Forecasting Residential Air Conditioning Loads

Shira Horowitz^{a,1,*}, Brandon Mauch^a, Fallaw Sowell^b

^aDepartment of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, PA 15213

^bTepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213

Abstract

A doubly censored Tobit model is used to forecast hourly air-conditioner usage for individual households. The model worked well over a wide range of temperatures, 9 to 38 degrees Celsius, making it possible to accurately forecast the electricity load for a variety of demand response applications including operational reserves for renewable energy integration. Individual models are simulated and summed to obtain aggregate forecasts and confidence intervals. The model allows for correlation between the individual shocks that occur in a region. This approach gives substantially more accurate results than the moving average method typically used for forecasting and measuring direct load control. Applying the model to data from three U.S. utilities produced mean square error values from 0.027 to 0.041 with average load values per customer ranging from 0.49 to 0.62 kW.

Keywords: direct load control, demand response, energy markets, electric load forecasting, tobit model.

1. Introduction

Demand response is the reduction of electric loads in response to market signals or system conditions. Electric grid operators and electrical utilities implement demand response to maintain grid reliability or provide electrical service at lower cost. One type of demand response is direct load control (DLC) where electrical appliances are remotely powered off. Grid operators must balance generation and load for reliability. From this perspective, load reduction (decrease in demand) is similar to generation increase (increase in supply). Unlike generation where the supply is deterministic (barring events that lead to a forced outage), the DLC resource is uncertain and must be forecasted. While generators are paid according to the quantity of energy supplied, DLC participants are paid based on the amount of load reduction². Load reductions cannot be directly measured; they are calculated by subtracting actual load during a DLC event from a

*Corresponding author

Email addresses: shira@alumni.cmu.edu (Shira Horowitz), bmauch@andrew.cmu.edu (Brandon Mauch), fs0v@andrew.cmu.edu (Fallaw Sowell)

¹Present address: PJM Interconnection, 955 Jefferson Avenue, Norristown, PA 19403, +1 610 666 2337.

²Load aggregators and large customers usually receive the market prices. However, residential customers usually receive a flat rate for participation from a load aggregator.

Preprint submitted to Applied Energy

June 14, 2014

customer's estimated load conditional on the DLC event not occurring. In this paper we propose a new method for forecasting and measuring DLC of residential ACs using a Tobit model with upper and lower censoring.

1.1. Direct Load Control

Effective DLC is widely used to reduce peak load, which delays the need to build power plants or transmission lines. In recent years it is also used as reserve capacity for contingencies in the grid. PJM, a northeastern regional transmission organization (RTO) in the US, provides 20% of its contingency reserves with DLC resources [1]. Supplying operating reserves with DLC lowers overall costs and allows increased active power output from low cost generators no longer needed to supply reserves [2]. DLC can also be used to adjust load as a means of balancing variability of wind and solar resources [3, 4, 5, 6]. The Department of Energy stated that increased reliance on electricity generation from wind and solar power is one factor that will drive demand response programs [7]. A German study analyzed the potential of DLC from different industries to provide operating reserves for wind power. It concluded that full use of the potential industrial DLC resources would result in operating cost savings of nearly 0.5 billion Euros [8].

Direct load control programs are tailored for different types of customers. Commercial and industrial customers agree to interrupt or curtail specific loads when called upon by the utility in exchange for reduced electric rates. Residential customers often agree to allow utilities limited control of thermal loads such as air conditioners and water heaters. Studies have shown significant economic benefits from incorporating interruptible load into dispatch decisions [9].

For residential customers, utilities and RTOs generally control air conditioner loads by cycling the power to the compressor. Power to the compressor unit is turned off for brief periods in order to temporarily reduce demand. Alternatively, thermal energy storage systems can be used to reduce AC use during peak hours by cooling a storage medium in the evenings and using the cold material to cool the air in a building [10].

Increased use of DLC will require more accurate load forecasting techniques that are easy to implement, like the method we develop in this work. Model accuracy is needed over a range of temperatures since DLC can be called for peak load reductions at high temperatures as well as contingency reserves at lower temperatures.

ACs are well suited for DLC since they can be powered off for short periods of time without much customer discomfort. A California utility surveyed customers during a pilot study and found the majority did not notice DLC events lasting 15 minutes or less [11]. ACs also comprise a large portion of residential loads (roughly 20% of residential electricity consumption) [12].

Advanced electric meters (i.e. smart meters) allow finer control over electric loads and provide more load data which will enable greater use of DLC in electric grids [13, 14, 15]. Two-way communication between the utilities and customers may also allow customers to offer load curtailments for specified prices [16]. As of 2011, 13.4% of all electricity customers had advanced meters [17], but the Department of Energy is providing funds to quickly increase this level [18]. There is a need for models to efficiently take advantage of this new data. This Tobit model captures the realistic situation that AC loads are bounded below by zero and above by the maximum energy consumption for a particular AC.

Recent changes in wholesale electric markets will also increase the use of DLC. In 2011, the Federal Energy Regulatory Commission (FERC) issued order number 745 which directs wholesale energy market operators to compensate demand side resources (eg. DLC) the full energy market price as long as dispatching the DR resource is cost-effective [19]. Each market operator sets a threshold price based on historical data which is used as the minimum price at which DR resources are compensated.

1.2. Load Forecasting

Accurate load forecasts are essential for efficient DLC. DR resources are paid the energy market locational marginal price for the reduced load based upon the customer baseline (CBL). This is an estimate for a counterfactual event, i.e. the expected load conditional on the DLC event not happening. Inaccuracies in the CBL lead to incorrect and unfair payments. Underpayments for DR resources discourage participation while overpayments lead to excessive charges on load serving entities who pay for the reductions. System planners need to accurately know how much load reduction to expect during a DLC event. Reducing uncertainty in the load forecasts will become more important as DLC resources provide more ancillary services to help balance the smart electricity grid.

Default CBLs differ across RTOs and independent system operators (ISOs) [20, 21]. Most are simple moving averages. In the PJM RTO, the default CBL is the average hourly load profile from the 4 highest load days of the previous 5 similar day types (weekdays, Saturdays, Sundays/holidays) [22]. The California ISO, the New York ISO and the New England's ISO calculate CBLs by averaging loads from the previous 10 similar days [23, 24, 25]. The Electric Reliability Council of Texas publishes 3 different default CBL calculations: a linear regression of energy consumption on covariates representing weather conditions, daylight hours, season and day of the week; a moving average of 8 of the previous 10 similar days; or a model that averages days with load profiles similar to the event day [26]. All ISO/RTOs accept alternative methods for CBL determination as long as it is approved. This paper presents such an alternative.

Broadly speaking, air conditioner load forecast models can be classified into two distinct categories: engineering models and statistical models. Both model types attempt to forecast AC load as a function of several variables, primarily: temperature and time of day. The most common are engineering models of a house that consist of a system of differential equations that capture the evolution of indoor temperature and the on/off cycles of the air conditioner compressor given weather variables such as temperature, solar radiation, etc. These models requires measuring the thermal characteristics and thermostat settings of each house for use as parameters [27, 28, 29, 30]. Another approach is to use maximum likelihood to estimate these parameters from historical data [31, 32, 33]. The latter method still requires knowledge of the thermostat setpoints. Unfortunately, these models are sensitive to changes in the physical properties of the residence such as home improvements.

There is comparatively less work on statistical models applied to AC load forecasts, especially residential. Statistical models do not directly model the dynamics of energy flows. Instead they capture trends in historical AC load data to predict future loads.

Parametric models of AC duty cycles have been used to estimate load reductions by comparing controlled and non-controlled AC data [34]. Autoregressive models have been used in AC forecasts for non-residential buildings [35], but not the highly variable residential data. Machine learning type models have been proposed to forecast building

energy consumption using support vector regression [36] and artificial neural networks [37, 38]. These types of models capture the non-linearities in energy demand, but are data intensive for each household. A recent proposal to forecast load reduction from AC DLC relies on fitting a model to load measurements at a feeder circuit level [39]. This method cannot forecast load for individual households and requires a large fraction of ACs on each feeder participate in DLC so that it can distinguish the signal from the noise. This is a concern for forecasts at lower temperatures.

This paper considers a doubly censored Tobit model to forecast hourly individual AC loads. The model uses ambient temperature and time of day as covariates. The individual forecasts are aggregated via simulation to create day-ahead hourly aggregate load forecast and forecast intervals. A Tobit model accounts for the non-linearities inherent in AC energy consumption while not requiring extreme amounts of data. Other nonparametric methods such as neural networks require vast amounts of data for each customer. The model developed in this paper may be adopted by RTOs such as PJM under their current rules and is therefore immediately applicable, while more complicated methods may take years to adopt.

The remainder of this paper is organized as follows: section 2 describes the dataset. Section 3 describes the Tobit model and the theoretical framework of the model. The results are in section 4 and section 5 covers the policy implications.

2. Data

Under a confidentiality agreement, a dataset was obtained from Pepco Holdings, Inc. The dataset contains AC energy consumption data, weather data and individual characteristics for the 536 residential ACs from July - October 2010. Due to data quality issues 69 units were discarded with the analysis performed on the remaining 467 units (details on data quality and cleaning protocol are in appendix Appendix B). Data loggers were installed on the ACs during the month of July so the initial date of data collection varies. The data loggers recorded current measurements for the compressor circuits. During installation, technicians took spot measurements of voltage and power which were used to convert the current measurements to power measurements, and determine the maximum power of the AC. Since this data was only collected at installation, we assume this voltage and power factor are constant throughout the data collection period.

The raw data were instantaneous current values recorded at three minute intervals. A constant current level is assumed during each three minute period to estimate energy consumption at the hourly time scale, i.e. the compressor was assumed to be on or off for the entire 3 minute period. Adding the three minute periods, the hourly estimates take on 21 discrete values. This discreteness does not have a significant effect on the final results which are aggregated over all units.

The ACs were serviced by one of three utilities: Potomac Electric Power Company (PEPCO), Delmarva Power and Light (Delmarva) and Atlantic City Electric (ACE). Hourly temperature data were collected from weather stations located near each utility's territory and were assumed to be uniform throughout each region; see table D.1.

[Table 1 about here.]

There were three missing temperature measurements which were interpolated from the adjacent hours.

All the customers in the data sets had signed up to participate in DLC programs. Some customers had only recently enrolled and their DLC was not active. Data on when individual customer’s DLC became active is unavailable. For each AC unit, dummy variables are used to model the impact of DLC events.

In order to participate in the DLC program, customers agreed to have either switches capable of remote operation installed on the AC compressor circuit or smart thermostats that could be adjusted remotely. Customers received notice 24 hours prior to a DLC event. During the time period covered in the data, 8 DLC events occurred ranging in duration from one to four hours. Customers had the option of overriding the signal if they wanted. However, this only occurred with one customer in the dataset. Summary statistics for the dataset are in table D.2.

[Table 2 about here.]

3. Framework

3.1. Tobit Model Applied to AC Data

Preference for AC usage is positively related to temperature. At higher temperatures, consumers want more cooling, even if their AC has reached its maximum capacity, while at cooler temperatures, consumers want less AC, and if it is cool enough, they may want heat. Actual AC electricity load is constrained by zero and its maximum capacity. These structural restrictions on observed AC energy consumption are modeled with a doubly censored Tobit model [40].

A model is needed to capture the AC electricity load over a range of temperatures. Most current residential DLC programs are concerned with peak shaving, i.e. reducing demand only when temperature are extremely high. The need to balance the grid during contingencies or to accommodate alternative energy sources can occur at almost any temperature, and, hence, a model is needed to accurately forecast the AC electricity load (possible savings) as a function of outdoor temperature over a wide temperature range – not just extreme values.

The desired electricity load for an individual AC i at time t (incremented hourly) is modeled as a latent variable $y_{i,t}^*$:

$$y_{i,t}^* = \sum_{h=1}^{24} (\beta_{D_h,i} D_{h,t} + \beta_{TD_h,i} D_{h,t} T_t) + \beta_{T^2,i} T_t^2 + \beta_{T1,i} T_{t-1} + \beta_{E,i} E_t + \beta_{P,i} P_t + \epsilon_{i,t} \quad (1)$$

where T_t is $\max(0, R_t - 18.3)$, R_t is the outdoor temperature in degrees Celsius during hour t , $D_{h,t}$ is an indicator variable for hour h (1 during hour h , 0 otherwise), E_t , an indicator variable for a DLC event during hour t (1 if there is a DLC event that hour, 0 otherwise), P_t is an indicator for the three hours immediately after an event (1 for each of the three hours after an event, 0 otherwise) and $\beta_{\chi,i}$ is a parameter for AC i for covariate χ . The error $\epsilon_{i,t} \sim N(0, \sigma_i^2)$ accounts for an unobservable shock.

The shifted temperature term and temperature squared account for non-linearities of the observed temperature range. The terms in the parentheses are hourly intercepts and linear responses to temperature to account for consumers’ diurnal activity cycle. The lagged temperature term is included to account for the home’s thermal inertia. A lag of

1 was included because the partial-autocorrelation function was insignificant for longer lags. DLC events in the dataset are accounted for with E_t . A DLC event may increase a customer's preference for AC immediately after the event. An indicator variable, P_t is included to account for this. For example, if an AC is cycled below its normal duty cycle for 2 hours during an event the temperature inside the home may rise above the desired temperature. When the event is over, the AC will have to increase its output beyond what it would normally be to account for the temperature rise. This is known as the snapback effect. Additional variables were considered such as lagged AC load, eg. humidity, however these variables were either strongly collinear with other covariates or did not improve the model's fit.

Certain ACs had one or more hours of the day that had very few uncensored $y_{i,t}$ values (eg. all the values at 3 a.m. for a particular AC were 0). This lack of observed variability prevents both identification and convergence of the optimization routine. Any hour with 3 or fewer uncensored values was combined with the previous or next hour (whichever had fewer uncensored values) until each bin contained greater than 3 uncensored values.

Combine the covariates into $\mathbf{X}_t$ and all $\beta_{\chi,i}$ into β_i to write (1) as

$$y_{i,t}^* = \mathbf{X}_t' \beta_i + \epsilon_{i,t}. \quad (2)$$

The observed data censor $y_{i,t}^*$ between 0 and λ_i , the capacity of the AC, to obtain $y_{i,t}$, the actual energy consumption of AC i during hour t :

$$y_{i,t} = \begin{cases} 0 & y_{i,t}^* \leq 0 \\ y_{i,t}^* & 0 < y_{i,t}^* < \lambda_i \\ \lambda_i & \lambda_i \leq y_{i,t}^*. \end{cases} \quad (3)$$

This model is estimated for each AC by maximum likelihood (ML). 52% of the forecasted values are lower censored and 3% are upper censored. The likelihood function is in appendix Appendix C. A point estimate for the forecast aggregated across customers is:

$$\hat{Y}_t = \sum_i y_{i,t}. \quad (4)$$

Robust standard errors allow for heteroscedasticity and autocorrelation and use Newey-West weights.

3.2. Forecasting Intervals

We simulate a load aggregator bidding DLC into the forward energy market the day before a DLC event occurs. For example, a load aggregator would place a bid in the forward market on August 14 for a DLC event that is to occur on August 15. The aggregator would have data up to and including August 13 to forecast load for the August 15 event. The forecasts reported in this paper are computed the same way. The August 15 forecast uses parameter estimates constructed with data up to and including August 13. The August 16 forecast uses data up to and including August 14. Each AC i , therefore, has different parameters $\hat{\beta}_{i,d}$ for each day d . The estimates, $\hat{\beta}_{i,d}$ and $\hat{\sigma}_{i,d}$ use all available $X_{i,t}$ where $t < d - 2$. To ensure at least two weeks of data, the forecasts were made from August 15 to October 1.

239 The starting values used for the nonlinear optimization routine for August 15 were
 240 the ordinary least squares estimates using data up to August 14. The starting values
 241 for each successive day $d > 1$ are the ML estimates from the previous day: $\hat{\beta}_{i,d-1}$ and
 242 $\hat{\sigma}_{i,d-1}$. Results were not sensitive to changes in the starting values.
 243 The forecasted latent variable for AC i at time t on day d is:

$$\hat{y}_{i,t}^* = \mathbf{X}_t' \hat{\beta}_{i,d} \forall t \in d. \quad (5)$$

244 Simulation is used to compute aggregate point forecasts and forecasting intervals. A
 245 simulated random error is added to the latent estimate and then censored to obtain M
 246 observed load estimates $v_{i,t,m}$ for each AC i at each time $t \in d$, $v_{i,t,m}^* = \hat{y}_{i,t}^* + e_{i,t,m}$

$$v_{i,t,m} = \begin{cases} 0 & v_{i,t,m}^* \leq 0 \\ v_{i,t,m}^* & 0 < v_{i,t,m}^* < \lambda_i \\ \lambda_i & \lambda_i \leq v_{i,t,m}^* \end{cases} \quad (6)$$

247 Sum across ACs to obtain M aggregated loads at each time, t

$$Y_{t,m} = \sum_{i=1}^N v_{i,t,m} \text{ for } m = 1, \dots, M. \quad (7)$$

248 ACs in the same electric utility are in the same geographic region. This raises the concern
 249 that their shock can be correlated. If ignored, the forecast intervals for the aggregate
 250 load will be too optimistic. The Gaussian errors are allowed to have correlation ρ across
 251 ACs for each t in each utility. The m^{th} simulation draws a random error, $e_{i,t,m}$ from
 252 $N(0, \Sigma_d)$. For N ACs in a single utility the covariance matrix is

$$\Sigma_d = \begin{bmatrix} \hat{\sigma}_{1,d}^2 & \rho \hat{\sigma}_{1,d} \hat{\sigma}_{2,d} & \cdots & \rho \hat{\sigma}_{1,d} \hat{\sigma}_{N,d} \\ \rho \hat{\sigma}_{2,d} \hat{\sigma}_{1,d} & \hat{\sigma}_{2,d}^2 & \cdots & \rho \hat{\sigma}_{2,d} \hat{\sigma}_{N,d} \\ \vdots & \vdots & \ddots & \vdots \\ \rho \hat{\sigma}_{N,d} \hat{\sigma}_{1,d} & \rho \hat{\sigma}_{N,d} \hat{\sigma}_{2,d} & \cdots & \hat{\sigma}_{N,d}^2 \end{bmatrix}. \quad (8)$$

253 The correlation ρ is selected to have the highest agreement between the observed fore-
 254 casting intervals and the theoretical forecasting intervals for the aggregate load. This
 255 can be thought of as a generalized method of moments (GMM) estimation problem. See
 256 Appendix C for details.

257 4. Results

258 The Tobit model described in section 3 is estimated for each AC in the sample. The
 259 fit of the individuals models are summarized in Figures D.2, D.3 and D.4. These figures
 260 display the upper and lower quartiles and the mean of the t-statistics for the models
 261 parameters. Figure D.2 displays the parameters for the twenty four hourly dummy
 262 variables. Figure D.3 displays the twenty four parameters for the hourly dummy variables
 263 crossed with the temperature. Figure D.4 displays the parameters for dummy variables
 264 for DR events, dummy variables for post DR events and the lagged temperature.

265 Because this paper is concerned with accuracy in forecasting, insignificant t-statistics
266 were not dropped from the model estimated for each individual customer. The deci-
267 sion to include a variable in the model estimated for each customer was based on the
268 improved forecast of the aggregate load. However, the vast majority of the t-statistics
269 were statistically significant. Most of the insignificant t-statistics were associated with
270 customers with short samples.

271 The estimated models are used to forecast AC load each hour from August 15, 2010
272 to October 1, 2010. The following steps produce hourly aggregate AC load forecasts for
273 each day during the simulation period.

- 274 1. Estimate the Tobit models for data collected up to day $d - 2$.
- 275 2. Calculate an hourly forecast for each AC for day d .
- 276 3. Simulate individual forecasts and aggregate to get expected load and forecasting
277 intervals.
- 278 4. Repeat 1 - 3 for $d + 1, d + 2, \dots$.

279 The flowchart in Figure D.1 illustrates the forecasting strategy.

280 [Figure 1 about here.]

281 [Figure 2 about here.]

282 [Figure 3 about here.]

283 [Figure 4 about here.]

284 Figure D.5 shows the actual and Tobit forecasted AC usage for a typical household
285 for a 6 day period.

286 [Figure 5 about here.]

287 Aggregate forecasts were produced using the simulation method described in section
288 3. For each electric utility we estimated a correlation coefficient, ρ , to represent the corre-
289 lation of the forecast errors between AC's in the same utility. This coefficient is necessary
290 to calculate aggregate forecasts for each utility since AC usage amongst customers in the
291 same region is not IID. Households in the same utility territory are assumed to be exposed
292 to the same weather patterns and random shocks (see section 3.2 and Appendix D).

293 Table D.3 shows the correlation coefficients for each electric utility. As expected, the
294 estimated correlation coefficients are significantly positive for each utility. If this had
295 been ignored, the confidence intervals for the forecasts would have been unrealistically
296 optimistic (i.e. too small). A set of accurate confidence intervals are graphed in Figure
297 D.6. These are representative 50% confidence intervals and the point forecasts with the
298 actual observed load for the PEPCO customers.

299 [Table 3 about here.]

300 [Figure 6 about here.]

301 The Tobit based forecasts are compared to the default CBL used in the PJM RTO
 302 since PEPCO is in PJM territory. The default CBL in PJM is the average hourly
 303 load profile from the 4 highest load days of the previous 5 similar day types (weekdays,
 304 Saturdays, Sundays/holidays) [22]. The mean squared error (MSE) is used to quantify
 305 and compare the forecast accuracy of the two methods:

$$MSE = \frac{1}{T} \sum_t (Y_t - \hat{Y}_t)^2 \quad (9)$$

306 where Y_t is the aggregated AC load at time t , $\hat{Y}_t$ is an estimator for Y_t and T is the
 307 total number of time intervals. Table D.4 shows the MSE for the default CBL and Tobit
 308 models for each region and has the average hourly customer load as a reference point.
 309 The Tobit model has a MSE that is an order of magnitude lower than the default CBL
 310 indicating that the forecast error is substantially lower and the Tobit model is more
 311 accurate. Representative examples of the Tobit based forecasts, currently used default
 312 CBL and actual values are given in Figures D.5 and D.6. Figure D.6 shows the actual
 313 load, Tobit based forecast and 50% forecast interval for the PEPCO utility. Figure D.7
 314 compares the default CBL forecast to the Tobit based forecast.

315 [Figure 7 about here.]

316 Since the Tobit model presented here uses hourly temperatures, it is important to see
 317 how well it performs over a range of temperatures. Figure D.8 shows the forecast errors
 318 plotted against the ambient temperature. At high temperatures there is a tendency to
 319 over-forecast. One possible explanation is that vacations occur more frequently at the
 320 end of summer. This is the time period where the forecasts were made, based on data
 321 collected earlier in the summer. A full summer of training data may improve the forecasts
 322 by allowing monthly indicators in the model.

323 [Figure 8 about here.]

324 [Table 4 about here.]

325 5. Policy Implications and Discussion

326 Load forecasts are an important tool in grid management to optimize resources. They
 327 are also indispensable in demand response programs where demand resources, unlike
 328 generators, must accurately estimate future loads in order to offer them into markets.
 329 As demand response increases in the U.S., forecasting, measurement, and verification for
 330 direct load control will become more important. We focused on AC loads, a substantial
 331 demand side resource in residential customers.

332 This paper introduces a censored regression model to predict AC loads in a utility
 333 territory. This model is much simpler to use compared to non-parametric methods and
 334 requires less data. Utilities and ISO's are already very familiar with regression models
 335 for load forecasts, so this incremental change is easy to implement in the near term. This
 336 model is more accurate than the moving averages used by most ISO's, and is simple,
 337 easy and cheap to implement. We intend to further refine the model in future work.

338 6. Acknowledgments

339 The authors thank Jay Apt for valuable comments and discussion. Pepco Holdings,
340 Inc. generously provided the data.

341 This work was supported in part by grants from the Alfred P. Sloan Foundation and
342 EPRI to the Carnegie Mellon Electricity Industry Center; the Doris Duke Charitable
343 Foundation, the Department of Energy National Energy Technology Laboratory, and the
344 Heinz Endowments to the RenewElec program at Carnegie Mellon University; the U.S.
345 National Science Foundation under Awards SES-0949710 and DGE-0750271; Portuguese
346 Foundation for Science and Technology (Fundação para a Ciência e a Tecnologia), number
347 SFRH/BD/33764/2009; and the Department of Energy under grants EE-OE0000300 and
348 DE-FOA-0000058.

349 This report was prepared as an account of work sponsored by an agency of the United
350 States Government. Neither the United States Government nor any agency thereof, nor
351 any of their employees, makes any warranty, express or implied, or assumes any legal
352 liability or responsibility for the accuracy, completeness, or usefulness of any information,
353 apparatus, product, or process disclosed, or represents that its use would not infringe
354 privately owned rights. Reference herein to any specific commercial product, process,
355 or service by trade name, trademark, manufacturer, or otherwise does not necessarily
356 constitute or imply its endorsement, recommendation, or favoring by the United States
357 Government or any agency thereof. The views and opinions of authors expressed herein
358 do not necessarily state or reflect those of the United States Government or any agency
359 thereof.

360 Appendix A. Glossary

361 **Customer baseline** An estimate of what a demand response resource would have con-
362 sumed had it not been participating in demand response. It is used to calculate
363 the quantity of energy reduced during a demand response event.

364 **Demand response** Changes in electric use by demand-side resources from their normal
365 consumption patterns in response to changes in the price of electricity, or to incen-
366 tive payments designed to induce lower electricity use at times of high wholesale
367 market prices or when system reliability is jeopardized [17].

368 **Direct load control** A demand response activity by which the program sponsor re-
369 motely shuts down or cycles a customer's electrical equipment (e.g., air condi-
370 tioner, water heater) on short notice. Direct load control programs are primarily
371 offered to residential or small commercial customers [17].

372 **Independent system operator (ISO) or Regional transmission organization (RTO)**
373 A federally regulated entity established to coordinate regional transmission and en-
374 sure the safety and reliability of the electric system [17].

375 **Reserves** A service in which demand resources are obligated to be available to pro-
376 vide demand reduction upon deployment by the system operator, based on reserve
377 capacity requirements that are established to meet reliability standards [17].

378 Appendix B. Data Cleaning Protocol

379 Six ACs had too few observations (fewer than 9 days; the remaining ACs all had
 380 more than 52 days) and were removed. The data from 2 ACs were collected at the wrong
 381 frequency (not three minute) and were removed. Some of the ACs showed they were
 382 rarely used (less than 3 hours during the entire summer). Data from these ACs were
 383 not analyzed since they provided no information for forecasts. Several ACs had 20A
 384 loggers even though the AC capacity was greater than 20A. If an AC logged data at 20A
 385 more than 10% of the time, we assume that it required a higher amperage logger, and
 386 discard the data. There were also several ACs that became stuck on a particular value.
 387 An AC was removed if it switched state (from off to on or vice versa) in fewer than 2%
 388 of its observations. Finally, we removed ACs that had unrealistically low readings (all
 389 observations below 3 amps). A summary of the AC data discarded is in table D.5.

390 [Table 5 about here.]

391 Appendix C. Tobit Model Derivations

392 This is a derivation of the Tobit model. Lower and upper bounds for censoring are
 393 represented here as a and b . For this paper, $a = 0$ and $b = \lambda$. To simplify the notation,
 394 drop i , the index for AC from the derivations. The latent variable is:

$$y_t^* = \mathbf{X}_t' \boldsymbol{\beta} + \epsilon_t. \quad (\text{C.1})$$

395 The censored variable is:

$$y_t = \begin{cases} a & y_t^* \leq a \\ y_t^* & a < y_t^* < b \\ b & b \leq y_t^*. \end{cases} \quad (\text{C.2})$$

396 Define the indicator variables:

$$I_t(a) = \begin{cases} 1 & y_t^* \leq a \\ 0 & a < y_t^*, \end{cases} \quad I_t(ab) = \begin{cases} 1 & a < y_t^* < b \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad I_t(b) = \begin{cases} 1 & b \leq y_t^* \\ 0 & y_t^* < b. \end{cases} \quad (\text{C.3})$$

397 Assume the latent variable, y_t^* has distribution $N(\mu, \sigma^2)$. The entire probability density
 398 of the lower censored region is applied at a , and the same for the upper censored region
 399 at b . The probability density function for the censored variable is:

$$f(y_t) = \begin{cases} \Phi\left(\frac{a-\mu}{\sigma}\right) & y_t^* \leq a \\ \frac{1}{\sigma} \phi\left(\frac{y_t-\mu}{\sigma}\right) & a < y_t^* < b \\ 1 - \Phi\left(\frac{b-\mu}{\sigma}\right) & b \leq y_t^*. \end{cases} \quad (\text{C.4})$$

400 The log-likelihood function $\ell(\mathbf{Y}_\tau, \mathbf{X}_\tau, \theta)$ can be expressed in terms of the vector of

parameters $\theta = [\beta' \sigma]'$ and τ , the length of the time-series

$$\begin{aligned}\ell(\mathbf{Y}_\tau, \mathbf{X}_\tau, \theta) &= \sum_{t=1}^{\tau} \ln(f(y_t)) \\ &= \sum_{t=1}^{\tau} \left[I_t(a) \ln \left(\Phi \left(\frac{a - \mu}{\sigma} \right) \right) + I_t(ab) \ln \left(\frac{1}{\sigma} \phi \left(\frac{y_t - \mu}{\sigma} \right) \right) + I_t(b) \ln \left(\Phi \left(\frac{\mu - b}{\sigma} \right) \right) \right]\end{aligned}\quad (\text{C.5})$$

Let $\mu = \mathbf{X}_t' \beta$ to obtain

$$\ell(\mathbf{Y}_\tau, \mathbf{X}_\tau, \theta) = \sum_{t=1}^{\tau} \left[I_t(a) \ln \left(\Phi \left(\frac{a - \mathbf{X}_t' \beta}{\sigma} \right) \right) - I_t(ab) \frac{1}{2} \ln (2\pi\sigma^2) \right. \quad (\text{C.6})$$

$$\left. - I_t(ab) \frac{(y_t - \mathbf{X}_t' \beta)^2}{2\sigma^2} + I_t(b) \ln \left(\Phi \left(\frac{\mathbf{X}_t' \beta - b}{\sigma} \right) \right) \right]. \quad (\text{C.7})$$

The gradient is $\nabla \ell(\mathbf{Y}_\tau, \mathbf{X}_\tau, \theta) = \left[\frac{\partial \ell}{\partial \beta'} \quad \frac{\partial \ell}{\partial \sigma'} \right]'$ where

$$\frac{\partial \ell}{\partial \beta} = \sum_{t=1}^{\tau} \left[I_t(a) \frac{\phi \left(\frac{a - \mathbf{X}_t' \beta}{\sigma} \right)}{\Phi \left(\frac{a - \mathbf{X}_t' \beta}{\sigma} \right)} \left(-\frac{\mathbf{X}_t}{\sigma} \right) + I_t(ab) \frac{(y_t - \mathbf{X}_t' \beta) \mathbf{X}_t}{\sigma^2} + I_t(b) \frac{\phi \left(\frac{\mathbf{X}_t' \beta - b}{\sigma} \right)}{\Phi \left(\frac{\mathbf{X}_t' \beta - b}{\sigma} \right)} \left(\frac{\mathbf{X}_t}{\sigma} \right) \right]$$

$$\frac{\partial \ell}{\partial \sigma} = \sum_{t=1}^{\tau} \left[I_t(a) \frac{\phi \left(\frac{a - \mathbf{X}_t' \beta}{\sigma} \right)}{\Phi \left(\frac{a - \mathbf{X}_t' \beta}{\sigma} \right)} \left(\frac{\mathbf{X}_t' \beta - a}{\sigma^2} \right) - I_t(ab) \frac{1}{\sigma} \right. \quad (\text{C.8})$$

$$\left. + I_t(ab) \frac{(y_t - \mathbf{X}_t' \beta)^2}{\sigma^3} + I_t(b) \frac{\phi \left(\frac{\mathbf{X}_t' \beta - b}{\sigma} \right)}{\Phi \left(\frac{\mathbf{X}_t' \beta - b}{\sigma} \right)} \left(\frac{b - \mathbf{X}_t' \beta}{\sigma^2} \right) \right]. \quad (\text{C.9})$$

The gradient contains the indeterminate form $\lim_{z \rightarrow -\infty} \frac{\phi(z)}{\Phi(z)} = \frac{0}{0}$. To evaluate the gradient we apply L'Hôpital's rule for $z < -38$ use $\frac{\phi(z)}{\Phi(z)} \approx -z$.

The HAC variance is based on [41]. The auto-covariance is

$$\gamma(\delta) = \frac{1}{\tau - \delta} \sum_{t=1}^{\tau - \delta} \nabla \ell(\mathbf{X}_t, \theta) \nabla \ell(\mathbf{X}_{t+\delta}, \theta)'. \quad (\text{C.10})$$

The Newey-West weights are expressed as $\omega(\delta) = 1 - \frac{\delta}{\Delta + 1}$ where $\Delta \leq \sqrt{\tau}$, we use $\Delta = \tau^{0.4}$. The variance of likelihood estimate with the HAC correction is expressed in terms of the auto-covariance and Newey-West weights:

$$\text{Var} \nabla \ell = \omega(0) \hat{\gamma}(0) + \sum_{\delta=1}^{\Delta} \omega(\delta) (\hat{\gamma}(\delta) + \hat{\gamma}(-\delta)'). \quad (\text{C.11})$$

The variance with HAC correction also uses the Hessian $H = \nabla^2 \ell$

$$Var(\theta) = (-H)^{-1} Var \nabla \ell (-H/\tau)^{-1}.$$

410

411 **Appendix D. Confidence Interval for ρ**

412 The confidence interval for ρ is obtained by inverting the GMM objective function,
413 also called the DM test in Newey and McFadden [42].

414 The parameter estimates for the individual Tobit models are obtained by maximum
415 likelihood. For a single utility collect all the estimates in $\hat{\psi}$, i.e. the ML $\hat{\theta}$ estimates from
416 each AC. The gradient for each AC, $\nabla \ell(\mathbf{Y}_\tau, \mathbf{X}_\tau, \theta)$, set to zero are the moment conditions
417 for these estimates. Together these are the scores and will be denoted $S(\psi)$. The moment
418 conditions are satisfied at the ML estimates. The estimated models and the normality
419 assumption on the errors (for the Tobit model) implies the forecast distributions for each
420 AC. These can be used to obtain the forecast distribution for the utility's aggregate load
421 for AC generation. These AC's are in the same utility and will experience correlated
422 shocks. The common correlation of ρ is assumed between every AC in the utility.

423 Instead of calculating the distribution of the aggregate load forecast analytically, the
424 correlated errors will be simulated. The implied individual forecasts will be summed
425 to simulate an aggregate forecast. The order statistics from the simulated aggregate
426 forecasts are a consistent estimates for the critical values for the forecast distribution.
427 For the population parameter value ρ_0 the observed aggregate load for the utility, $Y_{d,h}$,
428 will be in the simulated forecast interval with the correct probability. The 95%, 90% and
429 50% forecast intervals give moment conditions for the estimation of ρ

$$G(\psi, \rho) = \frac{1}{\tau} \sum_{d,h} \begin{bmatrix} \chi \left(\hat{Y}(\psi, \rho)_{d,h,[.025M]} \leq Y_{d,h} \leq \hat{Y}(\psi, \rho)_{d,h,[.975M]} \right) & - & .95 \\ \chi \left(\hat{Y}(\psi, \rho)_{d,h,[.05M]} \leq Y_{d,h} \leq \hat{Y}(\psi, \rho)_{d,h,[.95M]} \right) & - & .9 \\ \chi \left(\hat{Y}(\psi, \rho)_{d,h,[.25M]} \leq Y_{d,h} \leq \hat{Y}(\psi, \rho)_{d,h,[.75M]} \right) & - & .5 \end{bmatrix}. \quad (\text{D.1})$$

430 where $\hat{Y}(\psi, \rho)_{d,h,k}$ is the k^{th} order statistics of the simulated aggregate forecasts, $\chi(\cdot)$ is
431 the indicator function and M is the number of simulations. This paper used $M = 1000$.
432 The results did not change for larger M .

433 The GMM objective function can be written

$$Q(\psi, \rho) = \begin{bmatrix} S(\psi) & G(\psi, \rho) \end{bmatrix} \begin{bmatrix} I & 0 \\ 0 & W \end{bmatrix} \begin{bmatrix} S(\psi) \\ G(\psi, \rho) \end{bmatrix}. \quad (\text{D.2})$$

434 Because the Tobit parameters ψ are just identified with the scores from the MLE,
435 the weighting matrix only influences the last three moments that estimate ρ .

436 Optimal two-step GMM is used with $W = I$ in the first step. Because ρ does not
437 enter the scores, the ψ estimates do not change as ρ takes different values. The confidence
438 intervals for ρ are obtained by

$$\tau(Q(\hat{\psi}, \hat{\rho}) - Q(\hat{\psi}, \rho_0)) \sim \chi^2_{(1)} \quad (\text{D.3})$$

where $\hat{\psi}$ are the MLE Tobit estimates and $\hat{\rho}$ is the efficient second round simulated method of moments estimate. For this data set $\tau = (48 \text{ days}) \times (24 \text{ hours}) = 1152$.

References

- [1] PJM, PJM 2011 annual report, Tech. rep., PJM (2012).
URL <http://www.pjm.com/about-pjm/who-we-are/~media/about-pjm/newsroom/annual-reports/2011-annual-report.ashx>
- [2] D. Wang, S. Parkinson, W. Miao, H. Jia, C. Crawford, N. Djilali, Hierarchical market integration of responsive loads as spinning reserve, *Applied Energy* 104 (2013) 229–238.
- [3] D. Callaway, Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy, energy conversion and management, *Energy Conversion and Management* 50 (2009) 1389–1400.
- [4] S. Koch, F. S. Barcenas, G. Andersson, Using controllable thermal household appliances for wind forecast error reduction, in: IFAC Conference on Control Methodologies and Technology for Energy Efficiency, Portugal, 2010.
- [5] S. Newell, F. Felder, Quantifying demand response benefits in PJM, Tech. rep., The Brattle Group (2007).
URL http://http://conserveland.org/libraries/3/library_items/195
- [6] T. Broecker, J. Fuller, F. Tuffner, D. Chassin, N. Djilali, Modeling framework and validation of a smart grid and demand response system for wind power integration, *Applied Energy* 113 (2014) 199–207.
- [7] DOE, 20% wind power by 2030, Tech. rep., Department of Energy (jul 2008).
URL <http://www.20percentwind.org/20p.aspx?page=Report>
- [8] M. Paulus, F. Borggrefe, The potential of demand-side management in energy intensive industries for electricity markets in Germany, *Applied Energy* 88 (2011) 432–441.
- [9] H. Aalami, M. P. Moghaddam, G. Yousefi, Demand response modeling considering interruptible/curtailable loads and capacity market programs, *Applied Energy* 87 (2010) 243–250.
- [10] A. Arteconi, N. Hewitt, F. Polonara, State of the art of thermal storage for demand-side management, *Applied Energy* 93 (2012) 371–389.
- [11] M. Sullivan, J. Bode, B. Kellow, S. Woehleke, J. Eto, Using residential AC load control in grid operations: PG&E’s ancillary service pilot, Tech. rep. (jun 2012).
URL http://www.ercot.com/content/meetings/dswg/keydocs/2012/0920PM/Using_AC_in_Grid_Operation_-_PGE_ancillary_services_pilot.pdf
- [12] EIA, Annual energy outlook 2012, DOE/EIA-0383, Energy Information Administration (jun 2012b).
URL [http://www.eia.gov/forecasts/aeo/pdf/0383\(2012\).pdf](http://www.eia.gov/forecasts/aeo/pdf/0383(2012).pdf)
- [13] G. Strbac, Demand side management: Benefits and challenges, *Energy Policy* 36 (12) (2008) 4419 – 4426, foresight Sustainable Energy Management and the Built Environment Project.
doi:10.1016/j.enpol.2008.09.030.
URL <http://www.sciencedirect.com/science/article/pii/S0301421508004606>
- [14] K. Hamilton, N. Gulhar, Taking demand response to the next level, *IEEE Power & Energy Magazine* 8 (3) (2010) 60–65.
- [15] B. Shen, G. Ghatikar, Z. Lei, J. Li, G. Wikler, P. Martin, The role of regulatory reforms, market changes, and technology development to make demand response a viable resource in meeting energy challenges, *Applied Energy* doi:http://dx.doi.org/10.1016/j.apenergy.2013.12.069.
- [16] C. Woo, P. Sreedharan, J. Hargreaves, F. Kahrl, J. Wang, I. Horowitz, A review of electricity product differentiation, *Applied Energy* 114 (2014) 262–272.
- [17] FERC, Assessment of demand response and advanced metering 2011, Tech. rep., Federal Energy Regulatory Commission (2011).
URL <http://www.ferc.gov/legal/staff-reports/11-07-11-demand-response.pdf>
- [18] DOE, , Tech. rep., Department of Energy (jul 2012).
URL <http://energy.gov/sites/prod/files/Smart%20Grid%20Investment%20Grant%20Program%20-%20Progress%20Report%20July%202012.pdf>
- [19] FERC, Demand response compensation in organized wholesale energy markets, Docket No. RM10-17-000 Order No. 745, Federal Energy Regulatory Commission (mar 2011).
URL <http://www.ferc.gov/whats-new/comm-meet/2011/121511/E-4.pdf>
- [20] C. Grimm, Evaluating baselines for demand response programs, in: 2008 AEIC Load Research Workshop, San Antonio, TX, 2008.

- [21] KEMA, PJM empirical analysis of demand response baseline methods, Tech. rep., KEMA (apr 2011).
URL <http://www.pjm.com/~media/committees-groups/subcommittees/drs/20110613/20110613-item-03b-cbl-analysis-report.ashx>
- [22] PJM, 2011 DSR annual report, Tech. rep., PJM (2012).
- [23] CASIO, Business practice manual for metering, Manual Version 6, California Independent System Operator (2012).
URL <https://bpm.caiso.com/bpm/bpm/version/000000000000168>
- [24] NYISO, Emergency demand response program manual, , New York Independent System Operator (dec 2010).
URL http://www.nyiso.com/public/webdocs/products/demand_response/emergency_demand_response/edrp_mnl.pdf
- [25] ISONE, Market rule 1 section III.8, Tech. rep., Independent System Operator of New England (jun 2012).
URL http://www.iso-ne.com/regulatory/tariff/sect_3/index.html
- [26] ERCOT, Emergency interruptible load service default baseline methodologies, Tech. rep., Electric Reliability Council of Texas (2012).
URL http://www.ercot.com/content/services/programs/load/eils/keydocs/Default_Baseline_Methodologies_REVISED-FINAL.doc
- [27] D. Bargiotas, J. Birdwell, Residential air conditioner dynamic model for direct load control, Power Delivery, IEEE Transactions on 3 (4) (1988) 2119–2126. doi:10.1109/61.194024.
- [28] A. Molina, A. Gabaldon, J. Fuentes, C. Alvarez, Implementation and assessment of physically based electrical load models: Application to direct load control residential programmes, Generation, Transmission and Distribution, IEE Proceedings- 150 (1) (2003) 61 – 66. doi:10.1049/ip-gid:20020750.
- [29] M. Gustafson, J. Baylor, G. Epstein, Estimating air conditioning load control effectiveness using an engineering model, Power Systems, IEEE Transactions on 8 (3) (1993) 972–978. doi:10.1109/59.260904.
- [30] G. Beccali, M. Cellura, V. L. Brano, A. Orioli, Single thermal zone balance solved by transfer function method, Energy and Buildings 32 (2005) 1268–1277.
- [31] A. Pahwa, C. Brice, Modeling and system identification of residential air conditioning load, Power Apparatus and Systems, IEEE Transactions on 104 (6) (1985) 1418 –1425. doi:10.1109/TPAS.1985.319247.
- [32] S. El-Ferik, S. A. Hussain, F. M. Al-Sunni, Identification and weather sensitivity of physically based model of residential air-conditioners for direct load control: A case study, Energy and Buildings 38 (8) (2006) 997 – 1005. doi:10.1016/j.enbuild.2005.11.010.
URL <http://www.sciencedirect.com/science/article/pii/S037877880500232X>
- [33] S. Kamoun, R. Malhamé, Convergence characteristics of a maximum likelihood load model identification scheme, Automatica 28 (5) (1992) 885 – 896. doi:10.1016/0005-1098(92)90142-3.
URL <http://www.sciencedirect.com/science/article/pii/0005109892901423>
- [34] N. Ryan, J. Powers, S. Braithwait, B. Smith, Generalizing direct load control program analysis: implementation of the duty cycle approach, Power Systems, IEEE Transactions on 4 (1) (1989) 293 –299. doi:10.1109/59.32491.
- [35] Y. Penya, C. Borges, D. Agote, I. Fernandez, Short-term load forecasting in air-conditioned non-residential buildings, in: Industrial Electronics (ISIE), 2011 IEEE International Symposium on, 2011, pp. 1359 –1364. doi:10.1109/ISIE.2011.5984356.
- [36] L. Xuemei, D. Lixing, L. Yan, X. Gang, L. Jibin, Hybrid genetic algorithm and support vector regression in cooling load prediction, in: Knowledge Discovery and Data Mining, 2010. WKDD '10. Third International Conference on, 2010, pp. 527 –531. doi:10.1109/WKDD.2010.136.
- [37] M. Beccali, M. Cellura, V. Lo Brano, A. Marvuglia, Forecasting daily urban electric load profiles using artificial neural networks, Energy Conversion and Management 45 (18-19) (2004) 2879–2900.
- [38] M. Beccali, M. Cellura, V. Lo Brano, A. Marvuglia, Short-term prediction of household electricity consumption: Assessing weather sensitivity in a Mediterranean area, Renewable and Sustainable Energy Reviews 12 (2008) 2040–2065.
- [39] J. Eto, J. Nelson-Hoffman, E. Parker, C. Bernier, P. Young, D. Sheehan, J. Kueck, B. Kirby, The demand response spinning reserve demonstration—measuring the speed and magnitude of aggregated demand response, in: System Science (HICSS), 2012 45th Hawaii International Conference on, 2012, pp. 2012 –2019. doi:10.1109/HICSS.2012.39.
- [40] J. Tobin, Estimation of relationships for limited dependent variables, Econometrica 26 (1) (1958)

- 554 pp. 24–36.
555 URL <http://www.jstor.org/stable/1907382>
556 [41] A. B. Bernard, M. R. Busse, Consistent standard errors in panel Tobit with autocorrelation, Working
557 paper 03-25, Tuck Business School (Aug. 2003).
558 URL <http://ssrn.com/abstract=439061orhttp://dx.doi.org/10.2139/ssrn.439061>
559 [42] W. K. Newey, D. McFadden, Chapter 36: Large sample estimation and hypothesis testing, in: R. F.
560 Engle, D. McFadden (Eds.), Handbook of Econometrics, Vol. 4, Elsevier, 1986, pp. 2111–2245.

561 **Figure Captions**

562	D.1	Flowchart of forecasting strategy.	18
563	D.2	Median, upper-quartile, lower-quartile of the t-statistics for $\hat{\beta}_{D_h,i}$	19
564	D.3	Median, upper-quartile, lower-quartile of the t-statistics for $\hat{\beta}_{TD_h,i}$	20
565	D.4	Median, upper-quartile, lower-quartile of the t-statistics for $\hat{\beta}_{E,i}$, $\hat{\beta}_{P,i}$, $\hat{\beta}_{T1,i}$	21
566	D.5	Actual hourly AC usage and day-ahead forecast for a typical individual	
567		household.	22
568	D.6	Average hourly actual AC usage and day-ahead forecast for PEPCO with	
569		50% forecast intervals over a six-day period.	23
570	D.7	Comparison of day-ahead Tobit forecast and default PJM CBL forecast	
571		for PEPCO data over a six-day period.	24
572	D.8	Forecast errors plotted against the ambient temperature from August 15,	
573		2010 to October 1, 2010 in PEPCO.	25

PEPCO

574 **List of Tables**

575	D.1	Temperature statistics during the period July - September 2010 from each	
576		region where the air conditioners are located.	27
577	D.2	Summary statistics from AC data set.	28
578	D.3	Estimated correlation coefficients for each electric utility.	29
579	D.4	Mean squared errors.	30
580	D.5	Number of ACs discarded from the dataset.	31

Table D.1: Temperature statistics during the period July - September 2010 from each region where the air conditioners are located.

	Utility		
	PEPCO	ACE	Delmarva
Minimum temperature ($^{\circ}\text{C}$)	13	9	12
Maximum temperature	37	38	36
Mean temperature	24	23	23
Standard deviation	4	5	4

Table D.2: Summary statistics from AC data set.

Variables	PEPCO	Utility		Total
		ACE	Delmarva	
Number of total air conditioners	181	72	214	467
Air conditioners cycling at 50%	58	72	88	218
Air conditioners cycling at 75%	68	0	68	136
Air conditioners cycling at 100%	55	0	58	113
Air conditioner size < 2 kW	50	10	49	109
Air conditioner size ≥ 2 and < 3 kW	83	42	110	235
Air conditioner size ≥ 3 and < 4 kW	36	17	50	103
Air conditioner size ≥ 4 kW	12	3	5	20
Average air conditioner size	2.6	2.7	2.5	2.5

Table D.3: Estimated correlation coefficients for each electric utility.

	PEPCO	Utility ACE	Delmarva
Correlation Coefficient ρ	.19	.45	.26
95% Confidence Interval for ρ	(.10, .34)	(.25, .75)	(.15, .47)

Table D.4: Mean squared errors.

	Tobit	default CBL	ave load per customer (kW)
PEPCO	0.034	0.260	0.55
ACE	0.041	0.347	0.49
Delmarva	0.027	0.302	0.62

Table D.5: Number of ACs discarded from the dataset.

Data Problem	Number of ACs
Length of time less than 9 days	6
Incorrect time intervals	2
AC nearly always off	10
AC maxed out	11
Stuck values	13
Values unrealistically low	27
Total discarded	69