Take It From the Top! An Innovative Approach to Residential and Commercial Program Savings Estimation Using AMI Data

Stephen Grover, Evergreen Economics, Portland, OR
John Cornwell, Evergreen Economics, Portland, OR
Sarah Monohon, Evergreen Economics, Portland, OR
Ted Helvoigt, Evergreen Economics, Portland, OR

ABSTRACT

This paper presents a new energy savings estimation approach—referred to as the AMI Customer Segmentation (AMICS) model—that provides accurate impact estimates by taking full advantage of hourly AMI data. This approach differs from more traditional methods in that it automatically develops a large number of customer-specific regressions covering a wider range of customer types, weather conditions and time periods. The approach uses a type of hierarchical linear model—the random coefficients model—that allows savings estimates to be tailored more closely to individual customer characteristics. This is accomplished by first grouping customer consumption data into different categories based on energy use and weather conditions. Separate models are then estimated for each usage/weather category, which allows for separate load shape predictions for very specific customer types.

The AMICS model specification was tested using data from both residential and commercial HVAC efficiency programs in California. Using participant and AMI data from both of these programs, average daily load shapes were calculated for specific day types (weekday, weekend, seasonal) and used to estimate program impacts. When estimated load shapes were compared against a holdout sample of customers, the random coefficients model performed extremely well; load shape estimates were within 1 percent of the holdout sample. Besides producing accurate estimates of energy use, the automated categorization and modeling processes allow for separate savings estimates and load shapes to be developed easily for a variety of situations.

Introduction

As electric utilities transition to advanced metering infrastructure (AMI), a greater amount and richer source of consumption data are becoming available to evaluators. A single customer's metered data at one-hour intervals translate to over 700 data points per month, providing an opportunity for evaluators to better understand the impact that energy efficiency programs (and other factors) have on energy consumption during specific hours of the day, rather than a daily average derived from monthly data. A common concern among economists and other analysts working with monthly consumption data is that the aggregation conceals more than it reveals. The availability of short-interval meter data allows for potentially more accurate and robust models.

One of the key areas where AMI data have the potential to improve accuracy is in billing regression models used to estimate program energy impacts. Most of the literature to date has focused on using monthly consumption data, as these are typically all that have been available for estimating impacts at the program level. See the California Evaluation Framework (CPUC 2006) and the Uniform Methods Project (Agnew and Goldberg 2013) for a summary of the more traditional methods using monthly data. Other studies (particularly those in demand response programs) have utilized AMI data to estimate load shapes and demand impacts (Nexant 2014 for example), but these models typically have been developed manually for each specific situation and therefore have not been practical for addressing a large number of customer types and time periods. Other works such as Hsiao et al. (1989) provided an early application of the random coefficients model to energy efficiency, while Granderson et al. (2015) have begun to look at developing AMI regression models in a more systematic fashion. None of these past studies, however, have presented a method for efficiently developing a large number of models that are tailored to a wide range of customer types and time periods that take full advantage of the information contained in the AMI data.

To explore how AMI data could be used in billing regression models, the California investor-owned utilities[1] (IOUs) contracted with Evergreen Economics to conduct exploratory research, and a portion of these research results is presented in this paper. During the course of this research, it became apparent that an innovative new analysis method—one that we refer to as the AMI Customer Segmentation or AMICS model—has the potential to be a groundbreaking impact evaluation approach that fully utilizes the benefits of AMI data. As discussed in the remainder of this paper, we believe that the AMICS model represents a significant improvement over traditional billing regression models, as it provides an efficient method for tailoring impacts to specific customer conditions (e.g., day types, seasons, customer types). The AMICS model also proved to be very accurate when tested against a holdout sample of customers.

This current work focusing on commercial customers builds on an earlier application of the AMICS model to residential data,[2] and the results of the residential model are summarized here to provide context for the commercial modeling.

Analysis Methods

This paper presents the results of the AMICS modeling approach that utilizes data from the following residential and commercial HVAC programs:

- **SCE Residential Quality Installation (QI) Program Participant Data** – a dataset containing 1-hour interval whole house metered consumption on 2,039 homes that participated in the SCE QI Program between January 2012 and December 2014. The SCE QI Program is a California statewide program designed to achieve energy and demand savings through the installation of replacement split or packaged HVAC systems in accordance with industry standards. Program data include household and program participation information including the home climate zone and date of participation in the program.
- **SCE Commercial Quality Installation (CQI) Program Participant Data** – a dataset containing a mix of 1-hour and 15-minute interval whole building consumption data from 1,958 business customers that participated in the SCE CQI program from January 2014 to August 2016. The SCE CQI program is part of a California statewide program that targets HVAC installations in the commercial sector.

Each customer dataset was combined with weather data obtained from the National Oceanic and Atmospheric Administration (NOAA) to develop datasets with both energy consumption and weather data. Weather station data were selected based on proximity to each customer's zip code, matching climate zone, and availability of complete hourly data. Additional analysis was performed to

[1] The California investor-owned utilities include Pacific Gas and Electric (PG&E), Southern California Edison (SCE), San Diego Gas & Electric (SDG&E), and Southern California Gas Company (SoCalGas).

[2] See http://www.calmac.org/publications/AMI_Report_Volume_1_FINAL.pdf for the final report for the first residential application of the AMICS model.

identify unreasonably high or low temperature readings, based on the record high and low temperatures in each climate zone. Missing observations and temperatures identified as unreasonable were imputed using the next closest weather station if available; otherwise, they were imputed with the average of the preceding and following temperature reads.

Modeling Approach

With the most basic billing regression specification, the model uses monthly consumption data and produces a regression line that represents the average energy use or savings across all customers included in the sample, with the average often calculated at a monthly or annual level. Variations on the standard linear regression such as the fixed effects model can be developed to produce separate savings estimates for sub-groups of customers. With the availability of AMI data, however, modeling approaches need to be adapted to account for the additional information as well as the sheer volume of data. In theory, the traditional billing regression methods such as the fixed effects model can be adapted for use with hourly AMI data. However, the number of different coefficients (or separate models) required to capture the variations in energy use across hours, days and seasons—in addition to accounting for variation across customers—requires significant amounts of computer processing time, along with a separate process for efficiently evaluating the performance of different model specifications.

Rather than attempting to adapt a fixed effects model for use with a high volume of AMI data, this paper explores a different hierarchical modeling approach that first categorizes the data into manageable sub-groups and then uses regression analysis to estimate energy use within each group. The term "random coefficients model" refers to a type of linear hierarchical model that provides a distribution of model parameters across customer types and weather conditions.[3] The random coefficients model works by explicitly accounting for two separate sources of variability commonly found in interval energy-use data. The first, within-subject variability, represents the variation in energy usage *throughout the day* by an individual customer. The second, among-subject variability, represents the variation in energy use *across customers and varying weather conditions* experienced by each customer.[4]

To incorporate both types of variability and estimate customer load shapes, the random coefficients model utilizes a multi-stage process that is summarized in Figure 1.

[3] See Snijders and Bosker (2012) for a more detailed explanation of the random coefficients model, and a more general discussion of the different types of hierarchical linear regression models.
[4] See Helvoigt (2016) for a more detailed discussion of how these types of variability can be addressed in billing regression models.

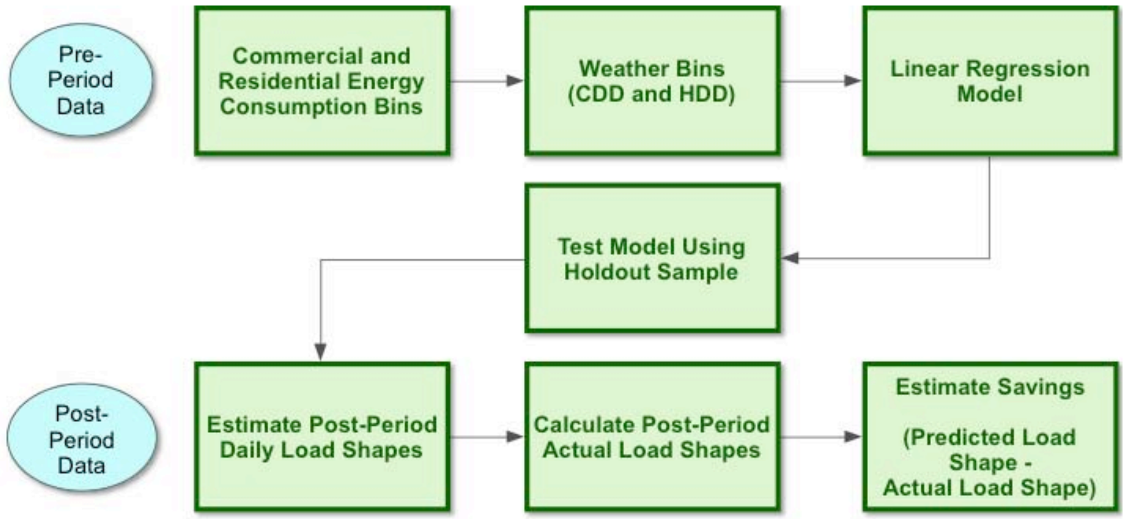

Figure 1. Summary of AMICS model savings estimation approach

First Stage – Binning Process. For the customers in the QI program, the first stage of the modeling approach uses a fixed effects regression model to create estimates of daily baseload electricity use for each customer, controlling for outside air temperature. The fixed effects model specification is as follows:

$$DailykWh_{i,t} = \alpha_i + \beta_1(CDD_{i,t}) + \beta_2(HDD_{i,t}) + \varepsilon_t$$

Where:

$DailykWh_{i,t}$ = Daily kWh consumption for customer i on day t.

$CDD_{i,t}$ = Cooling degree days (CDD) for customer i on day t.

$HDD_{i,t}$ = Heating degree days (HDD) for customer i on day t.

α_i = Customer specific constant (i.e., baseload weather normalized consumption)

β_1, β_2 = Coefficients estimated in the regression model

$\varepsilon_{i,t}$ = Random error assumed normally distributed

A characteristic of the fixed effects model is the estimation of a specific constant α_i, for every customer site. This constant varies by customer site and accounts for time-invariant effects on electricity consumption over the year. In the model specification above, the constant can be interpreted as site-specific baseload consumption after controlling for variation in outside air temperature (CDD and HDD, using a base temperature of 65 degrees Fahrenheit). Customers are then ranked in ascending order of baseload energy use and assigned to one of 20 “customer groups” based on each customer’s weather normalized energy usage, prior to program participation. In this way, customers with similar energy consumption are grouped together. Each group represents about 5 percent of total daily baseload electricity usage for the customers in our sample.[5] Because of this, the number of customers in each bin varies, but the amount of daily kWh each bin represents is approximately the same.

[5] The residential model uses 20 baseload energy usage groups, while the commercial model uses only 5. We selected a smaller number of baseload bins for commercial because these customers are segmented by two additional metrics. By the end of the binning process, we will have 20 residential customer groups (the 20 baseload

Two additional segmentation approaches are applied for the first stage binning of the commercial customers: total daily energy usage (including weather-dependent use) and load shape clustering. We found that these additional groupings are necessary to reduce the variation in energy usage among customers within each baseline usage group, as commercial customers are far more varied than residential in terms of their building sizes and operations. Our analysis found that segmenting commercial customers by the combinations of these three different metrics allows us to successfully identify separate categories of customers with similar energy usage patterns, without relying on additional information on business type or square footage. As an additional step, our final model drops the largest commercial customers (those with greater than 4,000 average daily kWh consumption) from the model. These very large customers have too much variation in their load shapes, and their inclusion in the model leads to less accurate forecasts for the remaining commercial customers. This final screen results in 31 of the largest customers being dropped from the analysis.

To create groupings on total daily usage, we calculate the average daily kWh consumption over a single pre-period year. Customers are then ranked in ascending order of total daily energy use and assigned to one of 12 total usage groups, with about 8 percent of customers in each bin.

With the CQI group, a cluster analysis is used to categorize commercial customers.[6] The cluster analysis is used to identify commercial customers with similar energy use during the pre-period. Cluster analysis is an unsupervised machine-learning algorithm designed to detect patterns in the data. The *k*-means clustering algorithm randomly assigns each customer's load shape to one of *k* clusters and then calculates the sum of the distance between each load shape and the centroid of the cluster it was assigned. Load shapes are then reassigned to the nearest cluster centroid and the process is repeated until the variation within each cluster cannot be improved. This *k*-means clustering is used to identify 14 unique clusters, each identifying a subset of commercial customers with similar average daily energy usage (magnitude) and load shape (hours of use) throughout the year. The benefit of using the cluster analysis is that similar customer groups can be created automatically from the AMI data, rather than relying on additional customer information such as hours of operation, building or business type that is often not tracked (or tracked accurately) by the utility.

Next, for both residential and commercial customers, every day of the study period is characterized (binned) in terms of the weather and day type. The weather groups are created by calculating cooling degree hours (CDH) for each hourly observation using a base temperature of 65 degrees Fahrenheit, and then taking the average of these hourly values to create a single cooling degree day (CDD) value for each home or business on each day (i.e., each "customer-day") in the study period, rounded up to the nearest integer. For models covering the heating season, this process is repeated to assign days to heating degree day (HDD) groups, again using a base temperature of 65 degrees Fahrenheit.[7] There are a total of 25 weather groups for both CDD and HDD, and categorizing days using outdoor temperature in this manner explicitly incorporates temperature into our modeling approach. To reflect possible differences in energy usage between weekends and weekdays, home-days are also binned based on the day type. Weekends are assigned to day type group 1, and weekdays are assigned to day type group 0.

Lastly, all groups are combined to create customer-day bins containing only one type of home or business on one type of day. These bins describe the customer-day groups in our sample based on the

---

groups) and 97 commercial customer groups (unique combinations of 5 baseload, 12 total usage, and 14 load cluster groups).

[6] We also will be exploring the value of using the cluster analysis for the residential customer as this research progresses, but this had not been completed at the time of this paper.

[7] An alternative model using 75 degrees Fahrenheit to define cooling days was explored in the residential analysis but did not have a significant effect on the estimation results.

residential or commercial group, weather group (CDD and/or HDD) based on the average daily weather value described above, and day type group (weekday versus weekend). Each participant remains assigned to just one customer group, but because temperature and day type changes day-to-day, each customer has customer-days that are assigned to many different bins.

Second Stage – Random Coefficients Model. For the next stage, 70 percent of the sample is randomly selected for use in the regression model to develop predicted hourly load shapes for each home-day bin using pre-period consumption data. The remaining 30 percent of the data is set aside as a holdout sample to test the performance of the predicted load shape. In this way, the predictive power of the model is tested against data that were not used to develop the model. For the residential and commercial QI datasets, the model is able to estimate load shapes within 1 percent accuracy for the holdout samples in both cases.

For the residential model, the average hourly kWh value is computed for the homes and businesses in each customer-day bin selected for modeling. These average hourly values of kWh represent the average load shape for each customer-day bin in the final regression model. For large datasets like the annual residential QI model, which has thousands of observations in a single bin, this approach cuts down on processing time without introducing bias for the resulting coefficients. If processing time is not a concern, all observations can be included in the model, as we did with the commercial model.

This modeling approach is used as it allows for the daily load shape (i.e., hourly kWh usage) of each customer-day bin to be estimated with a simple linear regression while accounting for covariance with other customer-day bin load shapes. Unlike a typical fixed effects regression, which produces a single set of coefficients and customer-specific constants, the random coefficients model produces a vector of regression coefficients for each home-day bin.

Two different regression specifications are explored for this stage. For the residential QI customers, the following model is used:

$$kW\_Hr_{i,t} = \sum_{j=1}^{5} \beta_{j,i}(ChangeH_{i,t}) + \sum_{k=1}^{5} \beta_{k,i}(ChangeH_{i,t} * H_{i,t}) + \varepsilon_{i,t}$$

Where:

$kW\_Hr_{i,t}$ = Mean kW consumption for homes in bin i during hour t.

$ChangeH_{i,t}$ = An array of dummy variables (0,1) representing hourly changepoints, taking a value of 1 if an hourly observation falls between two changepoints. In our final model, we use the changepoints 5am, 8am, 3pm, 6pm, 8pm, and midnight.

$ChangeH_{i,t} * H_{i,t}$ = An array of variables that interact the dummy changepoint variables with the hour of the day.

$\beta_{j,i}, \beta_{k,i}$ = Coefficients estimated in the model for homes in bin i.

$\varepsilon_{i,t}$ = Random error, assumed normally distributed.

For the commercial QI customers, a simpler specification is used that does not incorporate change points but instead has a single dummy variable for each hour of the day:[8]

$$kWh_{i,t} = \beta_{0i}H00_{i,t} + \beta_{1i}H01_{i,t} + \beta_{2i}H02_{i,t} + \beta_{3i}H03_{i,t} + ... + \beta_{23i}H23_{i,t} + \varepsilon_{i,t}$$

Where:

$kWh_{i,t}$ = Energy consumption, for customers in bin i during time interval t

$H00, H01...$ = An array of dummy variables (0,1) representing the hour of the day

$\beta_{0i}, \beta_{1i}...$ = Coefficients estimated by the model, for customers in bin i

ε = Random error, assumed normally distributed

For both the commercial and residential models, the coefficient estimates from the random coefficients model in each bin are then used to estimate consumption and eventually energy savings, as explained below in the third and final stage of the modeling process.

Third Stage – Savings Estimation. The final modeling stage requires that load shapes be calculated for the post-period using the results of the linear regression model. To accomplish this, the post-period data are subjected to the same binning process that is used with the pre-period data. Each individual home or business remains in the same customer energy usage group that it was assigned to in the pre-period, which helps isolate the effect of the program intervention occurring in the post-period by holding the expected general usage constant throughout the analysis period. Next, each day is assigned to a weather group (by CDD and/or HDD) and day type group (weekday or weekend).

After assigning each customer-day in the post-period to a customer-day bin, the predicted hourly pre-period kWh values for each customer-day bin are combined with the regression model results. This process results in a consumption estimate for each hour of the post-period in the hypothetical scenario where the customer did not participate in the program.

Once the forecasts of post-period usage are created (based on the pre-period consumption model and post-period weather data), they are compared with the actual post-period hourly kWh values. This is essentially comparing predicted consumption (which assumes that they did not participate in the program) to actual post-period consumption on days with the same weather conditions and day types. When actual post-period consumption falls below the predicted hourly kWh, this indicates energy savings during that hour attributable to the program. In essence, the estimated program savings is the difference between the predicted post-period hourly kWh and the actual post-period hourly kWh. This calculation could also be generalized to use long-term average weather data (rather than the specific post-period data used here) to create load shapes and subsequently energy savings estimates that assume more typical weather conditions.

Note that this process attributes the entire difference in actual versus predicted usage to the program intervention, which may or may not be appropriate depending on the program or market context. It also assumes that the existing site conditions are an appropriate baseline—any adjustments

[8] We have not yet re-run the residential models using the simpler regression specification to compare the results. This will be done as part of the second phase of the AMICS research. Given the accuracy achieved already with the residential model, we do not expect that a change to the simpler regression specification will significantly improve the prediction results.

to account for a standard practice or market baseline would need to be made outside the AMICS model.[9]

Analysis Results

The results of the AMICS model are presented here for both the residential and commercial sector applications. For the residential QI program, Figure 2 presents the annual impact results from the AMICS model. This graph compares the pre-period predicted load shape (red line) with the post-period actual load shape (blue), averaged across all households. Whenever the post-period load shape falls below the pre-period load shape, this indicates that savings were realized during that hour (green bars). The AMICS modeling approach results in approximately 7 percent annual savings attributable to the HVAC installed through the SCE QI Program.[10] Note also that this approach finds that the majority of savings is realized during the later part of the day, including during the peak hour periods of between 2 p.m. and 8 p.m., highlighted in yellow. The 95 percent confidence interval is shown for each estimate, and the error of the hourly predictions is greatest during the late afternoon and early evening, and smallest during the early hours of the morning.

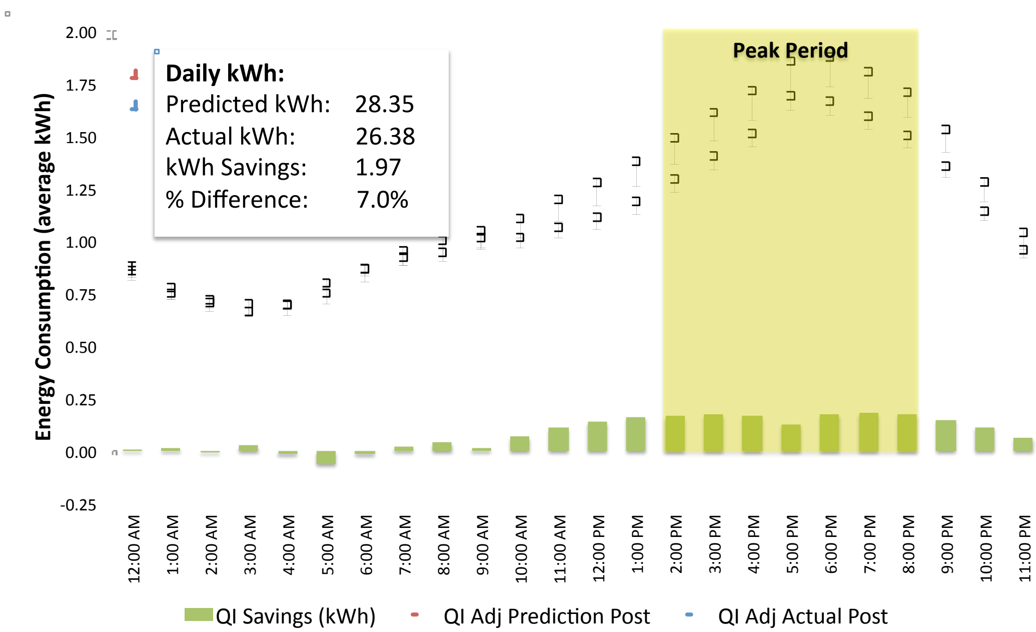

Figure 2. SCE QI overall annual post-period model, includes all months and day types

Figure 3 shows the model's predicted savings by season.[11] Note that this is an important feature of the AMICS modeling approach. As shown in the graph, most of the residential QI Program savings

[9] Both these issues are also present with the traditional billing regression methods and are not unique to the AMICS model.

[10] It was not possible to determine how much of the estimated savings come from the quality installation practices versus the new HVAC equipment from the data available for this study. In order to separate these impacts, the model would need to include a control group sample of customers who replaced their HVAC system but did not use a QI program contractor for the installation.

[11] Seasons are defined as: summer (July-September), fall (October-November), winter (December-February), spring (March-June).

occurred in the summer, which had an average daily savings of 5.3 kWh or 12.8 percent. Fall and spring had the next highest savings with 1.8 kWh and 1.2 kWh respectively, corresponding to 7.7 percent and 4.7 percent of the average daily kWh usage. Despite the variation in load shapes across seasons, the AMICS model is able to produce very accurate estimates in a variety of conditions. This ability for a single modeling process to match automatically a range of different load shapes is a key benefit of the AMICS approach.

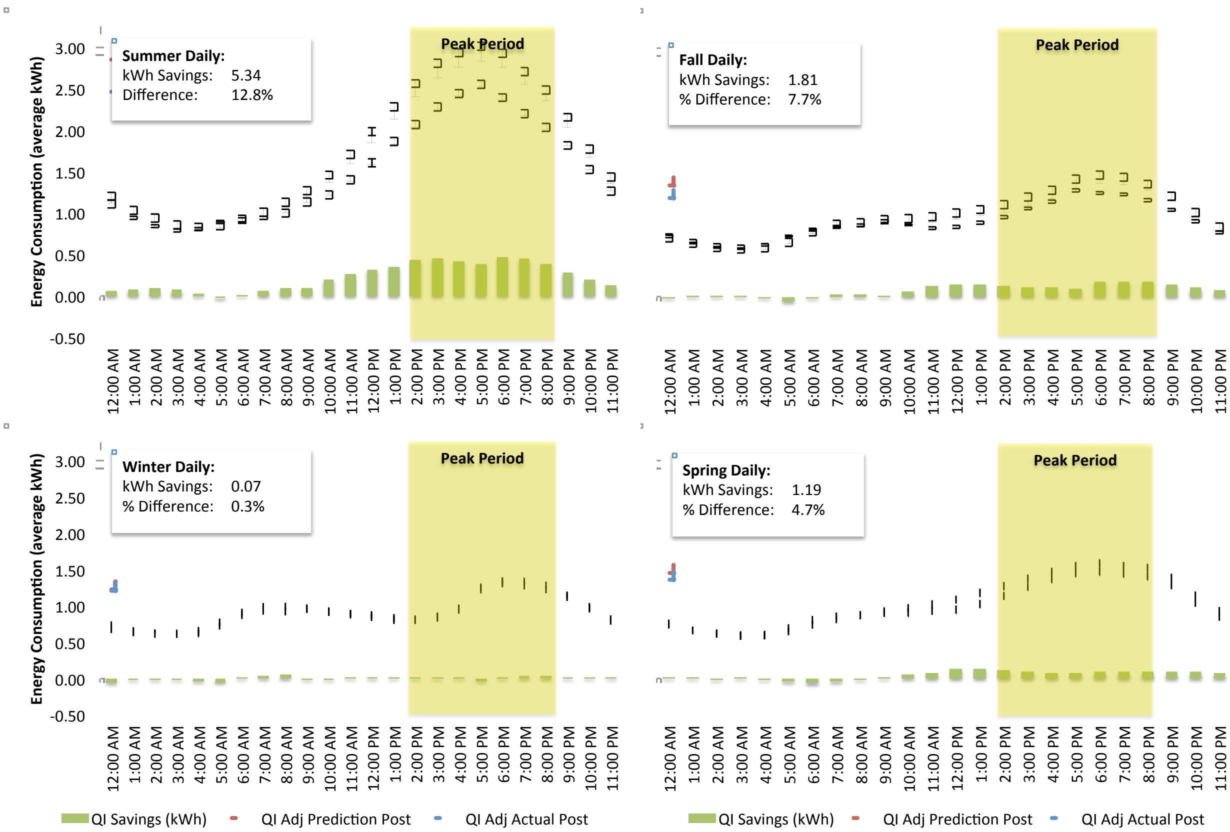

Figure 3. SCE residential QI annual model results by season

The AMICS model also showed promising results for the commercial QI program. When the model was tested against a holdout sample of customers, the predicted average daily load shape is within 1 percent of the actual usage for those customers that were not used to estimate the original model (Figure 4). As mentioned previously, 31 of the very largest commercial customers (i.e., those above 4,000 kWh average daily usage) are dropped from the analysis sample to improve the model prediction accuracy.

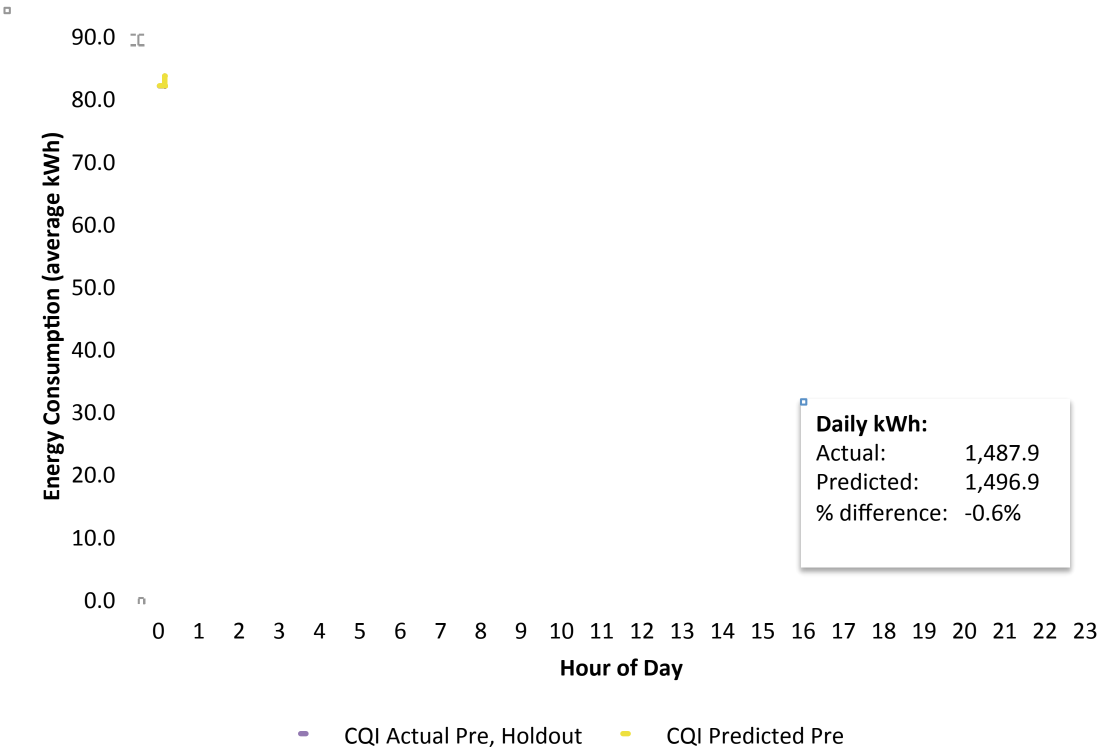

Figure 4. SCE residential QI annual model results by season

Figure 5 presents the average annual impacts estimates for the commercial QI group. For each hour of the day, the actual energy use was above the predicted value, resulting in no estimated savings for this program. For these particular customers, energy consumption increased in the post period, due to non-weather related factors that were not accounted for in the model. Future modeling work will explore alternative binning structures based on business type and customer size that may help incorporate some of these external factors into the model.

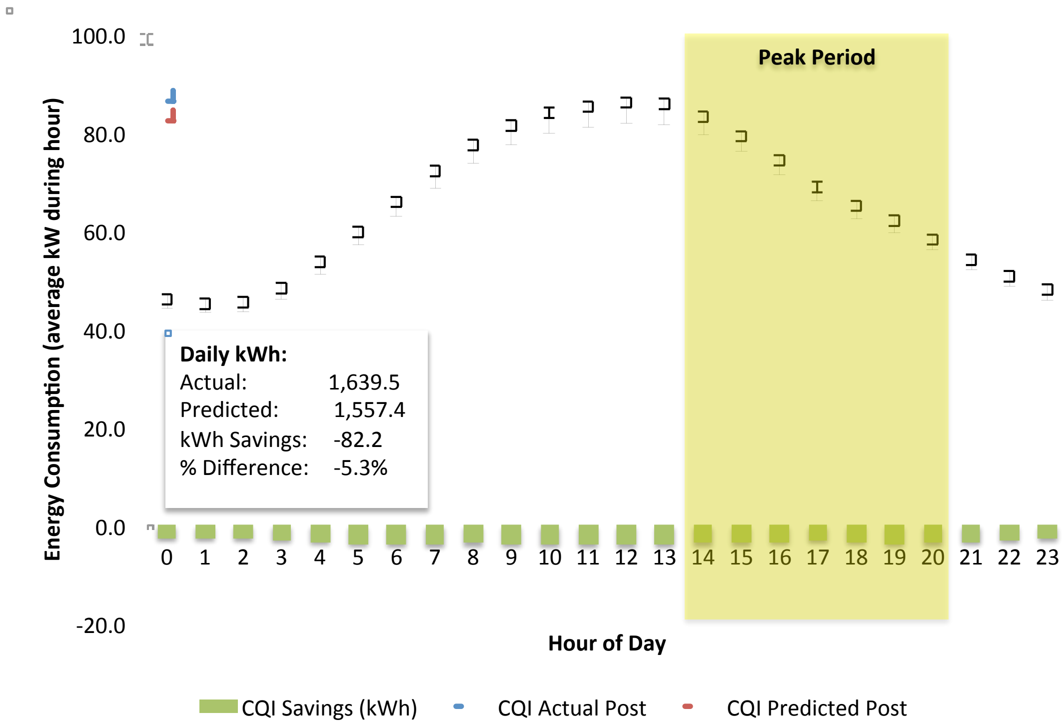

Figure 5. SCE residential QI annual model results by season

Summary and Conclusions

The results of this exploratory research demonstrate enormous potential for the AMICS model and represent a significant and positive departure from current approaches to analyzing AMI data and estimating program impacts. While analytically and conceptually more sophisticated than the traditional billing regression model, the additional complexity of the AMICS model is necessary to take full advantage of AMI data. As utilities continue to migrate their customers to interval meters, we believe it is necessary that evaluators embrace methods of analysis that fully exploit the abundant information contained in AMI data. The multi-stage approach also provides an efficient means for processing the high volume of AMI data and organizing it for use in the billing regression model in a manner that controls for important sources of variation.

The test application of the AMICS model using data from both the residential and commercial QI Programs provided some encouraging results. First, the AMICS modeling approach was able to produce very accurate predictions of load shapes, generally within 1 percent of actual use for a holdout sample for both commercial and residential customers. Second, the model was able to produce savings estimates that were consistent with expectations for the residential program. Finally, the predicted load shapes and impact estimates for the residential program at the daily and seasonal level were also consistent with expectations and, in the case of the seasonal models, were able to account for significant differences in load profiles across periods in the residential sector. While the model was able generally to match the load shapes for the commercial customers, it was not successful in producing any impact estimates.

Perhaps the greatest benefit of the AMICS model is the ability to estimate impacts for different customer types in an efficient manner. This is accomplished through an automated categorization process (i.e., bin assignments) that controls for much of the variation across customers and weather conditions. If only annual impact estimates are needed, then the fixed effects model is likely sufficient for estimating savings. However, a single annual savings number does not take full advantage of the information provided by the AMI data. The AMICS modeling approach, in contrast, uses the AMI information to create customer subcategories that help control for significant amounts of variation and ultimately allows for accurate load shape predictions. Once this process is completed, it provides a flexible method that enables different load shapes (and subsequently impact estimates) to be developed easily for a wide range of different time periods.

Future work will focus on refining the AMICS model for commercial customers. While the largest commercial customers may still require individual AMI data analysis, the performance of the AMICS model with the other commercial customers is promising as the model is able to estimate very closely the load shapes for a holdout sample of customers. Additional work is underway to create additional refinements of the model based on customer size and business type. A separate analysis is also underway that will utilize HVAC end-use metered data combined with the whole-building AMI data in the AMICS model to estimate both total consumption and the HVAC portion of total load.

References

Agnew, K. and M. Goldberg. 2013. "Chapter 8: Whole-Building Retrofit with Consumption Data Analysis Evaluation Protocol." *The Uniform Methods Project: Methods for Determining Energy Efficiency Savings for Specific Measures (UMP)*.

California Public Utilities Commission (CPUC). 2006. *The California Evaluation Framework.*

California Public Utilities Commission (CPUC). 2008. *Load Impact Estimation for Demand Response: Protocols and Regulatory Guidance.*

Evergreen Economics. 2016. *AMI Billing Regression Study Final Report.* Prepared for Southern California Edison.

Granderson, J, PN Price, D. Jump, N. Addy, and M. Sohn. 2015. "Automated Measurement and Verification: Performance of Public Domain Whole-Building Electric Baseline Models." *Applied Energy* 144: 106-113.

Granderson, J., S. Touzani, C. Custodio, S. Fernandes, M. Sohn, and D. Jump. 2015. *Assessment of Automated Measurement and Verification (M&V) Methods*. Lawrence Berkeley National Laboratory, LBNL#-187225.

Helvoigt, T., S. Grover, J. Cornwell, S. Monohon. 2016. "A Smart Approach to Analyzing Smart Meter Data." Presented at the ACEEE Summer Study, Asilomar, CA.

Hsiao, C., D. Mountain, M.W. Chan, K.Y. Tsui. 1989. "Modeling Ontario Regional Electricity System Demand Using a Mixed Fixed and Random Coefficients Approach." In *Regional Science and Urban Economics* Volume 19, Issue 4: 565-587.

Nexant. 2014. *2013 Load Impact Evaluation for Pacific Gas and Electric Company's SmartAC Program.* Prepared for Pacific Gas and Electric Company.

Snijders, Tom and R. J. Bosker. 2012. *Multilevel Analysis: An Introduction to Basic and Advanced Multilevel Modeling*. Thousand Oaks: Sage Publications.