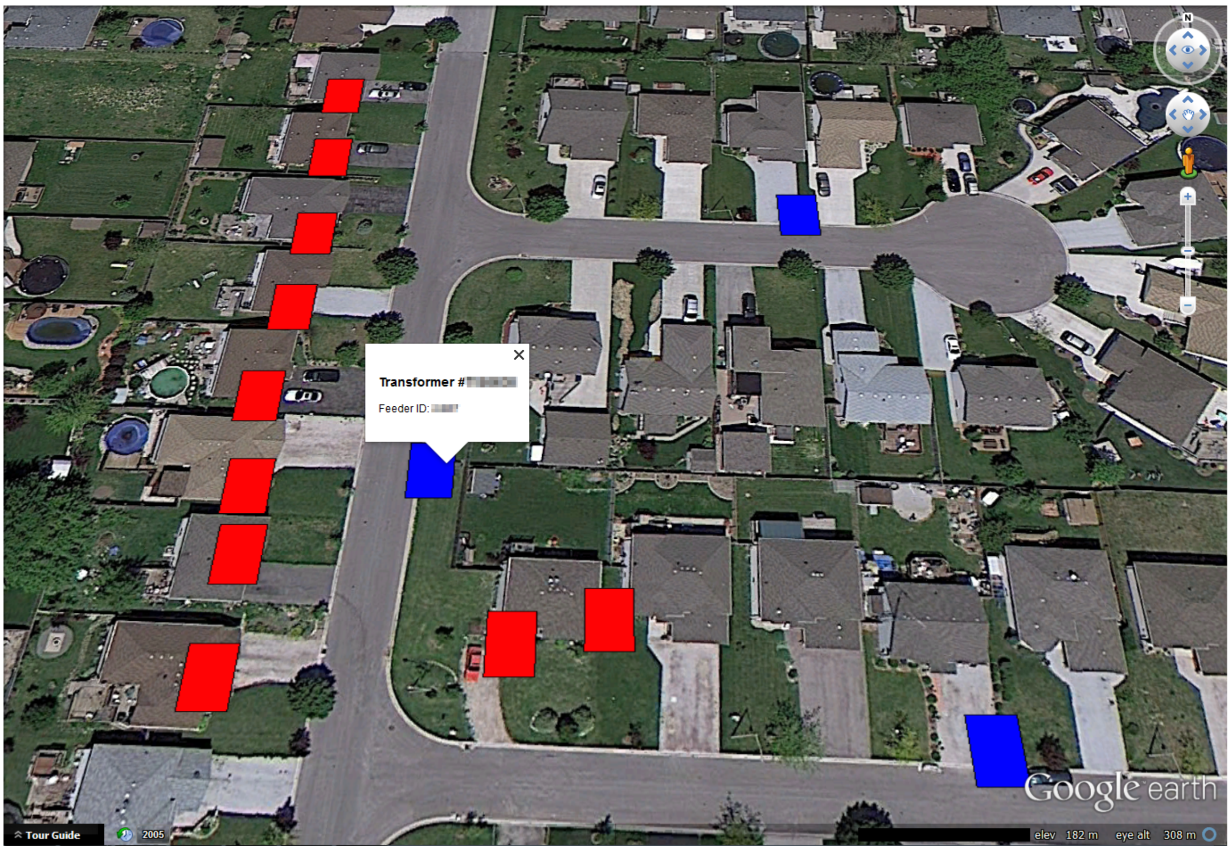

Figure 3.2: This screenshot from within Google Earth illustrates the meter labelling process. Meters are shown as red squares. Distribution transformers are shown as blue squares. In this screenshot, the meters for one selected transformer have been made visible. The meters for the two unselected transformers have been hidden.

customer change. A cursory look at the meter sample viewed in Google Earth shows that large apartment complexes are still adequately represented in the meter sample.

A second data quality issue is the presence of missing observations. From October 29, 2010 through February 28, 2011 only seven meters reported electricity demand. Those months have been removed from the sample so that the fitted model is not biased toward those seven customers' behaviour. In the remaining March 1, 2011 through October 17, 2012 sample, most meters have at least a few missing observations over the course of the sample period. Often, a meter's missing values occur as irregularly positioned gaps lasting multiple hours, such that data interpolation is not suitable. We choose to study the data in aggregate by deriving the average household demand in each hour from all households that reported observations during that hour. The normalized aggregate residential electricity demand time series is shown in (3.1). Going forward, we will refer to this as *aggregate electricity demand*.

$$\mathbf{y}_i = \frac{\sum_{j=1}^{J_{meters}} \mathbf{\Upsilon}_{i,j}}{\sum_{j=1}^{J_{meters}} I(\mathbf{\Upsilon}_{i,j} > 0)}, \quad i = 1, ..., N \tag{3.1}$$

Let $\mathbf{\Upsilon}$ be the $N \times J_{meters}$ dense matrix of smart meter readings from March 1, 2011 through October 17, 2012. $N = 14,328$, the number of hours in our sample time series after data cleaning. $J_{meters} = 20,556$, the number of meters remaining after data cleaning. Let the indicator function $I()$ return 1 if there exists a reading for meter j at hour i. As the function 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. We use the vector $\mathbf{Y}$ as the response variable for the remainder of this study. Figure 3.3 illustrates the number of meters reporting each hour.

3.2.1 Overview of the Aggregate Electricity Demand Sample

Figure 3.4 shows the sample of aggregate electricity demand plotted over time. The summer air conditioning requirements can clearly be seen during summer months. There is a less noticeable heating effect during winter. Figure 3.5 shows a density plot of $\mathbf{Y}$. 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.

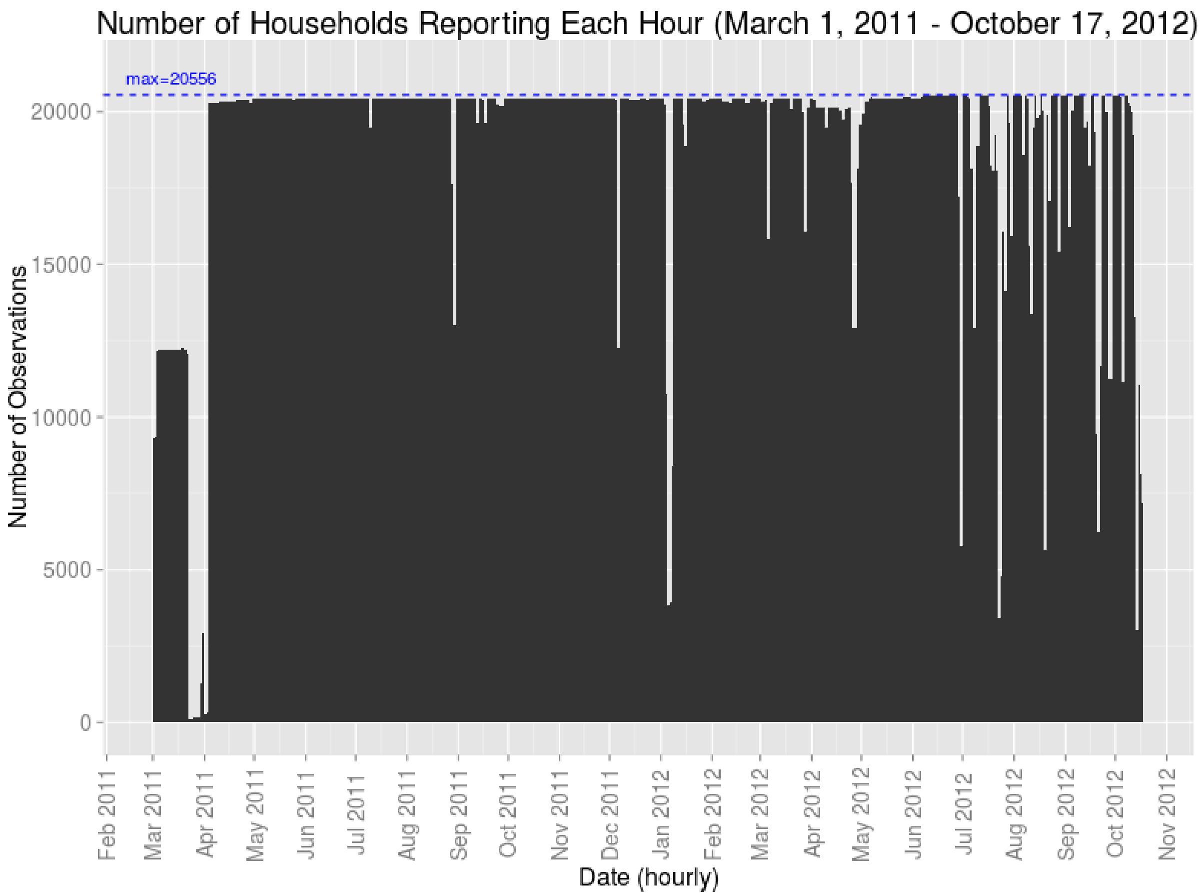

Figure 3.3: The number of meters reporting demand for each hour of the sample period. The maximum number of meters reporting each hour is 20,556. The minimum number reporting hourly during the sample period is 77 meters for a period in late March 2011.

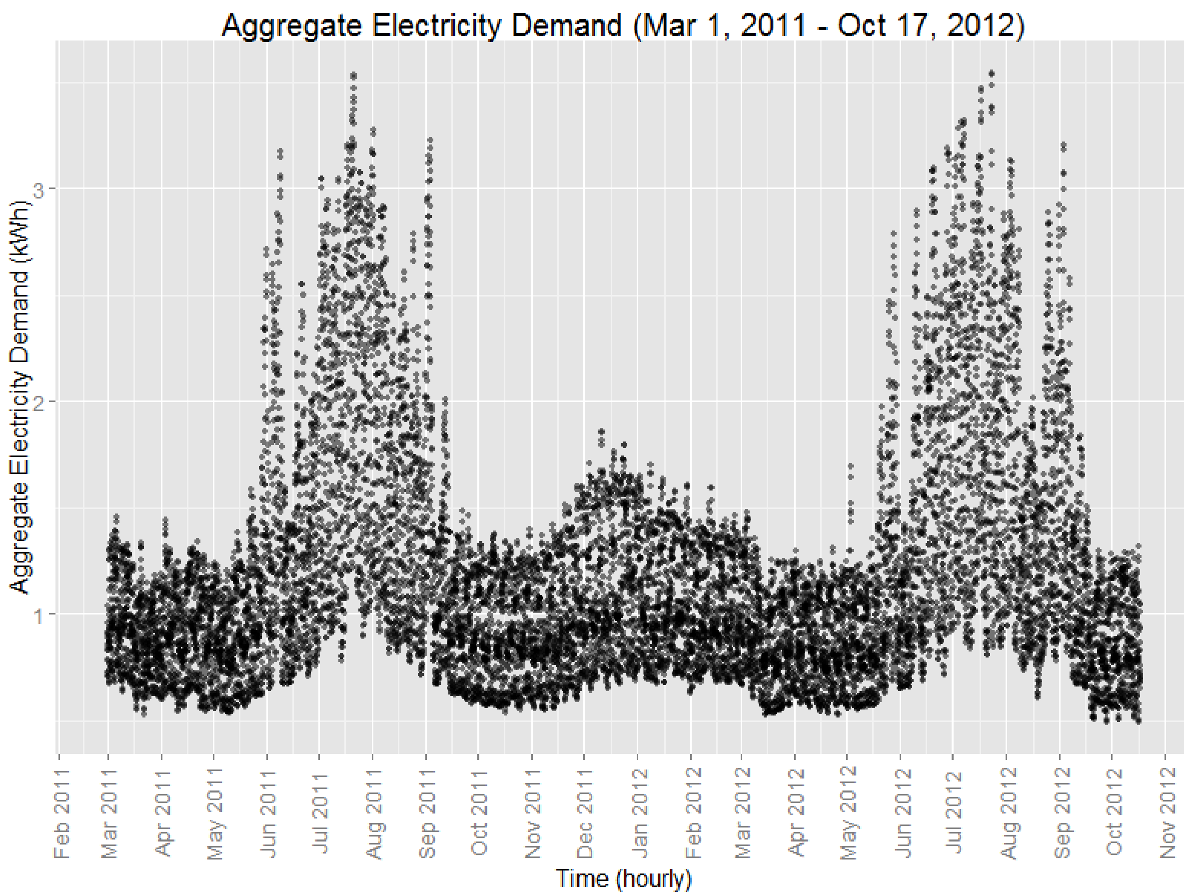

Figure 3.4: The cleaned sample of aggregate electricity demand plotted as a function of time. Opacity has been used to give a sense of observation density.

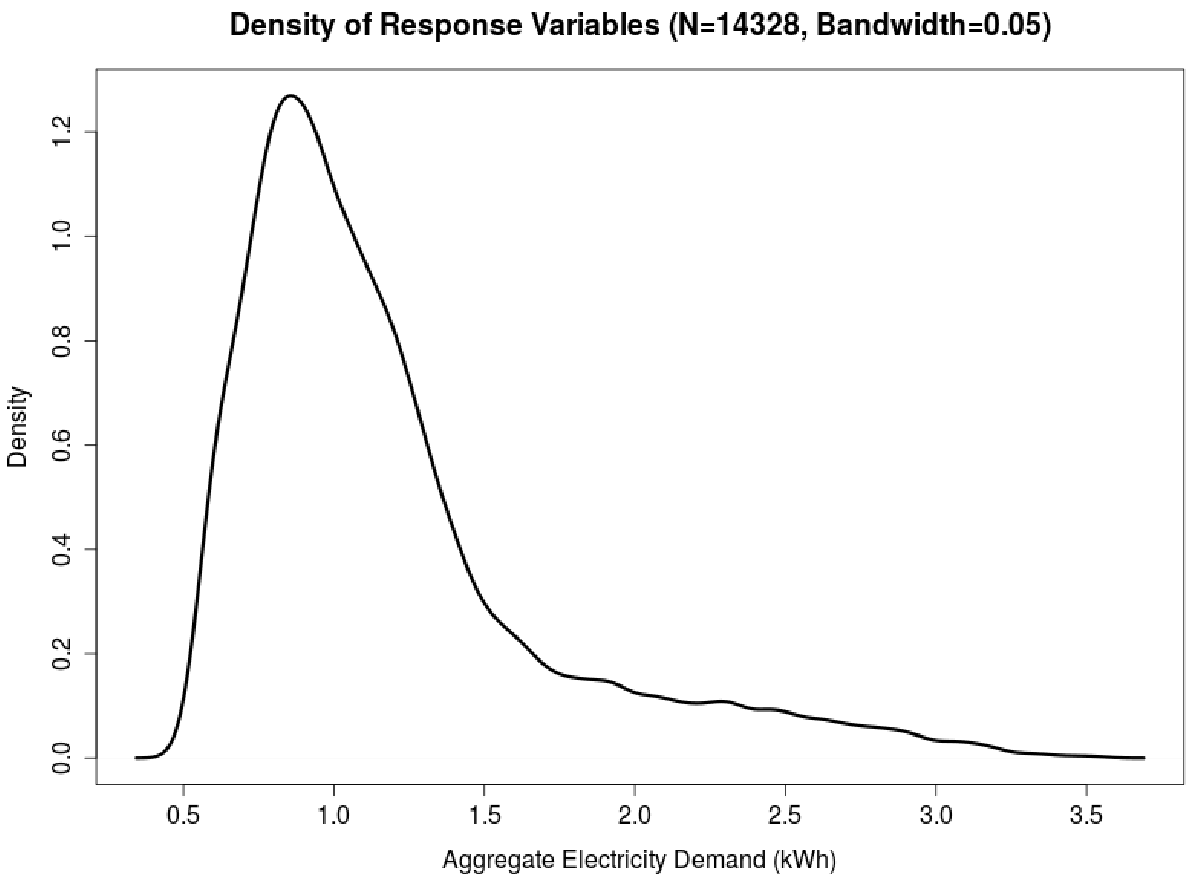

Figure 3.5: Density of the response vector $\mathbf{Y}$, a sample of aggregate electricity demand.

3.3 Temperature Transformation Comparison

This section of our analysis is a detailed comparison of existing temperature variable transformations that incorporate coincident weather observations, past temperature observations, and varying degrees of non-linearity. We will illustrate in section 3.9 that aggregate electricity demand is predominantly associated with changes in weather. By comparing all combinations of temperature variable transformations and selecting a well-performing model, a substantial amount of variance can be explained by weather. Based on our results, we will carry the well-performing model forward into chapter 4. After explaining as much variance associated with temperature as possible, the moderate effects of TOU pricing on aggregate electricity demand can be isolated and more accurately quantified.

Recall the multiple regression model used to estimate aggregate electricity demand (2.1). We restate the equation below for clarity:

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

Let $\hat{\mathbf{Y}}$ be an $N \times 1$ vector representing the estimate of $\mathbf{Y}$. N is the number of observations in our aggregate electricity demand sample. Let $\hat{\beta}_0$ be the estimated intercept term. We formulate 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 $\hat{\beta}$, $\hat{\omega}$, and $\hat{\theta}$ fit using ordinary least squares.

Our treatment of time and price explanatory variables is consistent with prior work. Accordingly, the explanatory variables in $\mathbf{X}$ and $\mathbf{V}$ remain fixed for each iteration of our comparison. 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. Note that categorical variables are modelled using contrasts and require one fewer degree of freedom than there are levels in the variable. Contrasts make use of the *baseline* value for each categorical variable that needs no indicator. For example, the baseline value for hour-of-day is `00:00`. All other hour-of-day coefficients are offset from that baseline [31]:

- $\mathbf{x}_{\bullet,p=1}$ 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.
- $\mathbf{v}_{\bullet,p=1}$ 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.

Explanatory variables in $\mathbf{X}$ and $\mathbf{V}$ are selected using forward selection and ANOVA as described in section 2.1.3. Sections 3.4 and 3.5 will describe each of the explanatory variables in $\mathbf{X}$ and $\mathbf{V}$.

We define three steps of temperature transformations which are used in conjunction with one another to generate $\mathbf{T}$:

1. **Coincident Weather Transformations**: dry-bulb temperature, *feels like* temperature
2. **Temporal Transformations**: current observation, lagged observations, CDH/HDH, moving average, and exposure-lag-response association
3. **Non-Linear Transformations**: switching regression, natural cubic splines

Our comparison iterates over all combinations of temperature transformations. Each iteration uses a different combination of transformation functions to generate the temperature transform basis matrix $\mathbf{T}$ while holding the matrices $\mathbf{X}$ and $\mathbf{V}$ fixed.

We also include several base 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 $\mathbf{T} = \tau$.

3.4 Temporal Explanatory Variables

Our treatment of temporal variables such as hour-of-day and working day are consistent with prior work. We will briefly describe the relationship of each temporal explanatory variable with electricity demand.

3.4.1 Hour of Day

First, we will define hour-of-day. The province of Ontario observes daylight savings time. In our sample, users shifted their clocks forward one hour on March 13, 2011 and March 11, 2012 as the province entered daylight savings. Clocks were shifted backward one hour on November 6, 2011 as the daylight savings ended. The hours used as explanatory variables are the hours customers base their routines around and that the province bases TOU price periods around. This seemingly obvious assumption has three implications for the structure of our model.

First, the periodicity of time series electricity demand is 24 hours. We considered modelling periodicity as a cyclic function using index i as an input. Using regression splines or a Fourier transform could potentially save many degrees of freedom in the model. However, daylight savings would essentially be a phase shift which requires special treatment.

Second, we wish to construct a descriptive, interpretable model. Particularly in chapter 4, when hour-of-day is interacted with utility rate seasons and TOU pricing, we wish to interpret effects associated with aggregate electricity demand for each hour. A basis matrix used to model a periodic function of i would be difficult and non-intuitive to interpret. As shown in [1], the alignment of Ontario's three TOU periods to demand patterns is a topic of interest and may not be optimally defined.

Third, when interacted with a working day indicator, the hour-of-day factor has clear and interpretable meaning.

For these reasons, hour-of-day is modelled as a categorical factor with 24 terms. In $\mathbf{X}$, hour-of-day is represented by 23 sparse columns with indicators for each hour. Following the convention of model contrasts, the hour 00:00 is considered the baseline and has no indicator column in $\mathbf{X}$. All other hours of the day are represented by indicator variables and deviate from this baseline case [31]. Figure 3.6 shows a box plot of electricity demand grouped by hour-of-day. See Appendix B for details regarding box and whisker plot interpretation. Examining the distribution for each hour shows the expected patterns of user activity within the home.

3.4.2 Working Days

It is intuitive to assume that residential electricity demand will differ by day-of-week. Demand within the home will typically be different on a Monday than a Saturday. It will also differ on holidays, when residential customers will not follow their typical weekday demand patterns. Following this intuition, we initially added day-of-week as a seven-category factor and an indicator variable flagging statutory holidays in the province. However, we are able to achieve the same level of explanatory power using only one degree of freedom by defining a *working day* indicator, rather than using seven degrees of freedom to model day-of-week and holiday. The working day indicator is defined similar to [40, 41] such that

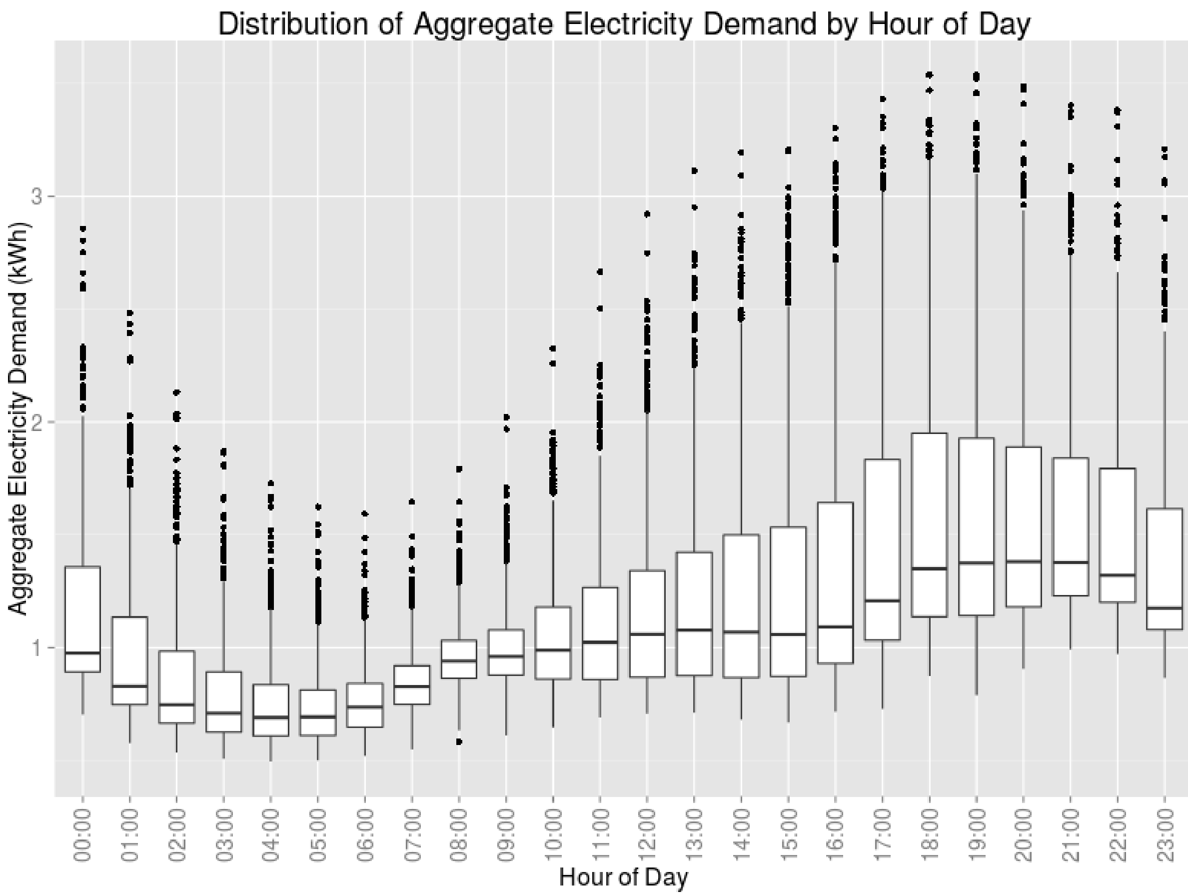

Figure 3.6: Plot of aggregate electricity demand grouped by hour. Note that this plot contains data from both working and non-working days.

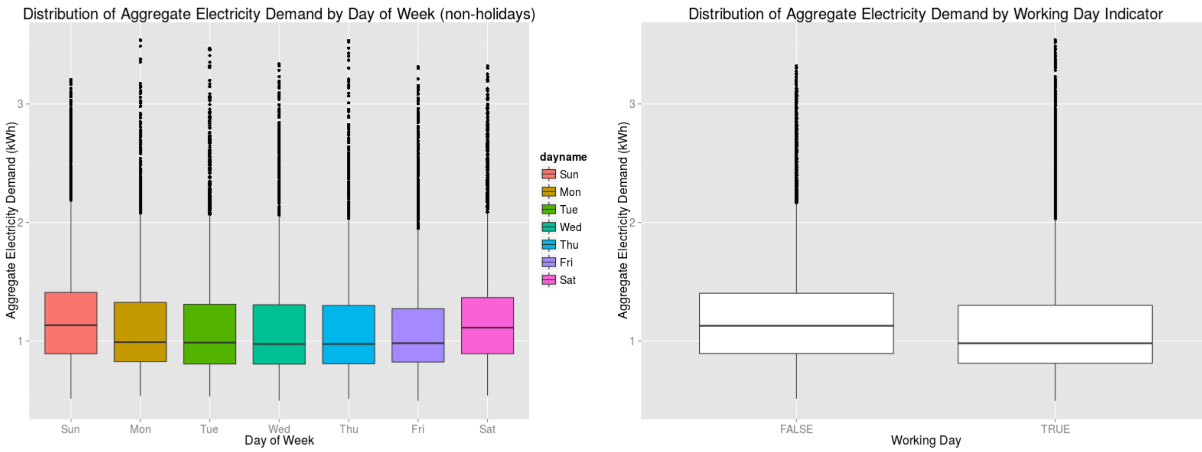

Figure 3.7: *Left*: Distribution of electricity demand grouped by day-of-week, for non-holidays only. *Right*: Distribution of electricity demand grouped by the working day indicator.

weekends and holidays are non-working days. All other weekdays are working days, shown in Figure 3.7. This in line with the local distribution company's definition of a working day and associated off-peak TOU prices. It allows for meaningful variable interactions to be defined in chapter 4 and reduces the degrees of freedom in our multiple regression model.

A statistically significant and intuitive *interaction* exists for working day and hour-of-day. An interaction between two categorical factors like 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 or baseline values. The main effects of each explanatory variable represent deviation from the sample mean and the two-way interactions represent deviation from their main effects.

For example, the baseline for hour-of-day is 00:00 and the baseline for working day is `working_day=FALSE` (i.e. non-working days). If observation i occurs at 00:00 on a non-working day, neither variable's main effects will be added to β_0. If observation i is 07:00 on a non-working day, only the coefficient estimate for 07:00 will be added to β_0 (i.e. main effects). Neither a main effect for working day nor an interaction for the two is added because the working day explanatory variable is at its baseline. If observation i is 07:00 on a working day. Then both the coefficient estimate for 07:00 and the coefficient estimate for `working_day=TRUE` will be added to β_0 (i.e. main effects). Additionally, an interaction indicator for the combination of 07:00 $\times$ `working_day=TRUE` exists and adds an interaction effect to β_0. An example of working day $\times$ hour-of-day interaction coefficient estimates is

Table 3.1: Coefficient estimates illustrating the intuition behind the hour-of-day × working day interaction.

Interaction Term	Coefficient Est.	p-val	Significance
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	

given in Table 3.1. The coefficient estimates will change slightly with each temperature transformation compared, but the sign, intuition, and statistical significance remain applicable. To interpret coefficient contrasts for an interaction term, one must keep in mind the baselines for each explanatory variable: 00:00 and `working_day=FALSE`. Aggregate electricity demand begins earlier on working days than the `working_day=FALSE` baseline. This activity is likely caused by residential customers preparing for work around 07:00 or 08:00 on working days, reflected by a positive coefficient estimate that is of noticeable effect size and has a statistically significant p-value. 10:00 through 17:00 on working days are associated with less aggregate electricity demand, likely because the majority of residential customers are at work. All coefficient estimates are negative, have meaningful effect size, and have statistically significant p-values.

3.5 Seasonality and Price

As suggested by Figure 3.4, 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, for temperature transformation comparison we are primarily interested in the explanatory power of each temperature transformation combination used to generate $\mathbf{T}$. Any explanatory variable that is collinear with the temperature transformation basis matrix masks its effects. Collinearity indicates the values of two explanatory variables increase and decrease together, fitting the effects of the underlying temperature phenomenon across two variables' coefficient estimates.

We check for *collinearity* using *variance inflation factor* (VIF) shown in (3.2) and the `car` package in R [19].

$$VIF(\hat{\beta}_j) = \frac{1}{1 - R^2_{\mathbf{X}_j|\mathbf{X}_{\bullet,j}}} \tag{3.2}$$

The term $R^2_{\mathbf{X}_j|\mathbf{X}_{\bullet,j}}$ represents the R^2 value of regressing $\mathbf{X}_j$ onto all other explanatory variables. As R^2 of one variable regressed onto others approaches 1, the VIF will become large, indicating collinearity. Table 3.2 shows VIF values for the main effects of a model fit with current, dry-bulb temperature, transformed using natural cubic splines (section 3.7.3). Categorical variables from $\mathbf{X}$ are also used. Generally, a $VIF > 5$ indicates collinearity, meaning that the term may mask the significance of other collinear terms [31]. In Table 3.2, none of the main effects considered are collinear.

Table 3.3 shows VIF values when a categorical variable for month is considered as well. In the second scenario, month and temperature are indicated to be collinear due

Table 3.2: VIF of explanatory variable main effects with no seasonality.

Explanatory Variable(s)	VIF	Degrees of Freedom
Natural Cubic Splines **T**	1.12	4
Hour-of-Day	1.12	23
Working Day	1.00	1

Table 3.3: VIF of explanatory variable main effects with a categorical variable for month.

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

to $VIF > 5$. The addition of a month explanatory variable would mask the effects of temperature.

Instead, we have considered two slightly similar but distinct categorical variables for utility rate season and a TOU billing indicator to be included in the explanatory variable matrix **V**. Utility rate season takes values "summer" or "winter" according to the local distribution company's billing season. It indicates a different, economic seasonality that the customer may react to. The TOU billing indicator, referred to as *TOU Active* indicates whether TOU billing has gone into effect. Together, these two economic indicators add model pricing effects that may influence residential customers. Table 3.4 of VIF statistics for this model. The addition of pricing seasonality is not collinear with **T**. The terms

Table 3.4: VIF of explanatory variable main effects with addition of utility rate season and a TOU billing indicator.

Explanatory Variable(s)	VIF	Degrees of Freedom
Natural Cubic Splines **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

also have statistically-significant p-values, so their inclusion in the temperature effects comparison are justified.

3.6 Conditioned Value Visualizations

All figures in sections 3.4 and 3.5 illustrating the relationships of temporal and price explanatory variables to aggregate electricity demand have been plots of the original, observed explanatory variables. As we visualize each type of temperature transformation through section 3.7, it becomes difficult to visualize the transformations using observed data. Examining Figure 3.8 one could argue that there is slight horizontal stratification among the plotted observations. For the purposes of visualizing the fitted regression line for a temperature transformation, it is desirable to filter out the effects of other explanatory variables. Because the effects from each explanatory variable are assumed to be independent, the response can be *conditioned* such that the effects from certain explanatory variables can be fixed at some hypothetical baseline. For example, if estimate $\hat{\mathbf{y}}_i$ occurred at 19:00, the coefficient estimate for hour 19:00 could be subtracted from $\hat{\mathbf{y}}_i$ so that the observation has been conditioned as if it occurred at hour-of-day's baseline 00:00. This manner of conditioning the response is done for all variables that are not of interest for the purposes of visualizing temperature transformations [7]. The application of conditioned values is used for visualization purposes only. It does not affect the model fitting process. Figure 3.9 shows the same aggregate electricity demand as Figure 3.8, but with hour-of-day conditioned to 00:00, working day conditioned to `FALSE`, TOU active conditioned to `FALSE`, and utility rate season conditioned to `summer`. All plots in sections 3.7.1, 3.7.2, and 3.7.3 will condition explanatory variables from $\mathbf{X}$ and $\mathbf{V}$ in this manner for the sake of visualizing the remaining variance fit by temperature transformations.

3.7 Temperature Variable Transformations

Sections 3.4 and 3.5 describe the temporal and electricity price explanatory variables in $\mathbf{X}$ and $\mathbf{V}$. In the remaining sections of this chapter, those explanatory variables in $\mathbf{X}$ and $\mathbf{V}$ will remain fixed. This section describes three steps of temperature transformations used to iterate all combinations of temperature variable transformations used to generate $\mathbf{T}$ in our comparison study.

Recall section 3.1.2. We begin the temperature transformation steps with an $N \times 1$ vector of outdoor, dry-bulb temperature observations τ. The algorithm used to generate $\mathbf{T}$ is shown in Algorithm 1. Using this algorithm, all combinations of temperature transformations will be compared.

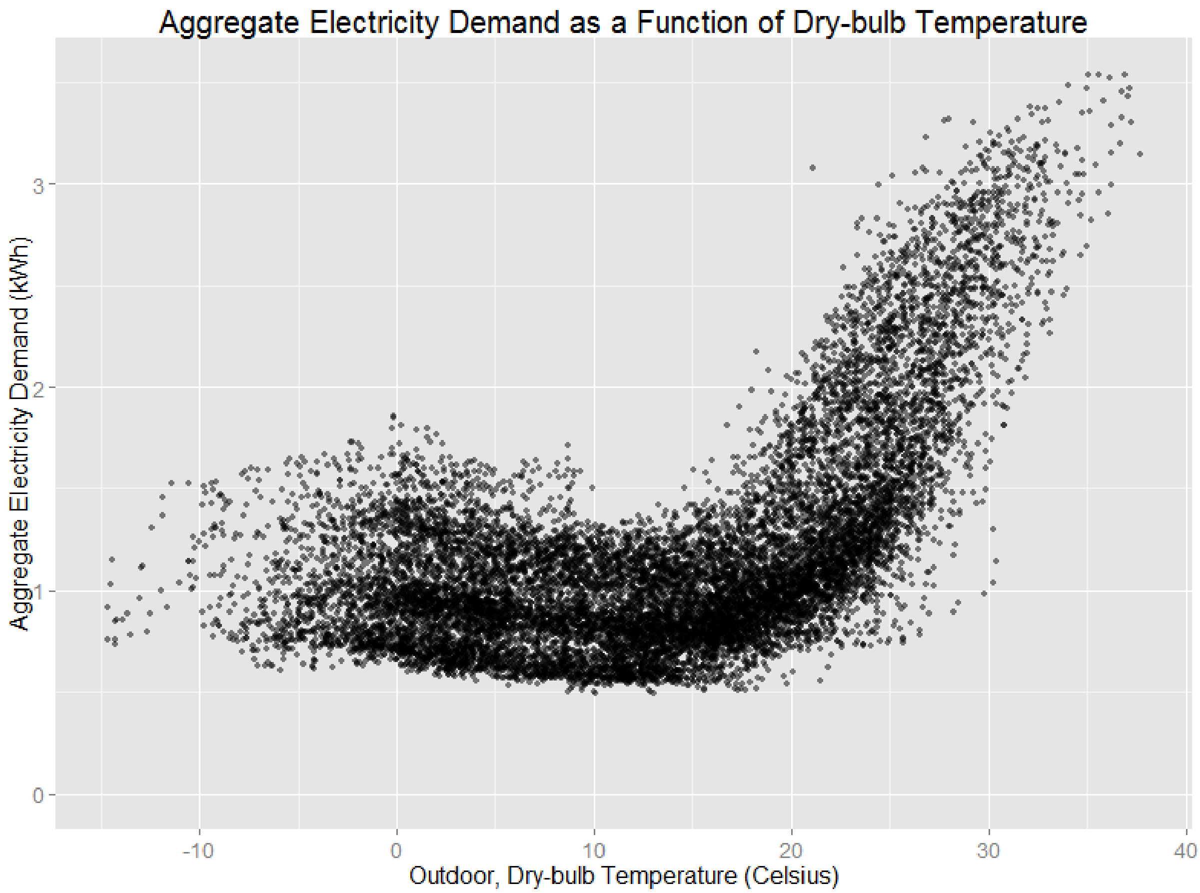

Figure 3.8: Aggregate electricity demand as a function of dry-bulb temperature. Points have been given 50% transparency to give a sense of density.

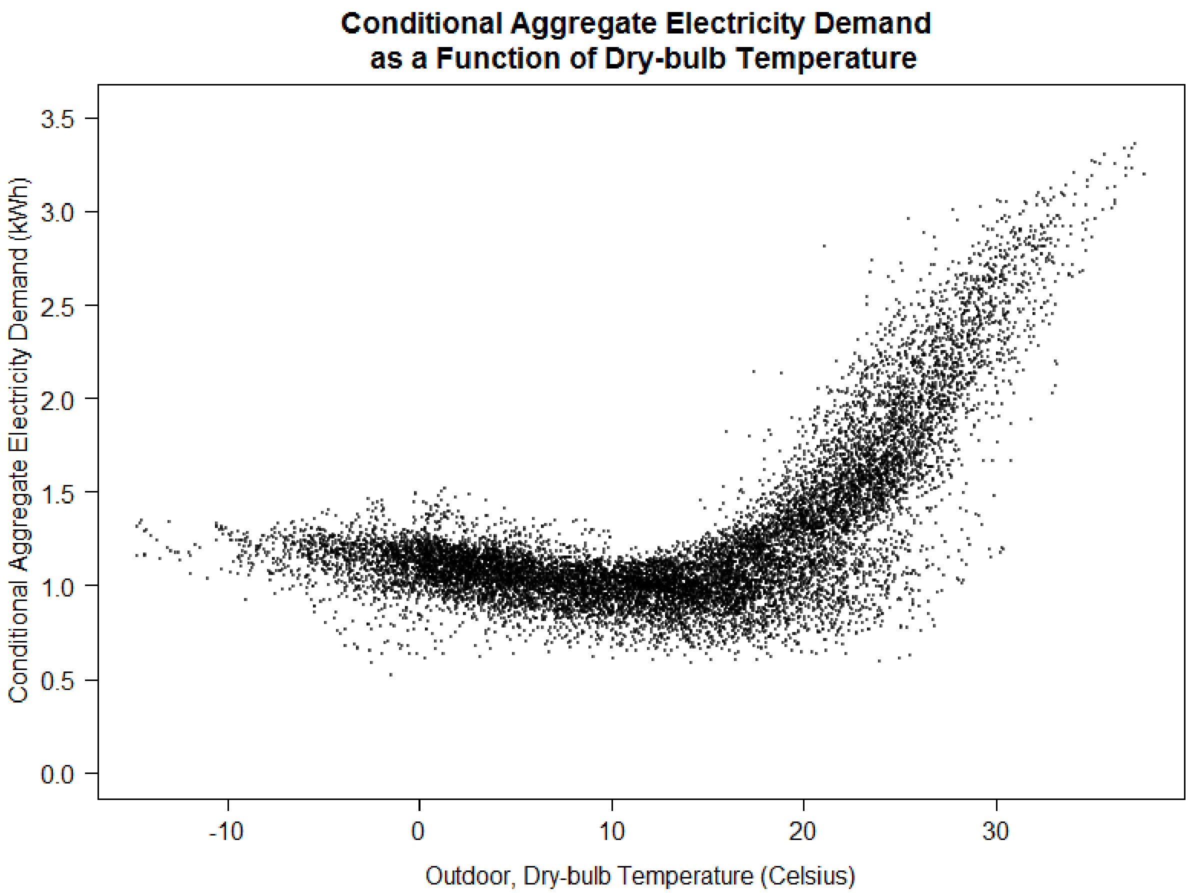

Figure 3.9: After fitting an interim model with explanatory variables in **X** and **V**, we can condition the response on each variables' baseline.

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

1. Transform dry-bulb temperature observations τ into the vector τ' using coincident weather observations.

2. Transform the vector τ' into the vector τ'' using a transformation which incorporates past observations. This transformation represents temperature's effects over time.

3. Finally, use the vector τ'' as input into a transformation which models the non-linear relationship between τ'' and aggregate electricity demand $\mathbf{Y}$. The result of this third step is the matrix $\mathbf{T}$ used in the multiple regression model.

3.7.1 Coincident Weather Transformation

If τ is left untransformed during the first step of the temperature transformation Algorithm 1, then $\tau' = \tau$. The output of the coincident weather transformation would be dry-bulb temperature observations.

Feels Like **Transformation**

Humidity has been incorporated as a component of temperature observations frequently in prior work for two reasons. First, due to the effects of relative humidity on heat transfer and human comfort, residential customers may be more inclined to use cooling controls within their home on humid days. Second, dehumidifiers and air conditioners must operate more frequently on humid days to remove moisture from the air to maintain a comfortable environment indoors.

Similarly, wind chill on extremely cold days may have an added effect beyond dry-bulb temperature observations due to human perception of temperature. They may be more inclined to stay indoors and use heating controls within their home.

The *feels like* temperature transformation replaces dry-bulb temperature observations with heat index or wind chill values where applicable. Algorithm 2 provides an overview of the *feels like* transformation and Figure 3.10 visualizes its output.

Algorithm 2 Overview of *feels like* temperature transformation. Formulation of heat index can be found at [60] and wind chill found at [14].

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

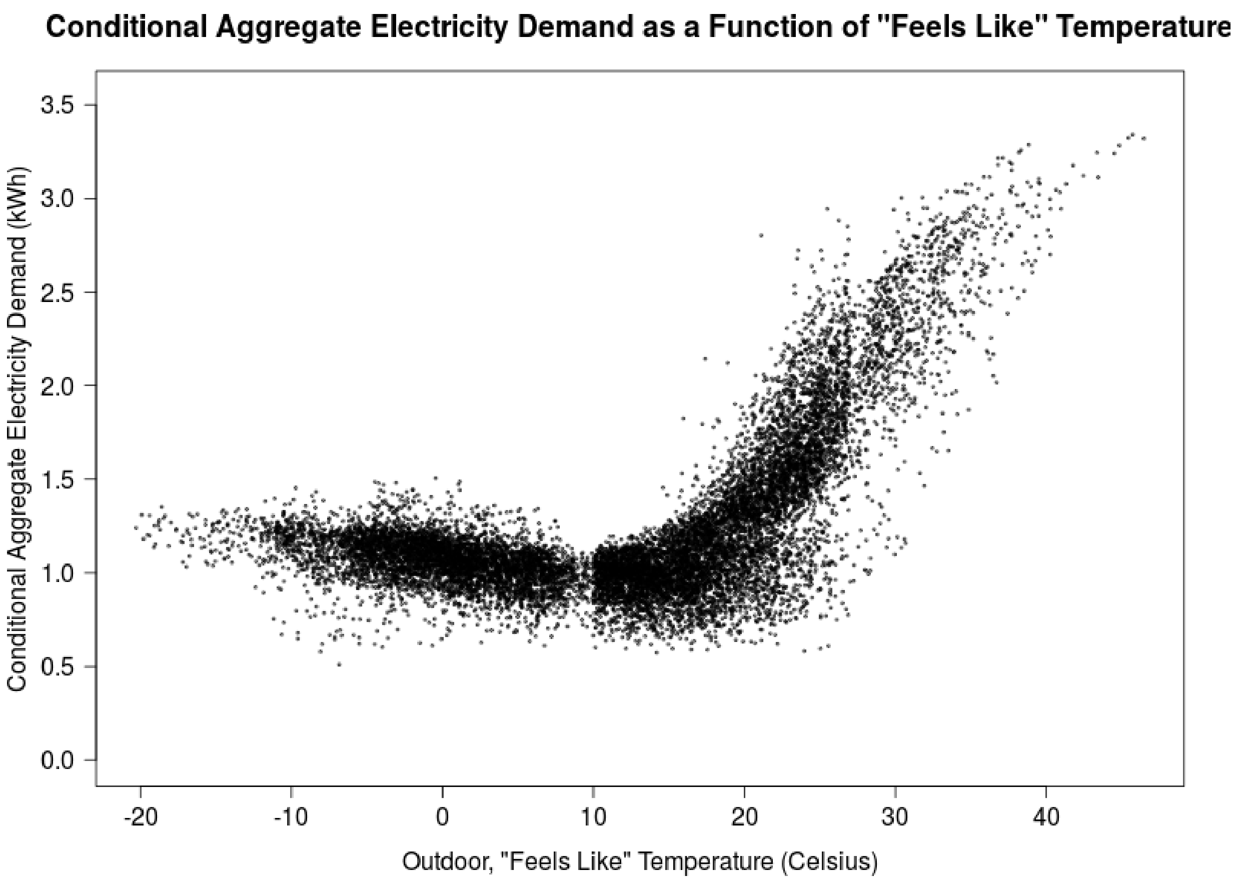

Figure 3.10: A scatter plot of *feels like* temperature observations plotted against aggregate electricity demand conditioned on other temporal explanatory variables.

Table 3.5: Up to 5 lags of dry-bulb temperature are correlated with the aggregate electricity demand at levels comparable to dry-bulb temperature at time i.

	Correlation with $\mathbf{y}_i$
τ_i	0.539
τ_{i-1}	0.551
τ_{i-2}	0.558
τ_{i-3}	0.558
τ_{i-4}	0.550
τ_{i-5}	0.533
τ_{i-6}	0.509
τ_{i-7}	0.477
τ_{i-8}	0.440
τ_{i-9}	0.400
τ_{i-10}	0.361
τ_{i-11}	0.328
τ_{i-12}	0.302

3.7.2 Temperature Effects Over Time

If τ' is left untransformed during the second step of the temperature transformation Algorithm 1, then $\tau'' = \tau'$. The output of the past weather observation transformation would be current observations from τ'.

Prior work mentioned in 2.2.2 and building management literature both suggest that lagged temperatures are important, due to heat stored in the building fabric [49]. To assess the importance of past temperature in predicting present electricity consumption, Table 3.5 shows the correlation of 0-12 lags of dry-bulb temperature τ 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. There may be a heat transfer delay through insulation, a build-up of temperature in building fabric, human perception of prolonged temperatures, or smoothing observations with a moving average may simply have better explanatory power.

Lagged Temperature Variables

The lagged observation transformation (3.3) 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 experienced to pass through a building's insulation. After the time delay, the household's thermal controls react.

$$\tau_i'' = \tau_{i-\ell}', \quad i = (1+\ell), ..., N \tag{3.3}$$

This interim transformation vector has ℓ fewer rows than the original vector of temperatures. Accordingly, rows $i = 1, ..., \ell$ from $\mathbf{Y}$, $\mathbf{X}$, and $\mathbf{V}$ must be removed from the sample. In our comparison we will compare all lags $\ell = 1, ..., 6$.

Temperature Moving Average

A *moving average* of recent temperatures, shown in (3.4), is included in our comparison to model the possibility that household thermal control systems are not reacting only to a specific temperature experienced at time i or some past time $i - \ell$. Instead, thermal control systems may be reacting to a number of recently experienced temperatures. The moving average also has the benefit of smoothing brief temperature extremes which, in the context of housing insulation and thermal controls, may not have descriptive value aside from their influence on the moving average. The variable L represents the number of recent temperatures used in the moving average.

$$\tau_i'' = \frac{\sum_{\ell=0}^{L-1} \tau_{i-\ell}'}{L}, \quad i = L, ..., N \tag{3.4}$$

The moving average transformation vector has $L - 1$ fewer rows than the vector τ' used as input. Accordingly, rows $i = 1, ..., (L-1)$ from $\mathbf{Y}$ and $\mathbf{X}$ must be removed from the sample. Figure 3.11 shows conditioned electricity demand plotted as a function of a six-hour moving average (i.e. $L = 6$) of dry-bulb temperature observations. $L = 6$ is selected empirically using the highest $Adjusted\, R^2$ as the selection criterion.

3.7.3 Non-Linear Temperature Effects

The temperature-sensitivity of electricity demand is dependent on the region of study. Ontario has four distinct seasons. In the context of electricity demand, summer and winter seasons are notable. Methods to cool residential households include central air conditioning,

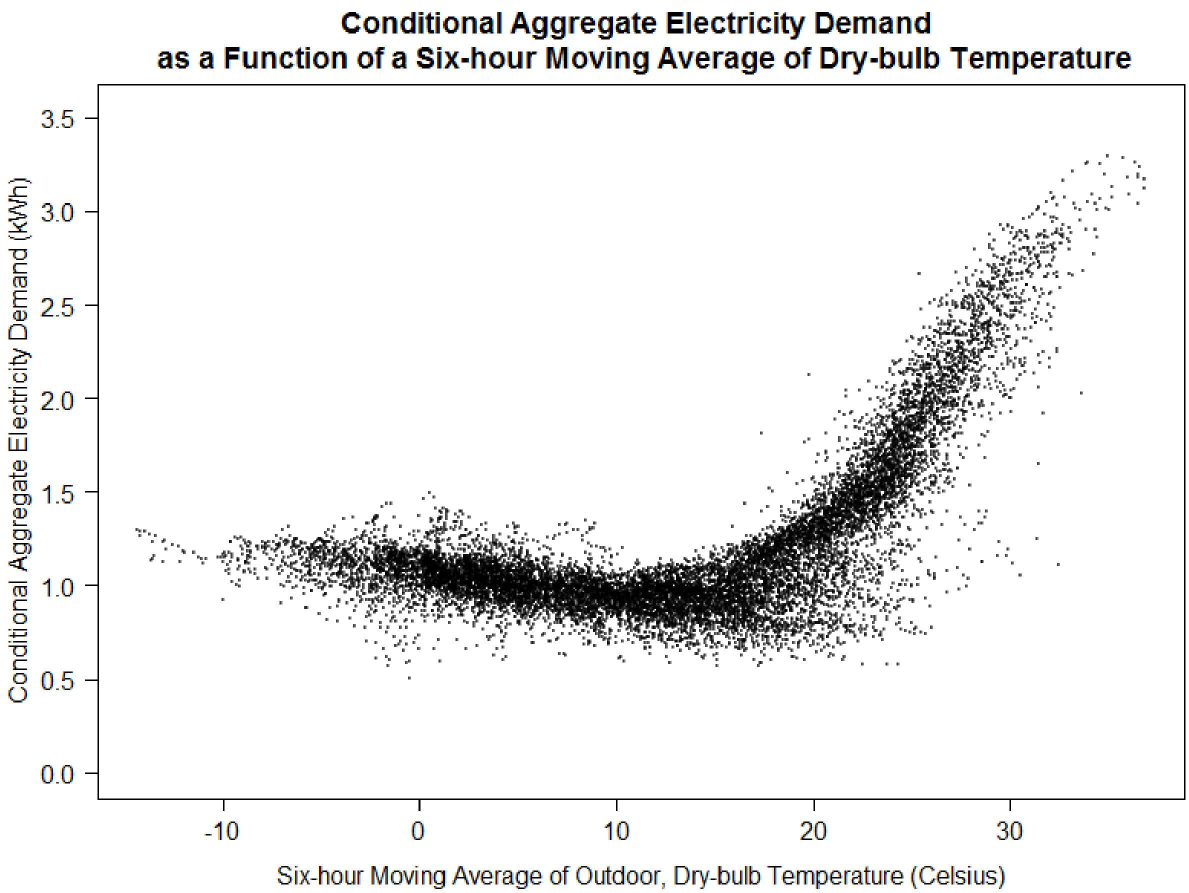

Figure 3.11: Conditioned aggregate electricity demand plotted as a function of temperature moving average with $L = 6$.

window unit air conditioning, and fans [45]. All of these cooling methods require electricity for use. Winter heating methods are roughly 66.7% natural gas, 13.9% wood, 10.4% electric, 5.4% heating oil, and 3.5% other [46]. Although Ontario's heating requirements are significant, only those homes with electric heating would increase electricity demand. There may also be increased electricity usage due to customers' tendency to stay indoors during during cold days. Given these details of thermal response methods in Ontario, we expect that residential cooling requirements have a more significant effect on electricity demand than heating requirements. Figure 3.8 suggests that this assumption is valid.

If τ'' is left untransformed during the third step of the temperature transformation Algorithm 1, then $\mathbf{T} = \tau''$. The output of the non-linearity transformation would be a one-column matrix of observations generated by the first two transformation steps. Figure 3.12 shows the coefficient estimate $\hat{\theta}$ fit for an untransformed vector of dry-bulb temperature observations. It easy to see that in a temperate region like Ontario, temperature's relationship with aggregate electricity demand is non-linear. The coefficient estimates for untransformed $\mathbf{T}$ has little meaning and is arguably misleading.

Switching Regression Transformation

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 [31]. In the context of electricity demand we can choose one break point and refer to it as a *switching regression* [41]. (3.5) shows the basis function which transforms τ'' into a column of $\mathbf{T}$ representing heating effects noticeable during low temperatures. Recall that the notation $\mathbf{t}_{i,1}$ represents the ith row and 1st column of the matrix $\mathbf{T}$.

$$\mathbf{t}_{i,1} = (\xi_{break} - \tau''_i)_+, \quad i = 1, ..., N \tag{3.5}$$

Similarly, (3.6) shows the basis function which transforms τ'' into a second column of $\mathbf{T}$ representing cooling effects noticeable during high temperatures.

$$\mathbf{t}_{i,2} = (\tau''_i - \xi_{break})_+, \quad i = 1, ..., N \tag{3.6}$$

Let the notation $(x)_+$ denote only positive values such that $(x)_+ := max(0, x)$. Let ξ_{break} be the temperature break point estimated using the `segmented` package in R [43, 44]. The fitted regression line for this switching regression transformation is shown in Figure 3.13.

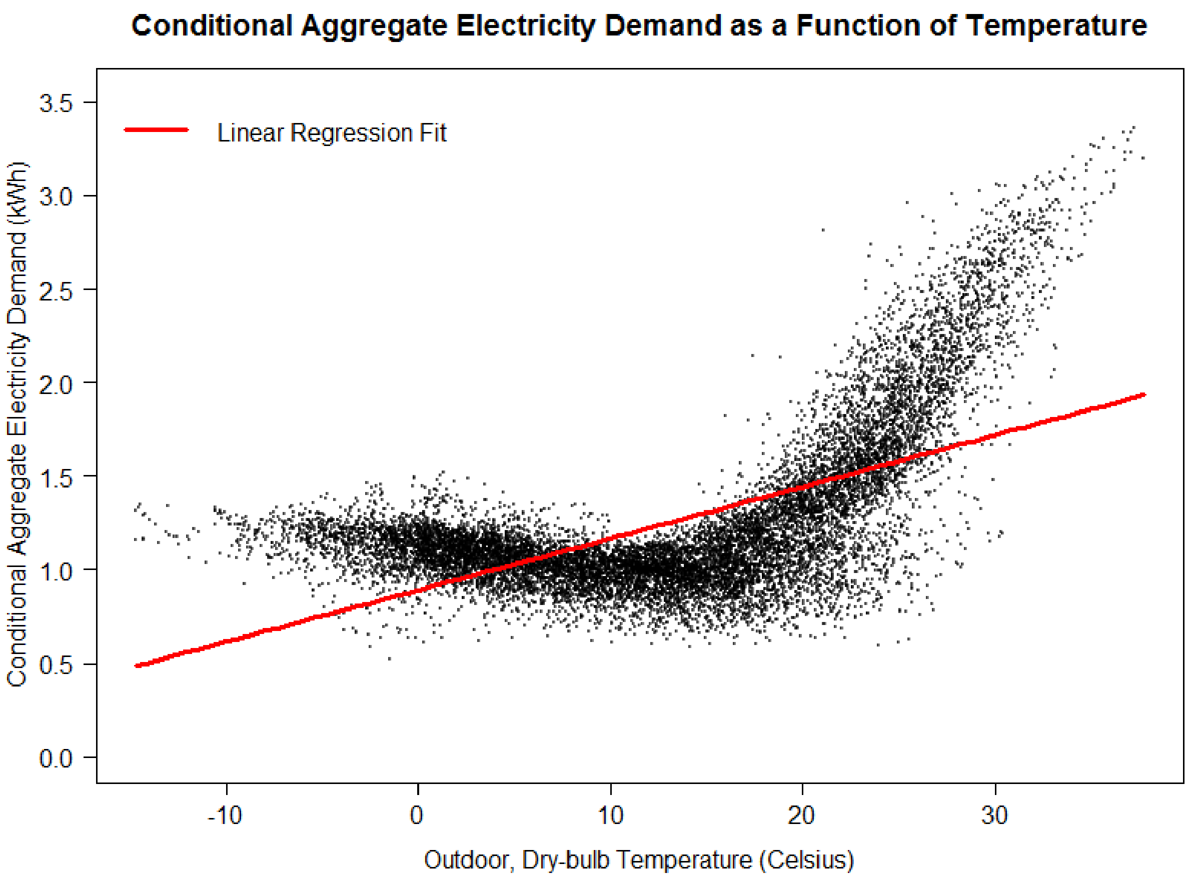

Figure 3.12: Linear regression line fit to untransformed, outdoor, dry-bulb temperature observations.

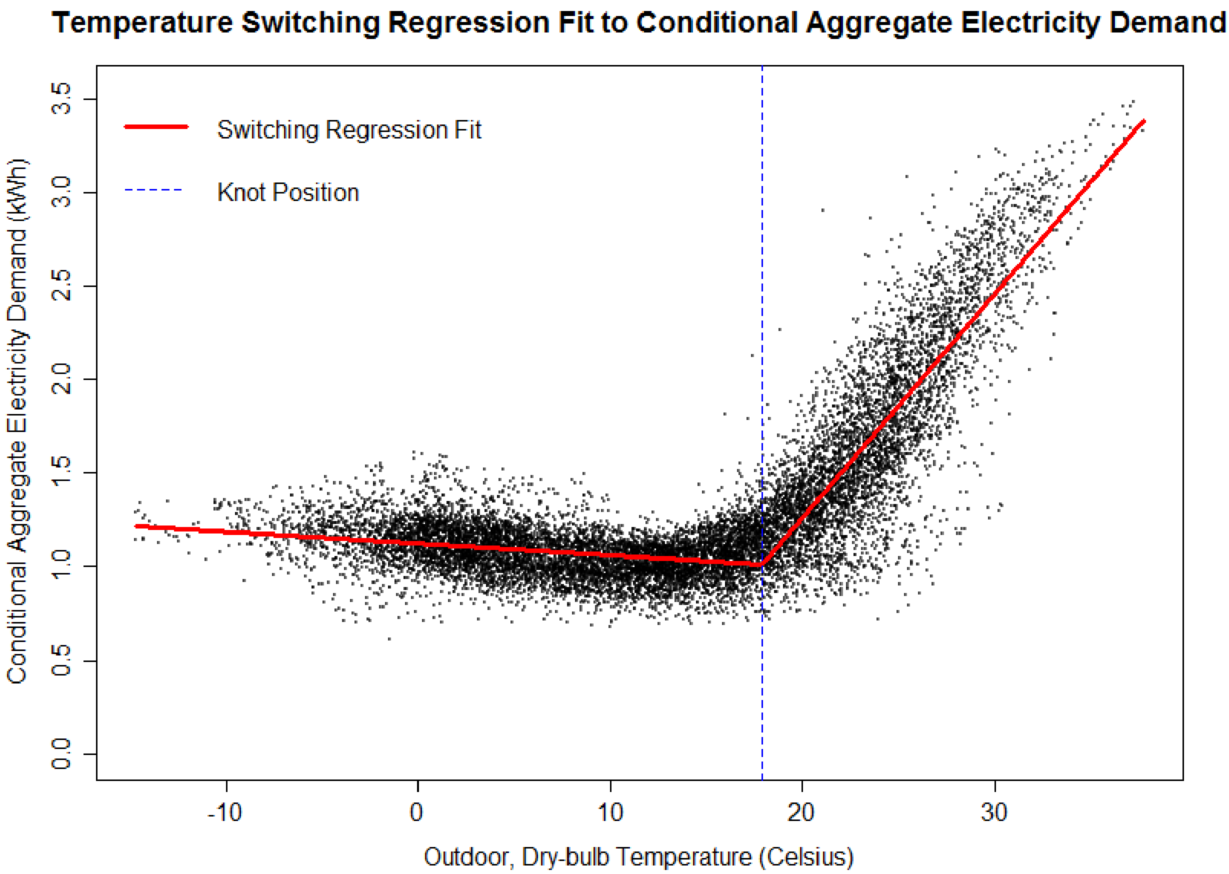

Figure 3.13: Fitted regression line for switching regression transformation of outdoor, dry-bulb temperature. Temperature break point at 17.9 °C.

Natural Cubic Splines Transformation

As we progress to increasingly non-linear transformations of τ'', the interpretation of the fitted model becomes less straightforward. Polynomial transformation of explanatory variables is a common econometric technique used to deal with non-linearity. However, high-degree polynomial transformations of explanatory variables occasionally result in unintuitive fits. The fit at explanatory variable extremes is particularly unstable. *Regression splines* are a common technique used to constrain the polynomial transformation [26].

Similar to switching regression, the goal of piecewise polynomial transformation is to break τ'' 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 basis function is used to transform observations in τ''. 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.

First, a number of knots are selected. Figure 3.14 illustrates $K = 3$ knots. Second, M is defined to be the order of desired continuity. 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' 1st derivative must be continuous at the knots. $M = 4$ places yet another restriction that the functions' 2nd derivative must be continuous at the knots. We have chosen order $M = 4$ splines, also known as *cubic splines*, which are widely used [26]. The first M columns of $\mathbf{T}$ represent the order of the spline (i.e. continuity restrictions), shown in (3.7).

$$\mathbf{t}_{i,m} = {\tau''}_i^{m-1}, \quad m = 1, ..., M \tag{3.7}$$

The subsequent K columns of $\mathbf{T}$ represent the polynomial basis function applied to each temperature region, shown in (3.8).

$$\mathbf{t}_{i,M+k} = (\tau''_i - \xi_k)_+^{M-1}, \quad k = 1, ..., K \tag{3.8}$$

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. Figure 3.14 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

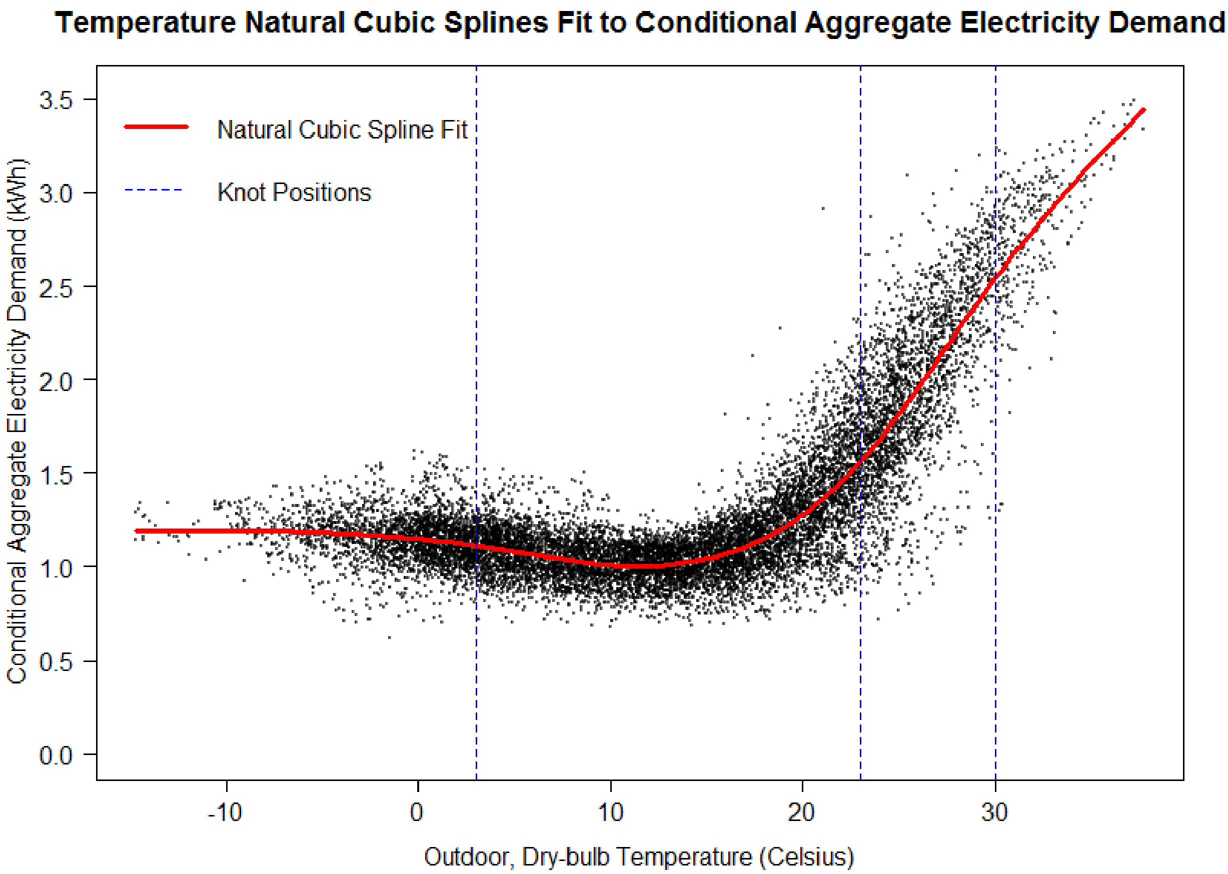

Figure 3.14: Natural cubic splines fit of outdoor, dry-bulb temperature fit to aggregate electricity demand. Knots are placed at 3 °C, 23 °C, and 30 °C.

as the selection criterion. A smooth transition between heating and cooling effects is visible around 17 °C. The lower tail exhibits heating effects saturation. The upper tail begins to show cooling effects saturation, but a plateau is never reached nor does it seem apparent in the data. This is a limitation of our data sample, having only two summers of weather. It is likely that cooling effects plateau at extremely hot temperatures not present during our sample period.

3.7.4 Complex Temperature Transformations

There are two complex temperature transformations which violate Algorithm 1. Both the heating/cooling degree-hour transformation and the exposure-lag-response transformation combine temperature transformation steps two and three. Both transform τ' directly to the matrix $\mathbf{T}$.

Cooling Degree-Hours and Heating Degree-Hours

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 (3.9).

$$\mathbf{t}_{i,1} = \sum_{\ell=0}^{L} (\xi_{break} - \tau'_{i-\ell})_+, \quad i = L, ..., N \tag{3.9}$$

Similarly, CDH is determined by summing the number of degrees above ξ_{break} during a window of L recent hours, shown in (3.10).

$$\mathbf{t}_{i,2} = \sum_{\ell=0}^{L} (\tau'_{i-\ell} - \xi_{break})_+, \quad i = L, ..., N \tag{3.10}$$

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. Because CDH and HDH values are approximately linear, shown in Figure 3.15, we do not fit a model using these values as input to a natural cubic splines transformation.

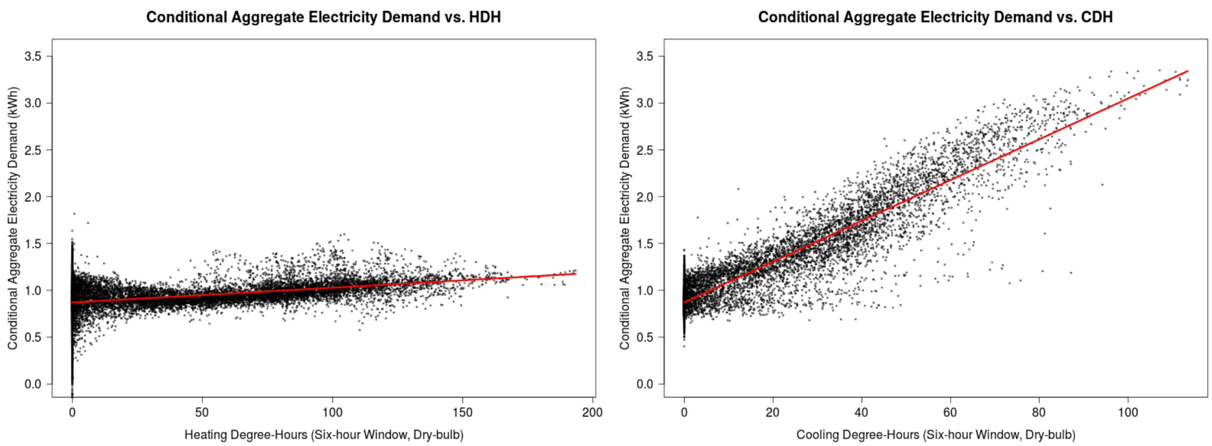

Figure 3.15: Both HDH and CDH window size is $L = 6$. *Left* Conditioned aggregate electricity demand plotted as a function of HDH. *Right* Conditioned aggregate electricity demand plotted as a function of CDH.

Exposure-Lag-Response Transformation

Recall Table 3.5 which shows the correlation of lagged temperature observations with electricity demand at time i. A finite *distributed lag* model was initially proposed in [3] to compute a weighted sum of past explanatory variable effects on a response variable. More recent implementations of this concept have come to be known as *distributed lag non-linear models* (DLNM) [21]. 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 τ'. 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*.

Let the exposure basis function be some non-linear transformation described in section 3.7.3. We transform τ' using a non-linear transformation but instead denote the resulting matrix as $\mathbf{Z}$, representing the exposure basis matrix with dimensions $N \times U$. Let U be the number of explanatory variables in the exposure basis matrix.

To represent the time dimension, we first define an $(L+1) \times 1$ vector $\pounds$ that represents lags of time, shown in (3.11).

$$\pounds = [0, ..., \ell, ..., L]^\top \tag{3.11}$$

A lag-response matrix $\mathbf{C}$ is created by transforming $\pounds$ with a lag basis function. For our temperature transformation comparison, we choose a cubic polynomial transformation as our lag basis function. The application of a polynomial transformation to $\pounds$ is shown in (3.12).

$$\mathbf{c}_{i,\bullet} = [\pounds_i^1, \pounds_i^2, ..., \pounds_i^D]^\top, \quad i = 1, ..., (L+1) \tag{3.12}$$

Let D be the degree of the polynomial transformation (i.e. the maximum exponent). Because we have chosen a cubic polynomial transformation, $D = 3$. Let the notation $\mathbf{c}_{i,\bullet}$ indicate all columns of row i of $\mathbf{C}$. Each column of $\mathbf{c}_{i,\bullet}$ is $\pounds_i$ raised to a higher power.

Assuming a maximum lag L, a lag dimension is added for each of the basis variables in $\mathbf{Z}$ to produce an $N \times U \times (L+1)$ array $\dot{\mathbf{R}}$. (3.13) shows the definition of the lag dimension in $\dot{\mathbf{R}}$ using elements from $\mathbf{Z}$.

$$\mathbf{r}_{i,j,\bullet} = [\mathbf{z}_{i,j}, ..., \mathbf{z}_{i-\ell,j}, ..., \mathbf{z}_{i-L,j}]^\top, \quad i = 1, ..., N; j = 1, ..., U \tag{3.13}$$

At the core of the exposure-lag-response association is the *cross-basis matrix*, which we will use as the temperature transformation $\mathbf{T}$. The cross-basis matrix simultaneously expresses the form of temperature exposure effects and non-linear weight of those effects on electricity demand a recent window of L hours. We first present a cross-basis array $\dot{\mathbf{H}}$ shown in (3.14).

$$\dot{\mathbf{H}} = (\mathbf{1}^\top \otimes \dot{\mathbf{R}}) \odot (\mathbf{1} \otimes G_{1,3}(\mathbf{C}) \otimes \mathbf{1}^\top) \tag{3.14}$$

Let $G_{i,j}$ be the operator permuting the indexes i and j and assume that a generic $i \times j$ matrix can be expressed as an $i \times j \times 1$ array. Let $\mathbf{1}$ be vectors of ones with appropriate dimensions. The notation $\otimes$ indicates the tensor product and $\odot$ indicates the entrywise product. $\dot{\mathbf{H}}$ will be an $N \times (U \cdot D) \times (L+1)$ array.

To create a cross-basis matrix to be used as $\mathbf{T}$ we sum along the third dimension of lags, shown in (3.15).

$$\mathbf{t}_{i,j} = \sum_{\ell=1}^{L+1} \mathbf{h}_{i,j,\ell}, \quad i = 1, ..., N; j = 1, ..., (U \cdot U) \tag{3.15}$$

In our comparison we used the package `dlnm` in R to create the exposure-lag-response cross-basis matrix [22, 21]. Figure 3.16 shows the effects of temperature over time as estimated by exposure-lag-response association. A natural cubic splines transformation is used as the

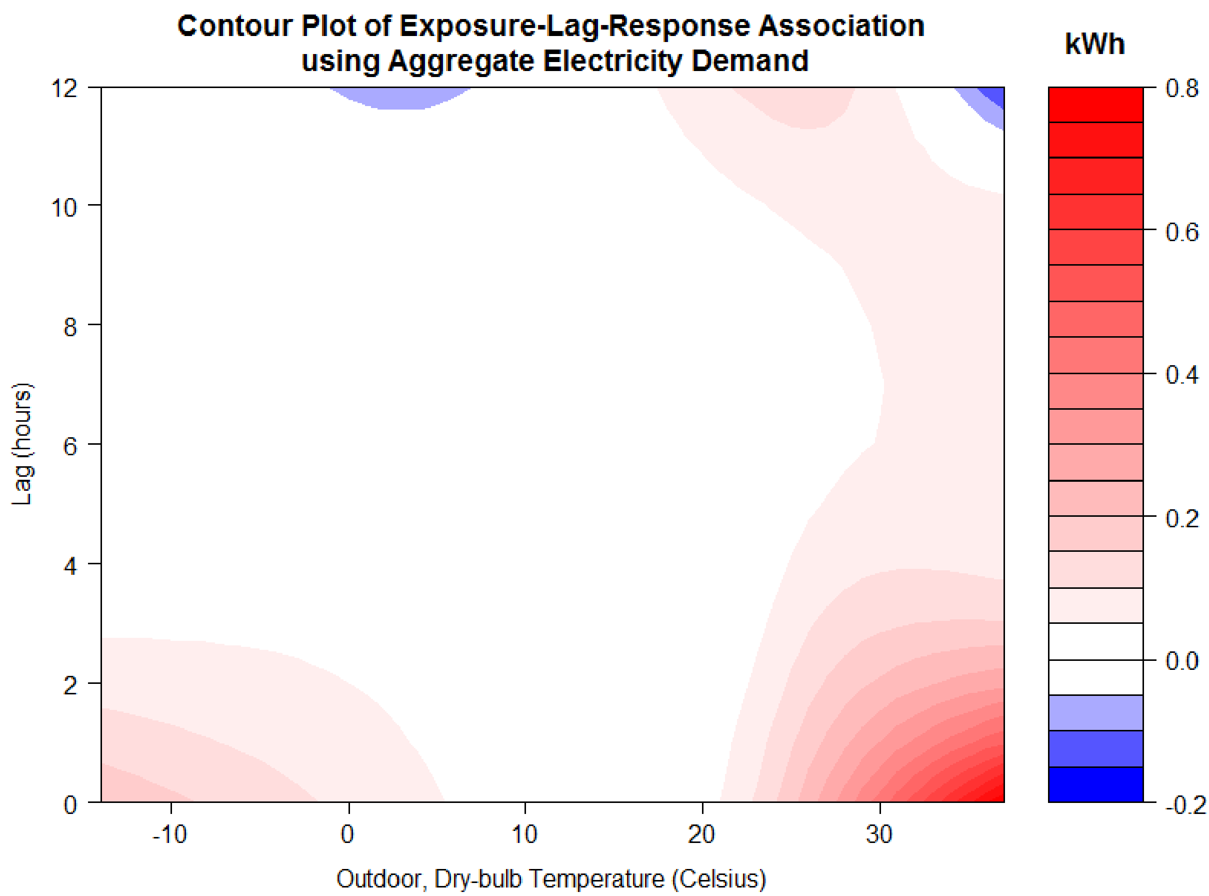

Figure 3.16: Exposure-lag-response association using a natural cubic spline exposure-response basis function, cubic polynomial lag-response basis function, and $L = 12$ lags considered.

exposure-response basis function and a cubic polynomial transformation is used as the lag-response basis function. The interpretation of this figure is that exposure to temperature has an additive effect on aggregate electricity demand over time. The effect associated with temperature at time i is added with effect associated with temperature at time $i-1$, and so on to $i-L$. In Figure 3.16 the maximum number of lags to consider effects for is set to $L = 12$. Any temperature exposures beyond that are assumed to have no effect on electricity demand at time i. Our interpretation of this transformation is discussed with greater detail in section 3.9.2.

3.8 Temperature Transformation Evaluation Criteria

3.8.1 Variance Explained

Our primary evaluation criteria will be a measure of variance explained, $Adjusted\, R^2$, defined in (2.8). We also check the value of BIC defined in (2.9) to guard against over fitting. 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.

3.8.2 Residual Analysis

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. Recall (2.2) which defines an $N \times 1$ vector of residuals $\mathbf{e}$. Ideally, $\mathbf{e}$ should be normally distributed, mean zero, and independent of each explanatory variable. To address heteroscedasticity increasing with electricity demand, the temperature transformation comparison was also run fitting models to log transformed electricity demand. The results did not yield any additional insight beyond the untransformed response $\mathbf{Y}$ and are briefly discussed in Appendix C.

3.9 Temperature Transformation Explanatory Power

The temperature transformation comparison results are shown in Table 3.6. The first three columns show how temperature transformations from the three categories are combined.

The *Adjusted* R^2 column is our primary evaluation criterion. It measures the amount of variance explained by the model. *BIC* and *DW* columns provide secondary measures of model complexity and serially correlated errors. The final two columns report predictive accuracy using the time series cross-validation process described in section 2.1.5. Discussion of the temperature transformation steps is provided in sections 3.9.1 through 3.9.3. Though combinations of temperature transformation steps each produce incremental improvements, the proportion of variance explained by all temperature transformations is notable when compared to the three baseline models: the null model, non-temperature explanatory variables only, and a model which applies no transformation to temperature observations. The $AdjustedR^2$ for a model which applies no temperature transformation is 0.580. Across all temperature transformations the resulting models' $AdjustedR^2$ values range from 0.790 to 0.911.

3.9.1 Discussion of Coincident Weather and Temperature

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 poorer out-of-sample predictive power than dry-bulb temperature. Our interpretation of this mixed result is that our *feels like* interim temperature transformation has little added value over simply using dry-bulb temperature observations.

3.9.2 Discussion of Delayed Temperature Effects

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. It has the highest correlation as shown in Table 3.5. 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.

Of the three temporal transformations that use a window of recent temperature observations, both CDH/HDH and the moving average transformations showed that a six-hour window of temperature observations yielded the highest *Adjusted* R^2. We set the number of lags used in exposure-lag-response to be six hours as well to be comparable with CDH/HDH

Table 3.6: Results of temperature transformation comparison.

Temporal Transform	**Weather Transform**	**Non-Linearity Transform**	**Adj.** R^2	**BIC**	**DW**	**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

and moving average transformations. 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.

Despite the prevalence of the CDH/HDH metric in literature, the more simplistic moving average transformation has greater explanatory power and predictive power. This may be caused by the smoothing effect that moving average has on the temperature explanatory variable.

Finally, we use the exposure-lag-response transformation in our analysis to evaluate a technique 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. Figure 3.16 was created using a natural splines form of temperature effects over 12 hours. As expected, low- and high-temperature extremes have a significant effect on aggregate electricity demand. However, the form of its impact over time is not intuitive. The most unintuitive result is that of $> 35\,°\mathrm{C}$ and $1 - 5\,°\mathrm{C}$ experienced 11-12 hours in the past reducing the estimated electricity demand at hour i. Considering the effect of temperate at time $i-12$ for other temperatures, one would expect that $> 35\,°\mathrm{C}$ and $1 - 5\,°\mathrm{C}$ have no effect on electricity demand. Our result seems symptomatic of overfitting the time dimension. The six-hour lag-response used for comparison to CDH/HDH and moving average is even more problematic, shown in Figure 3.17. This unintuitive result is likely a symptom of lag-response basis function having three degrees of freedom over six lags of time. For a short window of time one might select a simpler linear or identity lag-response basis function. However, they would have little explanatory value over moving average. For these reasons, we do not believe the exposure-lag-response transformation has added descriptive value for our case study.

3.9.3 Discussion of Non-Linear Structure

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, estimated using the `segmented` package in R, is $17.9\,°\mathrm{C}$ [44].

Natural cubic splines do provide more flexibility in modelling the temperature's non-

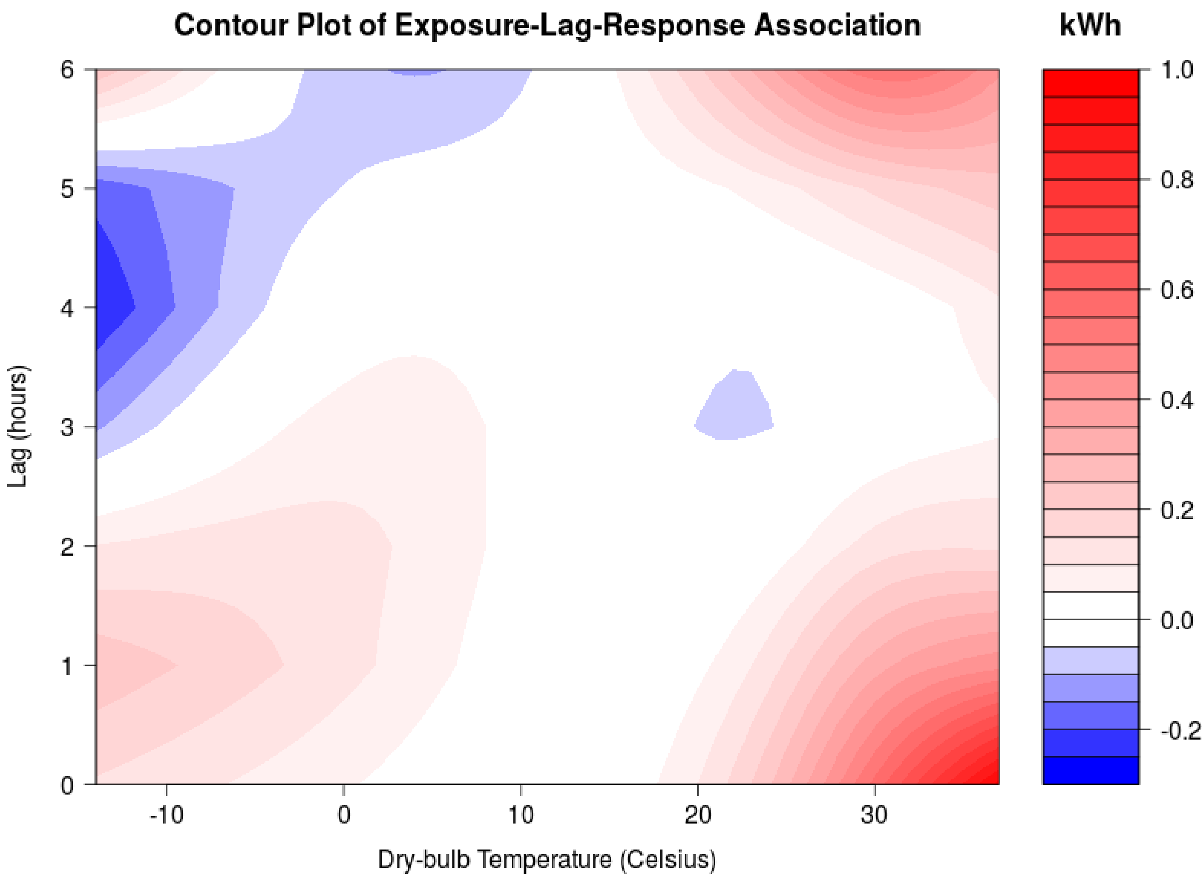

Figure 3.17: Exposure-lag-response association using a natural cubic spline exposure-response basis function, cubic polynomial lag-response basis function, and $L = 6$ lags of temperature exposure considered.

linear relationship with aggregate electricity demand and has higher $Adjusted\, R^2$ than switching regression. The assumption of linearity beyond the range of observations is well-suited to model saturation of heating effects at low temperatures, not reflected in switching regression. However, this assumption of linearity beyond observed high temperatures observed in our data is not well-suited to cooling effects saturation.

3.10 Residual Analysis

The first assumption of linear regression analysis is that errors be normally distributed with mean zero. Figure 3.18 shows the residuals of a multiple regression using a dry-bulb, six-hour moving average, natural cubic splines transformation to generate **T**. Residuals are normally distributed with mean=0.

The residual plot checks for statistical independence of the errors and the potential for them to be correlated with explanatory variables. We are primarily concerned with time and temperature. Figure 3.19 shows residuals as a function of dry-bulb temperature observations. This plot shows heteroscedasticity in the model's residuals, with increased variability at warm temperatures.

Figure 3.20 shows residuals as a function of time. The plot shows heteroscedasticity in the model's residuals, with greater variance associated with summer and winter seasons. There are also indications of greater variance associated with certain days. The variance associated with summer and winter seasons is likely a result of its heteroscedasticity associated with temperature and season's collinearity with dry-bulb temperature. This result is supported quantitatively by the Durbin-Watson statistic in Table 3.6. A $DW < 1$ value indicates positive serial correlation of residuals.

The final residual plot in Figure 3.21 shows residuals as a function of the estimated response variable, aggregate electricity demand. The plot illustrates that variance of residuals increases with larger values of the response variable. This too is likely a result of heterosedasticity associated with dry-bulb temperature. Because warm temperatures are associated with greater variance of residuals and also associated with higher electricity demand, it follows that greater residual variance is associated with higher electricity demand. Figures 3.19, 3.20, and 3.21 indicate that HAC standard errors, discussed in section 2.1.4, must be used when performing hypothesis tests in our TOU pricing case study.

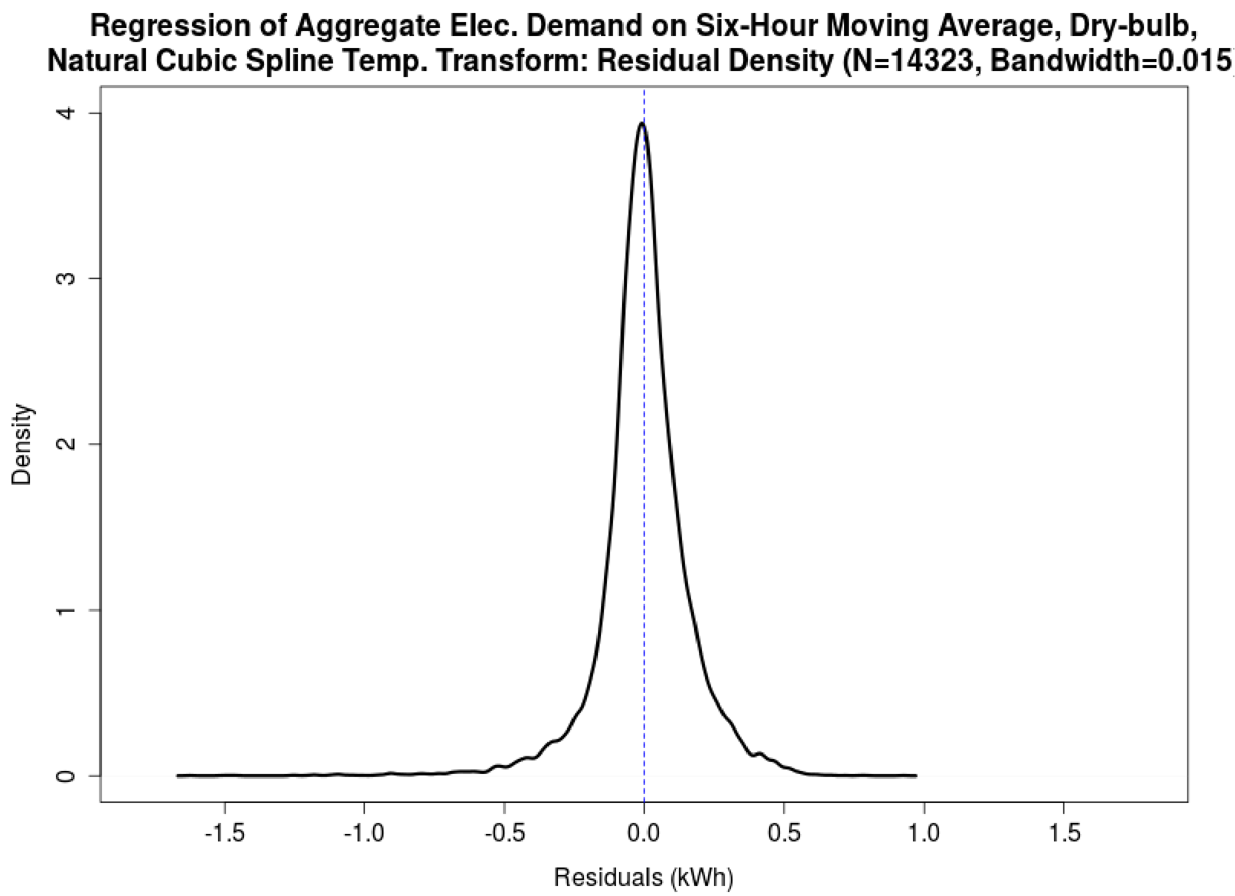

Figure 3.18: Density of residuals, resulting from a comparison model using dry-bulb, six-hour moving average, and natural cubic splines to generate the temperature transformation matrix $\mathbf{T}$.

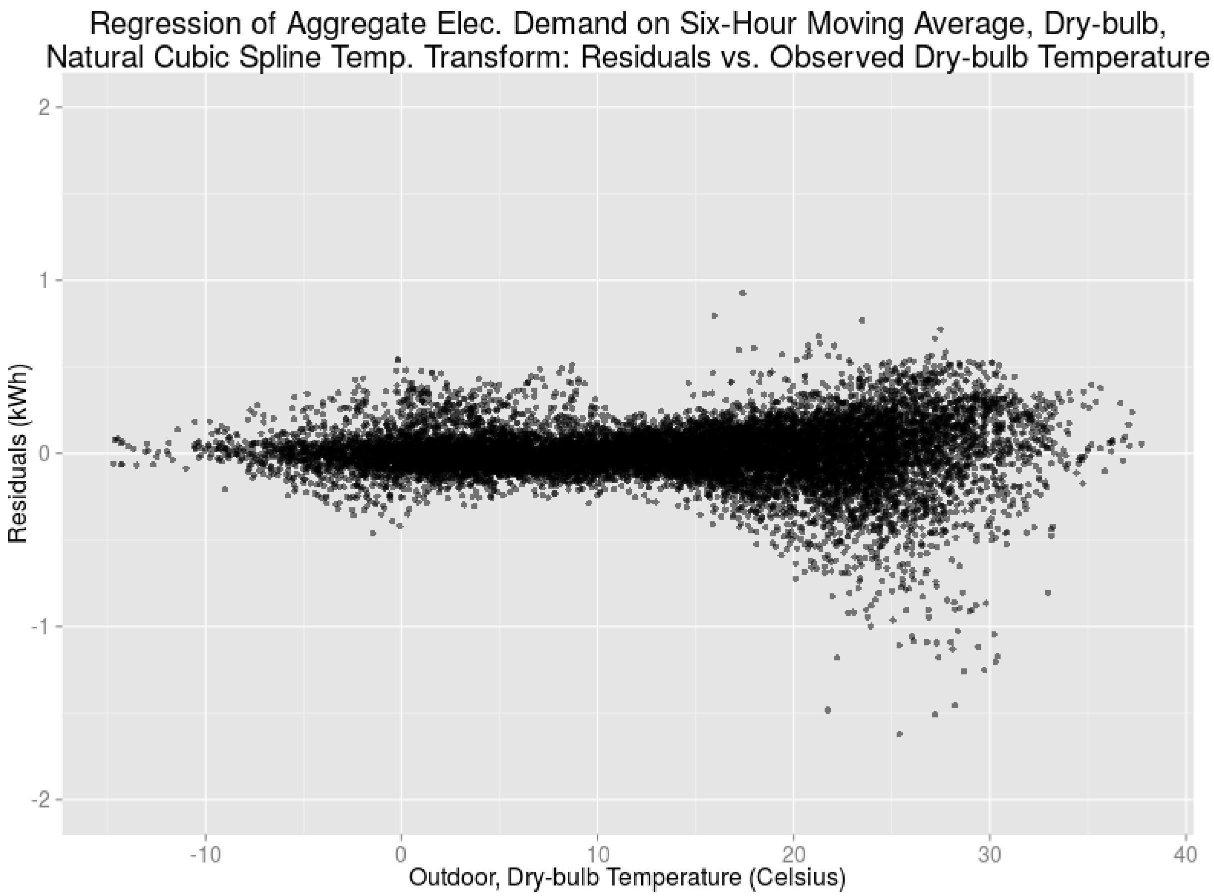

Figure 3.19: Residuals as a function of dry-bulb temperature observations, resulting from a comparison model using dry-bulb, six-hour moving average, and natural cubic splines to generate the temperature transformation matrix $\mathbf{T}$.

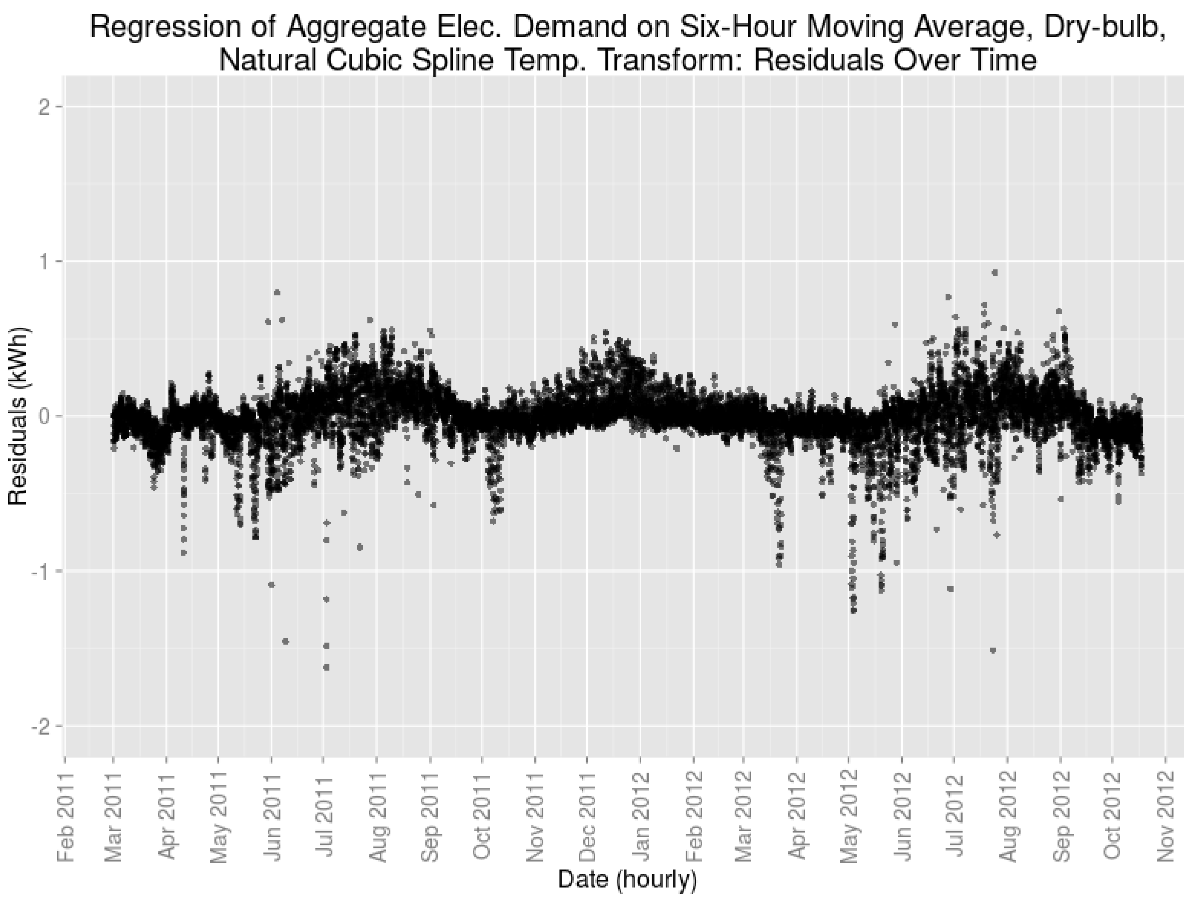

Figure 3.20: Residuals as a function of time, resulting from a comparison model using dry-bulb, six-hour moving average, and natural cubic splines to generate the temperature transformation matrix $\mathbf{T}$.

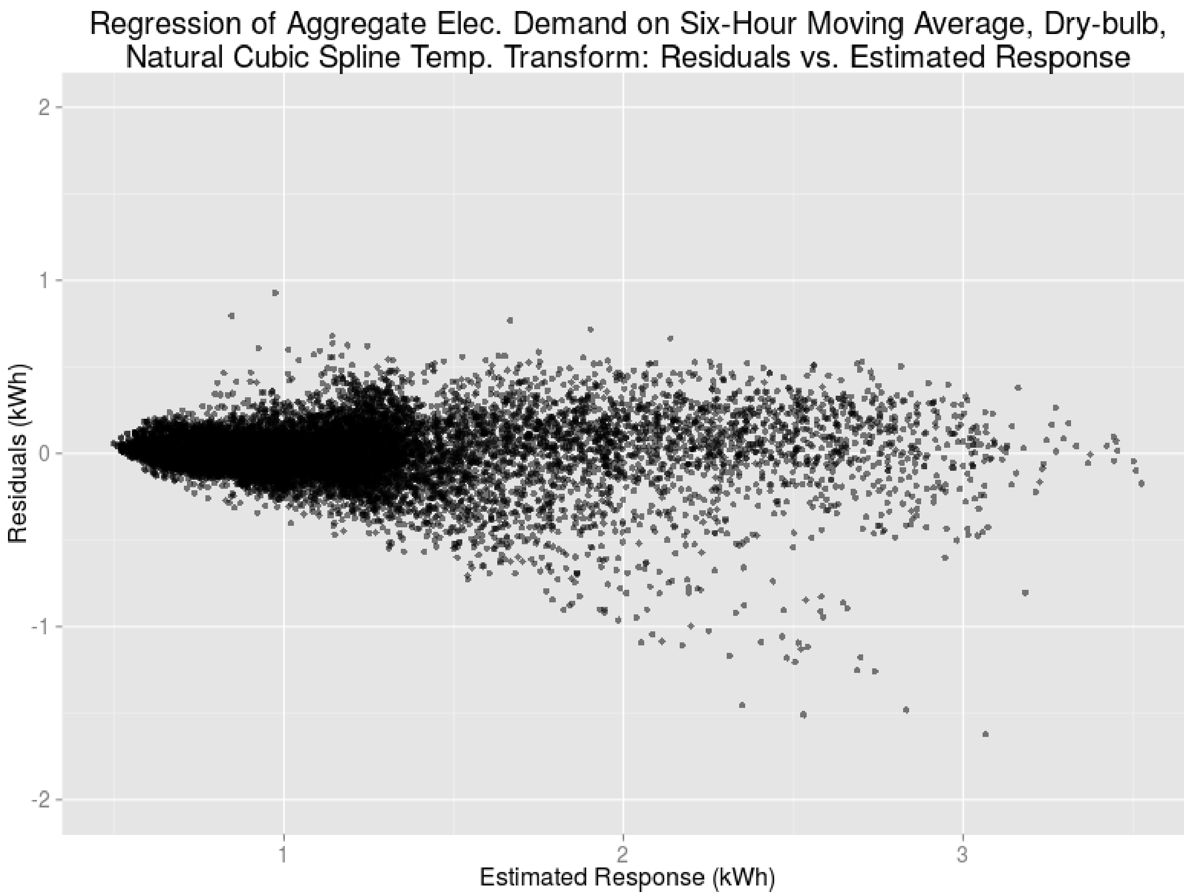

Figure 3.21: Residuals as a function estimated response, resulting from a comparison model using dry-bulb, six-hour moving average, and natural cubic splines to generate the temperature transformation matrix $\mathbf{T}$.

Chapter 4

Case Study: Effects of Mandatory Time-of-Use Billing in South West Ontario

In section 2.1 we describe the rationale for modelling time, price, and weather explanatory variables as the matrices **X**, **V**, and **T**. We held explanatory variable matrices **X** and **V** constant for our temperature transformation comparison in chapter 3. Based on our results in section 3.9 we conclude that the dry-bulb, six-hour moving average, natural cubic splines transformation **T** has the best balance of explanatory power and interpretability.

The motivation for this case study is to quantify change in aggregate electricity demand associated with mandatory TOU pricing. To model effects associated with TOU pricing with hourly fidelity, the matrices **X** and **T** will be held constant. We will use backward selection, ANOVA, and HAC standard errors (described in sections 2.1.3 and 2.1.4) to remove insignificant variables from a saturated explanatory variable matrix **V**. The remaining, significant explanatory variables in **V** are used as components of the multiple regression model in a "what if" analysis. We use our results from the "what if" analysis to quantify the change in demand associated with TOU pricing.

4.1 Results from Prior Work

A recent literature review performed by the authors of [50] reviews 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 -5.00%. An earlier TOU literature review [16] covering 12 TOU pilot studies concluded that TOU pricing induces a -3% to -6% change in residential on-peak demand.

We carry out a similar TOU literature review with a focus on mandatory TOU pricing implementations, similar to Ontario's TOU framework. Several of the studies in our review are impact analysis studies commissioned by the Ontario Energy Board. Notably, we quantify the impacts of TOU pricing in section 4.2 in a manner consistent with [48]. We also consider several controlled experiments and pilot studies. Table 4.1 summarizes the results of our literature review.

Table 4.1: Results from prior TOU studies.

Study	**Location**	**Pilot Study**	**Mandatory**	**Season**	**Total Change (%)**	**On-Peak (%)**	**Mid-Peak (%)**	**Off-Peak (%)**	**Weekend Off-Peak (%)**
[27]	Ontario, CAN	Yes	No	summer	-3.30	-3.70	*NR*	*NR*	*NR*
[35]	Arizona, USA	Yes	No	summer	-3.17	-8.84	-3.95	+2.86	*NA*
[55]	Ontario, CAN	Yes	No	summer	-6.00	-2.40 (*NS*)	*NR*	*NR*	*NR*
[62]	California, USA	Yes	No	full year	*NR*	-9.02	*NA*	+6.51	*NA*
[33]	Arizona, USA	Yes[a]	No	summer	*NR*	-20.65[d]	*NA*	*NR*	*NA*
[32]	north east USA	No	Yes[a]	summer	-3.14	-6.09	*NA*	-2.00	*NA*
[32]	north east USA	No	Yes[b]	summer	+0.39	+1.16	*NA*	+0.06	*NA*
[32]	north east USA	No	Yes[c]	summer	+2.64	+3.11	*NA*	+2.4	*NA*
[18]	Ontario, CAN	No	Yes	summer	0 to -0.45[e]	-2.60 to -5.70	decrease	increase	*NR*
[18]	Ontario, CAN	No	Yes	winter	0 to -0.45[e]	-1.60 to -3.20	decrease	increase	*NR*
[47]	Ontario, CAN	No	Yes	full year	0.66 (*NS*)	-2.80	-1.39	+0.16 (*NS*)	+2.21
[48]	Ontario, CAN	No	Yes	summer	0 to -0.10	-3.30	-2.20	+1.20	+1.90
[48]	Ontario, CAN	No	Yes	summer shoulder	*NR*	-2.20	-1.50	+1.50	+1.40
[48]	Ontario, CAN	No	Yes	winter	*NR*	-3.40	-3.90	-2.50	-1.20
[48]	Ontario, CAN	No	Yes	winter shoulder	*NR*	-2.10	-2.30	-1.10	+0.50 (*NS*)
[36]	ITA	No	Yes	Jan–Jun	*NR*	-0.83	*NA*	*NR*	*NA*
[37]	Anhui, CHN	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] three highest peak days of summer
[e] annual

Our first observation is that results from opt-in experiments and pilot studies such as [27, 35, 55, 62, 33] are often more pronounced than mandatory studies such as [18, 47, 48]. We expect our results to be less pronounced, similar to the latter studies.

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 [32] showed the opposite effect.

Finally, we observe that several tiered implementations of TOU to high-use customers first showed substantial flexibility to shift demand [32, 33]. However, [33] evaluated the effects of TOU pricing during the three highest peak days of summer. Though the authors do not classify it as such, this is very similar to critical peak pricing, in which more pronounced conservation effects are common.

4.2 Methodology

Based on our temperature transformation comparison results presented in section 3.9, we set the temperature transformation matrix $\mathbf{T}$ to be a dry-bulb, six-hour moving average, natural cubic splines transformation. We show (2.1) below again for clarity.

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

$\mathbf{Y}$ is an $N \times 1$ vector representing aggregate electricity demand. Because collinearity of temperature and seasonal explanatory variables are not an issue of concern when analyzing the effects of TOU, we add a categorical explanatory variable for month to $\mathbf{X}$, supported by ANOVA. Explanatory variables in the matrix $\mathbf{X}$ are then held constant.

We saturate the matrix $\mathbf{V}$ with explanatory variables related to TOU pricing. We include all two-way and three-way interactions combining a TOU billing indicator, working day, hour-of-day, and utility rate season. These variable interactions provide the necessary degrees of freedom to explain the effects of TOU billing for each hour of day. The categorical variables in $\mathbf{X}$ and $\mathbf{V}$ before backward selection 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.
- $\mathbf{v}_{\bullet,p=1}$ 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 $\mathbf{x}_{\bullet,p=25}$ are indicators representing the hour-of-day $\times$ rate season interaction.
- $\mathbf{v}_{\bullet,p=26}$ through $\mathbf{x}_{\bullet,p=48}$ 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 $\mathbf{x}_{\bullet,p=74}$ are indicators representing the hour-of-day $\times$ working day $\times$ rate season interaction.
- $\mathbf{v}_{\bullet,p=75}$ through $\mathbf{x}_{\bullet,p=97}$ are indicators representing the hour-of-day $\times$ working day $\times$ TOU active interaction.
- $\mathbf{v}_{\bullet,p=98}$ through $\mathbf{x}_{\bullet,p=120}$ are indicators representing the hour-of-day $\times$ rate season $\times$ TOU active interaction.
- $\mathbf{v}_{\bullet,p=121}$ is an indicator representing representing the working day $\times$ rate season $\times$ TOU active interaction.

We chose categorical variables representing the structure of utility rate seasons and TOU price periods rather than a continuous value representing the price paid for electricity each hour. To model customer price sensitivity to varying TOU rates, additional modelling techniques such as price elasticity analysis are required. Because our sample contains only two summers and TOU prices do not change during the course of one summer season, our ability to quantify price elasticity is limited. Categorical variables representing the TOU pricing structure are adequate and useful for our study's objective, the quantification of electricity demand conservation and peak-shifting.

4.2.1 Backward Explanatory Variable Selection

Using the explanatory variable selection process known as backward selection, initially described in section 2.1.3, we remove variables from the saturated matrix $\mathbf{V}$ to create more parsimonious model [31]. We will keep only the explanatory variable interactions added for mandatory TOU pricing that can be justified by ANOVA. The process of backward selection is shown in Algorithm 3. Using backward selection this manner, we drop the

Algorithm 3 Overview of backward variable selection using ANOVA with HAC standard errors.

1. Let $\mathbf{V}$ start as a matrix of all three-way and two-way interactions related to TOU pricing.
2. Run ANOVA over the coefficient estimates $\hat{\beta}$, $\hat{\omega}$, and $\hat{\theta}$ using HAC standard errors.
3. ANOVA F-tests are sorted in a descending order for each explanatory variable group.
4. If the p-value of a value of F is > 0.05, that explanatory variable is removed from $\mathbf{V}$.
5. The model estimate (2.1) is updated with the reduced explanatory variable matrix $\mathbf{V}$, and this backward selection algorithm returns to step 2. Otherwise, if all F-tests are significant at p-value ≤ 0.05, then the model is considered parsimonious and this backward selection algorithm ends.

following interactions from $\mathbf{V}$: working day $\times$ rate season $\times$ TOU active, hour-of-day $\times$ working day $\times$ TOU active, and working day $\times$ TOU active.

We also re-ran our temperature transformation comparison with the more complex matrix $\mathbf{V}$. $Adjusted R^2$ values are higher but the intuition and our conclusions from the analysis remain the same. The model used in this case study yields $Adjusted\, R^2 = 0.935$.

4.2.2 Counterfactual "What If" Analysis

An impact analysis commissioned by the Ontario Energy Board [48] formulates a "what if" counterfactual analysis that we will use to measure the effects of mandatory TOU pricing, shown in Algorithm 4. 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 not able to perform a "what if" counterfactual analysis for the winter rate season. We are only able to carry out the "what if" counterfactual analysis for summer utility rate season.

Algorithm 4 "What if" counterfactual analysis used to quantify the effects mandatory TOU electricity pricing.

1. Fit a model to the entire sample of data. In our study, the sample runs from March 1, 2011 – October 17, 2012.
2. For the summer utility rate season, select sample data 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.

4.3 Results from South West Ontario Data Set

We plot the estimated effects of TOU pricing for each hour of a working day in the summer utility rate season in Figure 4.1. Similarly, we plot the estimated effects of TOU pricing for each hour of a non-working day in the summer utility rate season in Figure 4.2. Table 4.2 shows the hourly effects averaged by TOU price period. Effect size is reported both in terms of kWh impact and percentage impact.

Table 4.2: Estimated change in the aggregate electricity demand for each price period under TOU pricing.

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

Finally, we report the estimated change to peak-to-average ratio which measures how extreme daily demand peaks are typically. Each day's peak-to-average ratio during a sample period is defined as the peak demand for each day divided by the average demand during that day. The result is a vector of daily peak-to-average ratios for each day in the sample. We can take the mean of daily peak-to-average ratios for to summarize the sample period. The 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, had TOU pricing been active. This represents an estimated change of -0.844% to the peak-to-average ratio, with a 95% confidence interval of ±0.6%.

4.4 Discussion

Because our sample of data is from one local distribution company in south west Ontario, we acknowledge that our results are only 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. Though we make an effort to accurately model the effects of temperature throughout chapter 3, our ability to model

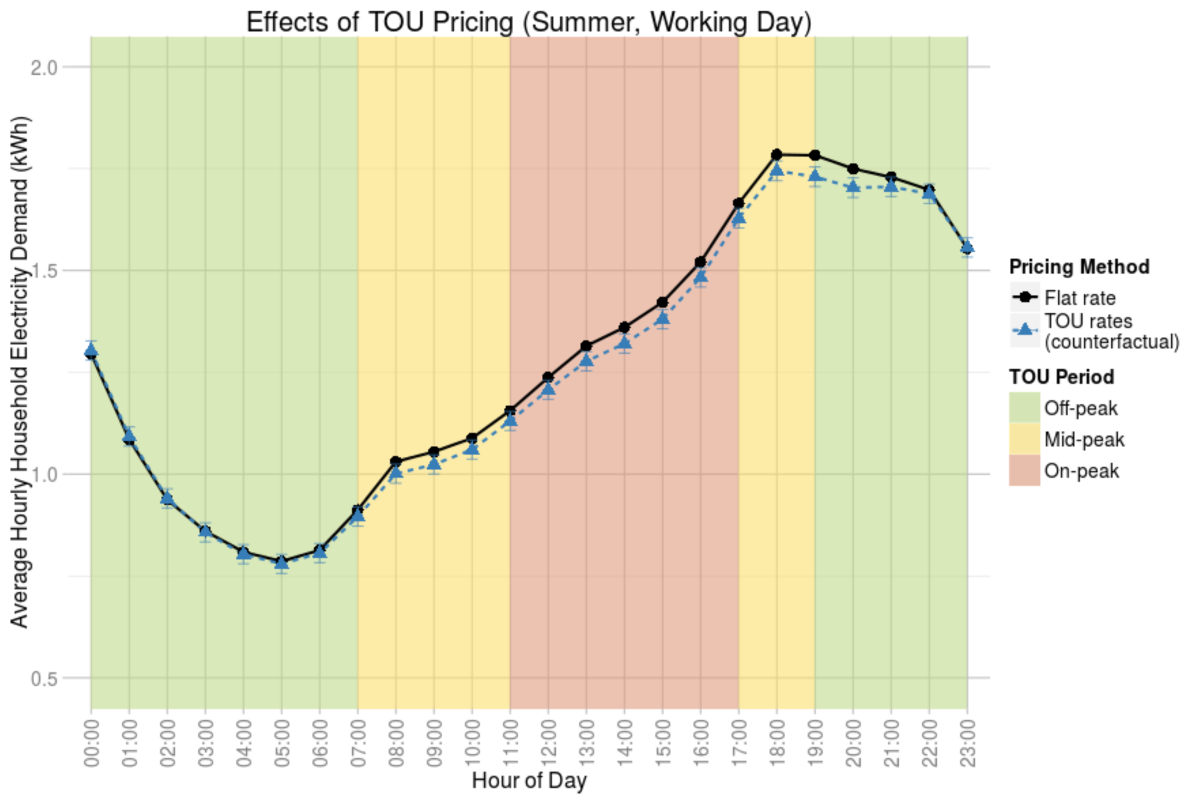

Figure 4.1: The hourly effects of a "what if" counterfactual analysis estimated using summer 2011 data from our sample. The observed data is the solid, black line, indicating the mean of observed demand for each hour of working days. The dotted blue line indicates the mean of estimated demand for each hour of working days, had TOU billing been in place. A 95% confidence interval is also plotted for each hour.

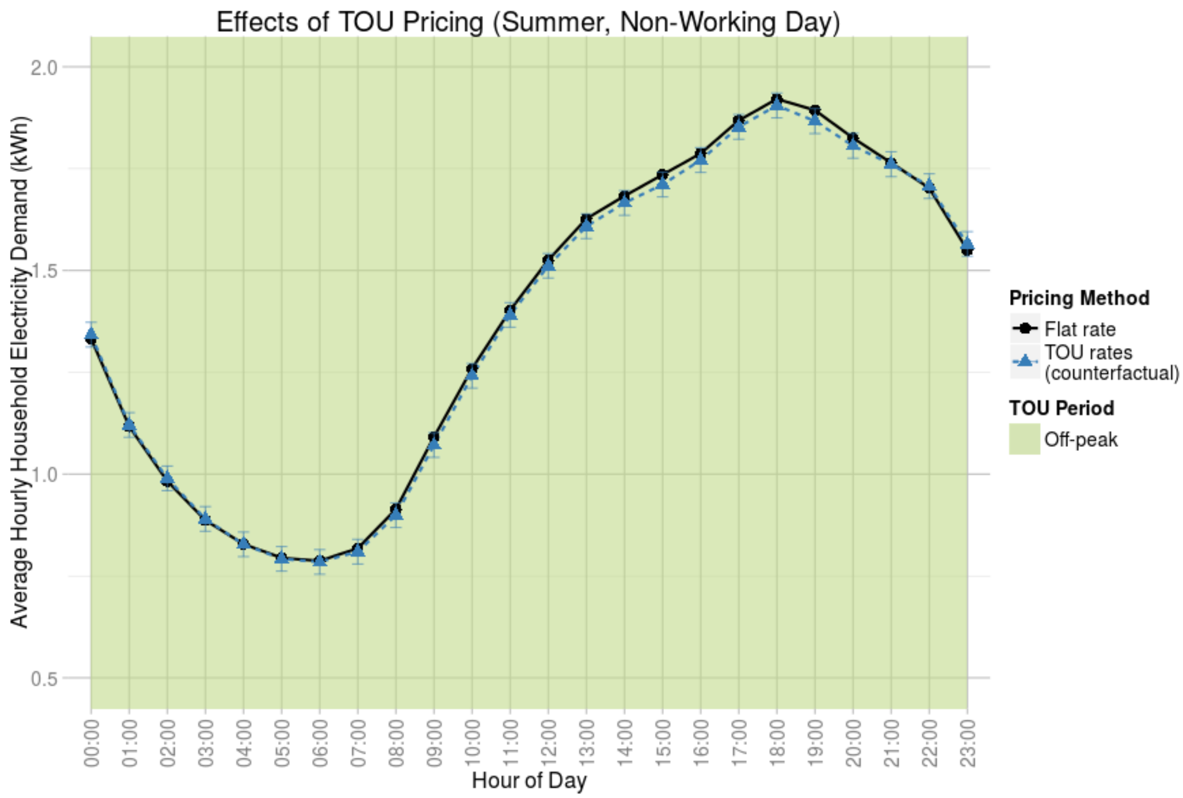

Figure 4.2: The hourly effects of a "what if" counterfactual analysis estimated using summer 2011 data from our sample. The observed data is the solid, black line, indicating the mean of observed demand for each hour of non-working days. The dotted blue line indicates the mean of estimated demand for each hour of non-working days, had TOU billing been in place. A 95% confidence interval is also plotted for each hour.

extremely warm temperatures are limited due to the range of observed temperatures in our sample, as discussed in section 3.7.3.

Both the slight decrease in peak-to-average ratio and the hourly demand reduction across all TOU price periods indicate that mandatory TOU pricing has addressed the Ontario Energy Board's goal of electricity conservation.

Analysis of electricity demand shifting is more complex. Examining Table 4.2, we can see that the majority of estimated summer demand reduction occurs in on-peak and mid-peak periods. When TOU effects are aggregated by price period, change during off-peak periods for both working days and non-working days is only -0.6% to -0.9% and is not statistically significant. The observed mean of electricity demand during those periods falls within the aggregated confidence interval of the counterfactual estimate. We conclude that conservation has been focused during on- and mid-peak periods, but demand from these periods has not been shifted to off-peak periods.

However, by examining hourly reduction in Figure 4.1 more closely, the off-peak hours 19:00 through 21:00 on working days are notable. In these three hours the the confidence interval for the estimated demand reduction does not contain the observed electricity demand. Together with the second mid-peak period, 17:00 through 18:00, there appears to be substantial demand reduction occurring in 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. Furthermore, [1] concludes that TOU periods may not be optimally aligned with provincial demand patterns, though that study uses residential, commercial, and industrial aggregate provencial data.

Extrapolating the results from Table 4.2 to all 20,556 residential customers in the local distribution company's service region, the hourly impact of TOU pricing is shown in 4.3.

Table 4.3: Change in aggregate electricity demand for each price period under TOU pricing extrapolated for the local distribution company's 20,556 residential customers.

	Hourly LDC Impact (MWh)	95% Confidence Interval (MWh)
Summer On-Peak	-0.72	± 0.49
Summer Mid-Peak	-0.62	± 0.49
Summer Off-Peak	-0.23	± 0.49
Summer Non-Working Day	-0.17	± 0.62

Chapter 5

Conclusions and Future Work

In this work we model changes in aggregate electricity demand associated with time, the price of electricity, and weather. We implement a detailed comparison of temperature variable transformations that incorporate coincident weather, past temperature, and non-linear transformations. Models are evaluated according to their explanatory power and interpretability so that our results may inform future modelling decisions.

Motivated by environmental concerns, we carry out a case study to evaluate the effectiveness of TOU electricity pricing in south west Ontario. Using results from our temperature transformation comparison, we select a well-performing transformation to model the predominant effects of weather on electricity demand. Many additional explanatory variable interactions related to TOU pricing are added to the model so that change in electricity demand associated with TOU pricing can be modelled with hourly fidelity.

5.1 Conclusions

- A six-hour moving average of dry-bulb temperature observations has significant explanatory power and is easily interpretable. Its non-linear relationship with electricity demand is best modelled using a natural cubic splines transformation.

- Temperature transformation using exposure-lag-response association generated mixed results. It has slightly greater explanatory power than other variable transformations, but its intended purpose, to model the weight of an exposure effect over time, is not easily interpretable. We conclude that this negative impact on model interpretability does not justify the moderate increase in variance explained.

- During the summer rate season, TOU pricing is associated with electricity conservation across all price periods. Additionally, the peak-to-average ratio of electricity demand changed -0.8% from 1.441 to 1.429.
- The average demand change during on-peak and mid-peak periods is -2.6% and -2.4% respectively. Change during working day and non-working day off-peak periods is -0.9% and -0.6% but is not statistically significant. We interpret this to mean that electricity demand is not being shifted to off-peak periods, but that conservation is focused during on- and mid-peak periods.
- Our TOU findings are consistent with prior TOU evaluations carried out within the province, though less pronounced than pilot studies. Our study is limited by the time period of data sampled. The estimated impact of TOU prices is based on data from the first summer after TOU pricing was implemented by the local distribution company. It is possible that customers had not been subject to TOU prices long enough to change their electricity demand habits. Conversely, it is also possible that because TOU pricing was still new, their response to the policy may be overstated.

5.2 Future Work

Beyond this analysis of TOU pricing's effects on residential customers in aggregate, we intend to study the effects of TOU pricing on individual customers or groups of customers.

5.2.1 Neighbourhood Comparison

Each meter in our data set has longitude and latitude information. Using this precise location data, a spatial component could be incorporated into the analysis. We consider the possibility of grouping customers by neighbourhood using postal code information or Statistics Canada's dissemination area coding (block-level). Does TOU pricing affect neighbourhoods differently? Can this be paired with household income bands available in Statistics Canada data? Does TOU pricing affect customers differently based on affluence?

5.2.2 Differences by Discretionary Electricity Usage

Customers may also be stratified according to their annual electricity consumption. A number of groups could be created based on the magnitude of customer's electricity use.

Do those who have greater discretionary electricity use have a different reaction to TOU prices than those with less?

5.2.3 Effects of Time-of-Use on Customers Grouped by Demand Pattern

Customers may also be grouped using some clustering technique to analyze their electricity demand time series. Grouping customers by demand patterns (i.e. demand shape) might yield some easily interpretable groups. For example, one group might be customers who are active within the household all day such as stay-at-home parents. A different group may be customers with demand peaks before and after the work-day, such as single working professionals. A random effects model could then be used to quantify how TOU pricing affects each group differently. Are some customer group demand patterns more flexible than others? How has each type of customer shifted their electricity demand? Are some groups more responsible for demand shifting than others? Do the effects of TOU pricing on each group have any future policy implications?

5.2.4 Irish Residential Data Sample

The Commission for Energy Regulation in Ireland has made a sample of 30-minute household electricity demand available for study through the Irish Social Science Data Archive. There are 5,000 labelled residential and commercial customers observed in an experimental setting from July 2009 through December 2010 [10]. Using this data we would like to both analyze the effects of temperature on residential electricity demand in Ireland and also confirm the ability of our temperature transformation comparison to describe an entirely new data set. Unfortunately, there no location information tied to each meter in the sample and participants may be located anywhere in the country. The primary concern for modelling the effects of temperature on this data is whether a weighted average temperature for the country is sufficient. Initial investigation of weather time series from four cities Cork, Connaught, Dublin, and Shannon suggests that temperature across the four locations is fairly homogeneous. A weighted average temperature for the country may be adequate for our proposed study.

APPENDICES

Appendix A

Review of Multiple Linear Regression

We wish to understand the relationship between a response variable, represented by the vector $\mathbf{Y}$, and three matrices of explanatory variables $\mathbf{X}$, $\mathbf{V}$, and $\mathbf{T}$. The explanatory variables in $\mathbf{X}$, $\mathbf{V}$, and $\mathbf{T}$ have some functional relationship with the response variable $\mathbf{Y}$, shown by (A.1).

$$\mathbf{Y} = f(\mathbf{X}, \mathbf{V}, \mathbf{T}) + \epsilon \tag{A.1}$$

The function f represents the systematic information that $\mathbf{X}$, $\mathbf{V}$, and $\mathbf{T}$ provide about $\mathbf{Y}$. The irreducible error term ϵ represents unmeasured explanatory variables or error introduced when assuming a form for f. Ideally, ϵ should be independent of $\mathbf{X}$, normally distributed, and have mean zero.

Our study assumes that the underlying relationship has a linear form, shown in (A.2).

$$f(\mathbf{X}, \mathbf{V}, \mathbf{T}) = \beta_0 + \mathbf{X}\beta + \mathbf{V}\omega + \mathbf{T}\theta \tag{A.2}$$

Let β_0 be the intercept term. All explanatory variables' effects are offset from this value. $\mathbf{X}$ is an $N \times P_a$ matrix of a type of explanatory variables. β is a $P_a \times 1$ vector of coefficients reflecting the additive effect of each variable in $\mathbf{X}$. $\mathbf{V}$ is an $N \times P_b$ matrix of a second type of explanatory variables. β is a $P_b \times 1$ vector of coefficients reflecting the additive effect of each variable in $\mathbf{V}$. $\mathbf{T}$ is an $N \times P_c$ matrix of a third type of explanatory variables. θ is a $P_c \times 1$ vector of coefficients reflecting the additive effect of each variable in $\mathbf{T}$. The reason for splitting explanatory variables into three matrices and three vectors of coefficients is simply for clarity in our study.

Substituting the linear functional form (A.2) into (A.1) yields a parametric, multiple linear regression model, shown in (A.3). This is also simply referred to as *multiple regression* and is referred to as such throughout our study.

$$\mathbf{Y} = \beta_0 + \mathbf{X}\beta + \mathbf{V}\hat{\omega} + \mathbf{T}\theta + \epsilon \tag{A.3}$$

Note that this equation is shown in matrix notation for clarity and brevity. The expanded form of the ith row is shown in (A.4)

$$\mathbf{y}_i = \beta_0 + \mathbf{x}_{i,1}\beta_1 + ... + \mathbf{x}_{i,P_a}\beta_{P_a} + \omega_1\mathbf{v}_{i,1} + ... + \omega_{P_b}\mathbf{v}_{i,P_b} + \theta_1\mathbf{t}_{i,1} + ... + \theta_{P_c}\mathbf{t}_{i,P_c} + \epsilon_i \tag{A.4}$$

Let i index each hour in a sample of time series data and $\mathbf{y}_i$ be the ith observation of the response variable. Each explanatory variable $\mathbf{x}_{i,p_a}$ and its coefficient β_{p_a} models part of a systematic, linear relationship with $\mathbf{y}_i$. We define the notation $\mathbf{x}_{\bullet,j}$ to represent the jth column and all rows of $\mathbf{x}$. The jth coefficient β_j represents the additional contribution of variable $\mathbf{x}_{\bullet,j}$ to $\mathbf{y}$ [26]. As an example of this additive relationship, consider one observation of electricity demand $\mathbf{y}_i$ where the hour-of-day at index i is 19:00. At 19:00, residential customers are typically home from work and active within their house, resulting in increased electricity demand. $\mathbf{y}_i$ will deviate positively from the sample mean β_0. If $\mathbf{x}_{i,j}$ is an indicator variable for hour 19:00, then β_j will be positive, reflecting the systematic relationship of 19:00 with electricity demand.

In practice, the true underlying function f driving the response variable is often unknown. We can only create an estimate for f using the samples of observations $\mathbf{X}$, $\mathbf{V}$, $\mathbf{T}$, and $\mathbf{Y}$. This estimate is shown in (A.5)

$$\begin{aligned} \hat{\mathbf{Y}} &= \hat{f}(\mathbf{X}, \mathbf{V}, \mathbf{T}) \\ \hat{\mathbf{Y}} &= \hat{\beta}_0 + \mathbf{X}\hat{\beta} + \mathbf{V}\hat{\omega} + \mathbf{T}\hat{\theta} \end{aligned} \tag{A.5}$$

The estimated underlying function $\hat{f}$ produces $\hat{\mathbf{Y}}$, the estimate for $\mathbf{Y}$. $\hat{f}$ is comprised of coefficient estimates represented by $\hat{\beta}_0$, the $N \times P_a$ vector $\hat{\beta}$, the $N \times P_b$ vector $\hat{\omega}$, and the $N \times P_c$ vector $\hat{\theta}$. Note that the estimated model has no error term because an estimate cannot account for irreducible error. Instead, we can only measure the residual differences between values of $\mathbf{Y}$ and $\hat{\mathbf{Y}}$, shown in (2.2). We use these residuals to quantify the best estimates of $\hat{\beta}$, $\hat{\omega}$, and $\hat{\theta}$ using ordinary least squares.

Appendix B

Box Plot Interpretation

B.1 Description and Diagram

When plotting many observations which can be grouped into categorical factors, scatter plots will likely have many overlapping points, reducing the amount of information they can convey. One option may be to use opacity or stutter points along one axis to give a sense of density. Another, more common solution is to make use of box plots.

For categorical groups of data, box plots are a useful tool for conveying additional information about the distribution of data in each category. Figure B.1 is a diagram of a box and whisker plot as it relates to the probability density function of a normal distribution. The bold line in the middle of the box is the median of the data for a given category. The top and bottom of the box indicate the 25th and 75th percentiles of the data, known as the 1st and 3rd quaretiles. The range between the 1st and 3rd quaretiles is known as the inter-quartile range and represents 50% of the data for a given category. The thin lines extending out of each end of the box are *whiskers*. When there are no outliers, the whisker extends out the maximum or minimum value. When there are outliers, the whisker extends only to 1.5 times the interquartile range, roughly two standard deviations. Any data points outside this range are considered outliers for each category [11].

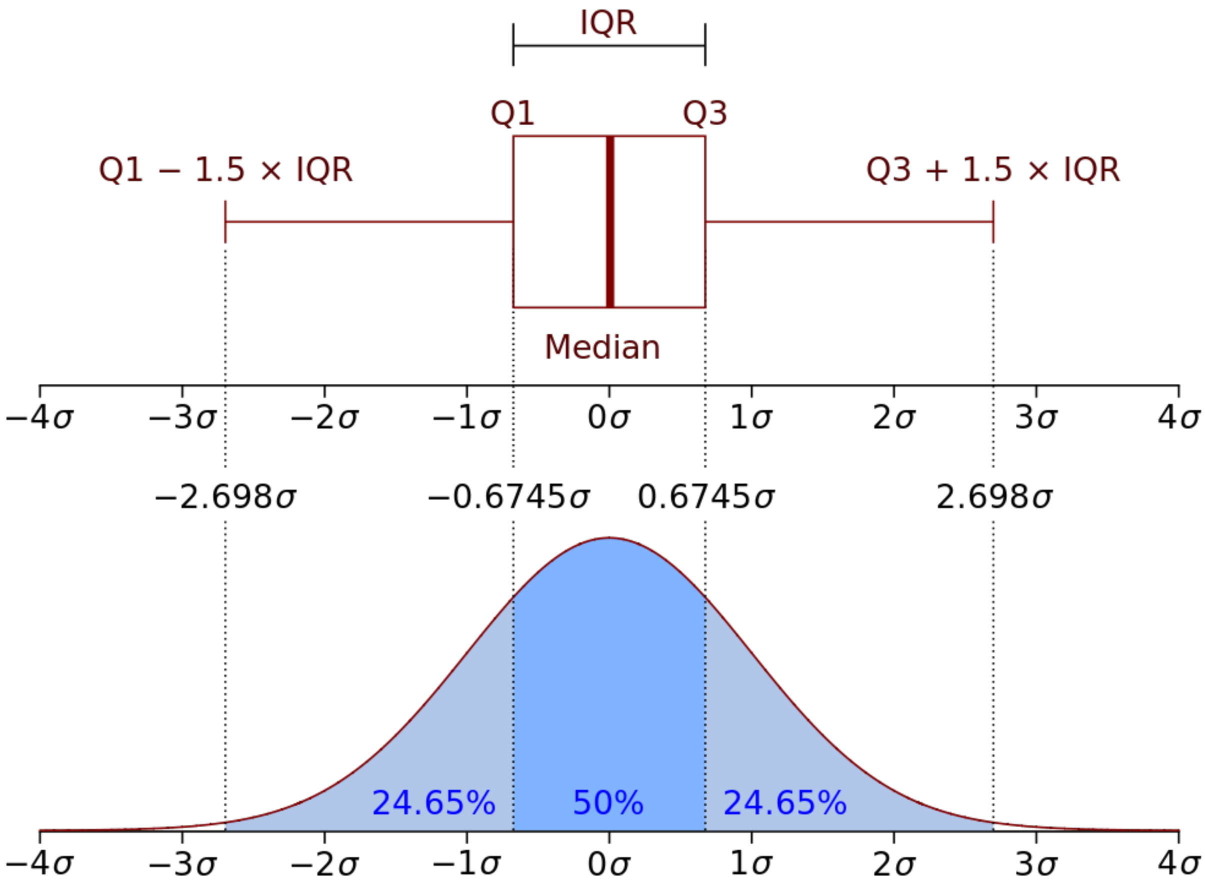

Figure B.1: Diagram of a box plot and its relationship to the probability density function of a normal distribution. (Creative-Commons 2.5 image [34], edited by Reid Miller)

Appendix C

Temperature Transformation Comparison Using Log Transformed Aggregate Electricity Demand

C.1 Formulation

Fitting a model using log transformed response variables is often used to address heteroscedasticity in residuals. We repeat the set of comparisons described in section 3.3 using a log transformed response, shown in (C.1).

$$\mathbf{y}'_i = ln(\mathbf{y}_i), \quad i = 1, ..., N \tag{C.1}$$

The estimate for $\mathbf{Y}'$ is shown in (C.2).

$$\hat{\mathbf{Y}}' = \hat{\beta}_0 + \mathbf{X}\hat{\beta} + \mathbf{V}\hat{\omega} + \mathbf{T}\hat{\theta} \tag{C.2}$$

A log transformation results in a greater amount of shrinkage of the larger response values, reducing heteroscedasticity [31]. When dealing with the log transformed model, all reported statistics and plotted data will be backtransformed for interpretation.

C.2 Results

Table C.1 shows the results of fitting (C.1) as described in section C.1. The results compare each temperature transformation's explanatory power when estimating a log transformed response. All summary statistics have been backtransformed for comparison to the matching temperature transformation used to estimate an untransformed response. Note that BIC values between Tables 3.6 and C.1 are not comparable due to the different response variables.

Table C.1: Results of temperature transformation comparison estimating log transformed response as described in section C.1.

Temporal Transform	**Weather Transform**	**Non-Linearity Transform**	**Adj. R^2**	**BIC**	**DW**	**Avg. MAE (kWh)**	**Avg. MAPE (%)**
			ideal=1	ideal=low	ideal=2	ideal=0	ideal=0, max=100
Null Model (i.e. intercept only)			-0.030	13412.8	0.068	0.419	33.58
Non-Temperature Explanatory Variables Only			0.434	3110.2	0.044	0.364	32.17
None (i-0)	None (Drybulb)	None (Linear)	0.614	-645.8	0.061	0.249	20.83
None (i-0)	None (Drybulb)	Switching Regression	0.864	-14244.5	0.189	0.152	13.00
None (i-0)	None (Drybulb)	Natural Splines	0.867	-14495.8	0.185	0.149	12.76
None (i-0)	Feels Like	Switching Regression	0.849	-14153.2	0.192	0.158	13.36
None (i-0)	Feels Like	Natural Splines	0.868	-14698.5	0.195	0.150	12.83
i-1	None (Drybulb)	Switching Regression	0.872	-15444.0	0.201	0.140	12.11
i-1	None (Drybulb)	Natural Splines	0.877	-15782.0	0.199	0.137	11.88
i-1	Feels Like	Switching Regression	0.854	-15265.1	0.205	0.147	12.51
i-1	Feels Like	Natural Splines	0.879	-16012.1	0.210	0.138	11.88
i-2	None (Drybulb)	Switching Regression	0.869	-15832.0	0.218	0.138	11.89
i-2	None (Drybulb)	Natural Splines	0.877	-16220.8	0.217	0.135	11.62
i-2	Feels Like	Switching Regression	0.849	-15589.8	0.222	0.144	12.27
i-2	Feels Like	Natural Splines	0.880	-16501.8	0.228	0.135	11.60
i-3	None (Drybulb)	Switching Regression	0.860	-15482.3	0.223	0.143	12.00
i-3	None (Drybulb)	Natural Splines	0.869	-15890.6	0.222	0.139	11.70
i-3	Feels Like	Switching Regression	0.839	-15225.8	0.224	0.147	12.32
i-3	Feels Like	Natural Splines	0.874	-16215.4	0.233	0.138	11.64
i-4	None (Drybulb)	Switching Regression	0.846	-14594.0	0.220	0.151	12.36
i-4	None (Drybulb)	Natural Splines	0.855	-15004.0	0.217	0.147	12.00
i-4	Feels Like	Switching Regression	0.824	-14355.8	0.222	0.153	12.62
i-4	Feels Like	Natural Splines	0.861	-15350.6	0.227	0.145	11.94
i-5	None (Drybulb)	Switching Regression	0.830	-13439.1	0.198	0.159	12.76
i-5	None (Drybulb)	Natural Splines	0.839	-13814.2	0.196	0.155	12.37
i-5	Feels Like	Switching Regression	0.810	-13248.3	0.201	0.161	13.02
i-5	Feels Like	Natural Splines	0.844	-14160.6	0.205	0.154	12.30
i-6	None (Drybulb)	Switching Regression	0.812	-12083.5	0.184	0.168	13.25
i-6	None (Drybulb)	Natural Splines	0.819	-12406.2	0.182	0.164	12.80
i-6	Feels Like	Switching Regression	0.794	-11952.1	0.186	0.170	13.53
i-6	Feels Like	Natural Splines	0.825	-12729.9	0.190	0.162	12.75
CDH/HDH (L=6)	None (Drybulb)	Switching Regression	0.883	-17514.8	0.181	0.129	10.90
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	0.862	-17222.4	0.178	0.135	11.33
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	0.884	-17654.8	0.190	0.129	11.05
Moving Avg. (L=6)	None (Drybulb)	Natural Splines	0.893	-18131.7	0.191	0.126	10.83
Moving Avg. (L=6)	Feels Like	Switching Regression	0.861	-17359.6	0.184	0.135	11.44
Moving Avg. (L=6)	Feels Like	Natural Splines	0.896	-18476.8	0.195	0.125	10.76
Lag-Response: Cubic Polynomial (L=6)	None (Drybulb)	Exposure-Response: Switching Regression	0.876	-17771.8	0.200	0.138	11.89
Lag-Response: Cubic Polynomial (L=6)	None (Drybulb)	Exposure-Response: Natural Splines	0.906	-19754.9	0.219	0.120	10.11
Lag-Response: Cubic Polynomial (L=6)	Feels Like	Exposure-Response: Switching Regression	0.841	-17538.5	0.197	0.143	12.12
Lag-Response: Cubic Polynomial (L=6)	Feels Like	Exposure-Response: Natural Splines	0.906	-19790.8	0.220	0.121	10.14

References

[1] Adedamola Adepetu, Elnaz Rezaei, Daniel Lizotte, and Srinivasan Keshav. Critiquing time-of-use pricing in Ontario. In *2013 IEEE International Conference on Smart Grid Communications, SmartGridComm 2013*, pages 223–228, 2013.

[2] Dennis Aigner. The Residential Electricity Time-of-Use Pricing Experiments: What Have We Learned? In David A. Wise and Jerry A. Hausman, editors, *Social Experimentation*, chapter 1, pages 11 – 54. University of Chicago Press, Chicago, 1985.

[3] Shirley Almon. The distributed lag between capital appropriations and expenditures. *Econometrica: Journal of the Econometric Society*, 33(1):178–196, 1965.

[4] Omid Ardakanian, Negar Koochakzadeh, Rayman Preet Singh, Lukasz Golab, and S. Keshav. Computing Electricity Consumption Profiles from Household Smart Meter Data. In K. Selçuk Candan, Sihem Amer-Yahia, Nicole Schweikardt, Vassilis Christophides, and Vincent Leroy, editors, *Energy Data Management*, pages 140–147, Athens, 2014.

[5] Alok Bhargava, Luisa Franzini, and Wiji Narendranathan. Serial Fixed Correlation and Effects the Model. *Review of Economic Studies*, 49(4):533–549, 1982.

[6] Benjamin J. Birt, Guy R. Newsham, Ian Beausoleil-Morrison, Marianne M. Armstrong, Neil Saldanha, and Ian H. Rowlands. Disaggregating categories of electrical energy end-use from whole-house hourly data. *Energy and Buildings*, 50:93–102, July 2012.

[7] Patrick Breheny and Woodrow Burchett. *visreg: Visualization of Regression Models*, 2015. R package version 2.1-1.

[8] Alexander Bruhns, Gilles Deurveilher, and Jean-Sébastien Roy. A non-linear regression model for mid-term load forecasting and improvements in seasonality. In *Power*

Systems Computation Conference, number August, pages 1–8, Liege, 2005. Curran Associates, Inc.

[9] José Ramón Cancelo, Antoni Espasa, and Rosmarie Grafe. Forecasting the electricity load from one day to one week ahead for the Spanish system operator. *International Journal of Forecasting*, 24(4):588–602, October 2008.

[10] Commission for Energy Regulation. CER Smart Metering Project, 2010.

[11] Michael J Crawley. *The R Book*. John Wiley & Sons, Ltd., Chichester, 2nd edition, 2013.

[12] Robert F. Engle, Clive William John Granger, John Rice, and Andrew Weiss. Semiparametric Estimates of the Relation Between Weather and Electricity. *Journal of the American Statistical Association*, 81(394):310–320, 1986.

[13] Environment Canada. Environment Canada — Climate: Hourly Data Report. http://climate.weather.gc.ca, February 2015.

[14] Environment Canada. Environment Canada — Wind Chill Calculator. http://www.ec.gc.ca/meteo-weather/default.asp?lang=En&n=0F42F92D-1 http://www.ec.gc.ca/meteo-weather/default.js, April 2015.

[15] Julian J. Faraway. *Practical Regression and Anova using R*. 2002.

[16] Ahmad Faruqui and Sanem Sergici. Household response to dynamic pricing of electricity: a survey of 15 experiments. *Journal of Regulatory Economics*, 38(2):193–225, August 2010.

[17] Ahmad Faruqui, Sanem Sergici, and Lamine Akaba. Dynamic pricing of electricity for residential customers: the evidence from Michigan. *Energy Efficiency*, 6(3):571–584, February 2013.

[18] Ahmad Faruqui, Sanem Sergici, Neil Lessem, Dean Mountain, Frank Denton, Byron Spencer, and Chris King. Impact Evaluation of Ontarios Time-of- Use Rates: First Year Analysis. Technical report, Ontario Power Authority, Toronto, 2013.

[19] John Fox and Sanford Weisberg. *An R Companion to Applied Regression*. Sage, Thousand Oaks CA, second edition, 2011.

[20] Luiz Friedrich, Peter Armstrong, and Afshin Afshari. Mid-term forecasting of urban electricity load to isolate air-conditioning impact. *Energy and Buildings*, 80:72–80, 2014.

[21] Antonio Gasparrini. Distributed lag linear and non-linear models in R: the package dlnm. *Journal of Statistical Software*, 43(8):1–20, 2011.

[22] Antonio Gasparrini, Ben Armstrong, and Mike G. Kenward. Distributed lag non-linear models. *Statistics in Medicine*, 29(21):2224–2234, 2010.

[23] Google. Google Earth. http://www.google.com/earth/, March 2015.

[24] Andrew Harvey and Siem Jan Koopman. Forecasting Hourly Electricity Demand Using Time-Varying Splines. *Journal of the American Statistical Association*, 88(424):1228–1236, 1993.

[25] Thomas Haslwanter. Residuals for Linear Regression Fit, 2013.

[26] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* Springer, 2nd edition, 2005.

[27] Hydro One. Hydro One Networks Inc. Time-of-Use Pricing Pilot Project Results. Technical report, Hydro One, Toronto, 2008.

[28] Rob J. Hyndman and George Athanasopoulos. *Forecasting: principles and practice.* OTexts, October 2013.

[29] Rob J. Hyndman and Shu Fan. Density forecasting for long-term peak electricity demand. *IEEE Transactions on Power Systems*, 25(2):1142–1153, 2010.

[30] IEEE Standards Association. IEEE Guide for Loading Mineral- Oil-Immersed Transformers and Step-Voltage Regulators, 2012.

[31] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An Introduction to Statistical Learning with Applications in R.* Springer, New York, 4th edition, 2013.

[32] Katrina Jessoe, David Rapson, and Jeremy Blair Smith. The Effect of a Mandatory Time-of-Use Pricing Reform on Residential Electricity Use. In *American Economic Association Annual Meeting*, pages 1–54, San Diego, 2013. American Economic Association.

[33] Loren Kirkeide. Effects of Three-Hour On-Peak Time-of-Use Plan on Residential Demand during Hot Phoenix Summers. *Electricity Journal*, 25(4):48–62, 2012.

[34] Chen-Pan Liao. Boxplot vs PDF, 2011.

[35] Dale P. Lifson and Allen K. Miedema. A comparative analysis of time-of-use electricity rate effects: the Arizona experiment. *Energy*, 6(5):403–408, May 1981.

[36] Simone Maggiore, Massimo Gallanti, Walter Grattieri, and Michele Benini. 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)*, volume 5, pages 1–4, Stockholm, 2013. Institution of Engineering and Technology.

[37] Li Mei and Wan Qiulan. Study on TOU Price Implantation Effect on Huainan Resident. In *2011 International Conference on Computer Distributed Control and Intelligent Environmental Monitoring*, number May 2007, pages 2221–2224. IEEE, February 2011.

[38] Reid Miller. Github Repository (r24mille): Environment Canada Weather History. https://github.com/r24mille/env_canada_weather_history, November 2014.

[39] Reid Miller. Github Repository (r24mille): Local Distribution Company Regression. https://github.com/r24mille/ldc_regression, June 2015.

[40] Frits Møller Andersen, Helge V. Larsen, and Trine Krogh Boomsma. Long-term forecasting of hourly electricity load: Identification of consumption profiles and segmentation of customers. *Energy Conversion and Management*, 68:244–252, 2013.

[41] Julián Moral-Carcedo and José Vicéns-Otero. Modelling the non-linear response of Spanish electricity demand to temperature variations. *Energy Economics*, 27(3):477–494, May 2005.

[42] Dean C. Mountain and Evelyn L. Lawsom. A Disaggregated Nonhomothetic Modeling of Responsiveness to Residential Time-of-Use Electricity Rates. *International Economic Review*, 33(1):181–207, 1992.

[43] Vito M.R. Muggeo. Estimating regression models with unknown break-points. *Statistics in Medicine*, 22:3055–3071, 2003.

[44] Vito M.R. Muggeo. segmented: an R Package to Fit Regression Models with Broken-Line Relationships. *R News*, 8(1):20–25, 2008.

[45] Natural Resources Canada. Residential Sector Ontario Table 4: Space Cooling Secondary Energy Use and GHG Emissions by Cooling System Type — Office of Energy Efficiency. http://oee.nrcan.gc.ca/corporate/statistics/neud/dpa/showTable.cfm?type=CP&sector=res&juris=on&rn=4&page=0, February 2012.

[46] Natural Resources Canada. Residential Sector Ontario Table 5: Space Heating Secondary Energy Use and GHG Emissions by Energy Source — Office of Energy Efficiency. http://oee.nrcan.gc.ca/corporate/statistics/neud/dpa/showTable.cfm?type=CP&sector=res&juris=on&rn=5&page=0, February 2012.

[47] Navigant Research and Newmarket Tay Power Distribution. The Effects of Time-of-Use Rates on Residential Electricity Consumption. Technical report, Ontario Energy Board, Toronto, 2010.

[48] Navigant Research and Ontario Energy Board. Time of Use Rates in Ontario, Part 1: Impact Analysis. Technical report, Ontario Energy Board, Toronto, 2013.

[49] Guy R. Newsham, Benjamin J. Birt, and Ian H. Rowlands. A comparison of four methods to evaluate the effect of a utility residential air-conditioner load control program on peak electricity use. *Energy Policy*, 39(10):6376–6389, October 2011.

[50] Guy R. Newsham and Brent G. Bowker. The effect of utility time-varying pricing and load control strategies on residential summer peak electricity use: A review. *Energy Policy*, 38(7):3289–3296, July 2010.

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

[52] Ontario Energy Board. OEB announces electricity price plan for residential, low-volume and designated consumers. http://www.ontarioenergyboard.ca/documents/communications/pressreleases/2005/press_release_110305.pdf, March 2005.

[53] Ontario Energy Board. Smart Meter Deployment and TOU Pricing Monitoring Report: August 2010. http://www.ontarioenergyboard.ca/oeb/_Documents/SMdeployment/SM_Monitoring_Report_August2010.pdf, 2010.

[54] Ontario Energy Board. Smart Meter Deployment and TOU Pricing Monitoring Report: August 2012. http://www.ontarioenergyboard.ca/oeb/_Documents/SMdeployment/Monthly_Monitoring_Report_August2012.pdf, 2012.

[55] Ontario Energy Board, IBM, and eMeter Strategic Consulting. Ontario Energy Board Smart Price Pilot Final Report. Technical Report July, Ontario Energy Board, Toronto, 2007.

[56] Konstantinos Papakostas and Nicolaos Kyriakis. Heating and cooling degree-hours for Athens and Thessaloniki, Greece. *Renewable Energy*, 30(12):1873–1880, 2005.

[57] Angel Pardo, Vicente Meneu, and Enric Valor. Temperature and seasonality influences on Spanish electricity load. *Energy Economics*, 24:55–70, 2002.

[58] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2013.

[59] Fred Ramsey and Daniel Schafer. *The Statistical Sleuth: A Course in Methods of Data Analysis*. Cengage Learning, 3rd edition, 2012.

[60] Lans P. Rothfusz. The Heat Index Equation, 1990.

[61] Galit Shmueli. To Explain or to Predict? *Statistical Science*, 25(3):289–310, August 2010.

[62] Kenneth Train and Gil Mehrez. Optional time-of-use prices for electricity: econometric analysis of surplus and pareto impacts. *The RAND Journal of Economics*, 25(2):263–283, 1994.

[63] Achim Zeileis. Econometric computing with hc and hac covariance matrix estimators. *Journal of Statistical Software*, 11(10):1–17, 11 2004.