

2006-07 MFA WEATHERIZATION PROGRAM GAS BILLING ANALYSIS

Task 5.1: Draft Final Report

March 17, 2009

Prepared for:



Public Service Company of New Mexico

Alvarado Square, MS-0506

Albuquerque, NM 87158-0506

Attn: Steve Bean

Energy Efficiency manager

Prepared by:



KEMA, Inc.

466 Geary, 4th Floor

San Francisco, CA 94102

Table of Contents

TASK 3: DATA ANALYSIS	1
TASK 3.1: BILLING DATA RESOLUTION	ERROR! BOOKMARK NOT DEFINED.
TASK 3.2: WEATHER NORMALIZATION	2
TASK 3.3 BILLING ANALYSIS	5

Task 3: Billing Data Analysis

Billing analyses were conducted to estimate weather normalized annual gas consumption (NAC) for each home on the gas billing file, both before and after implementation of the weatherization on the premise. This was done so that an evaluation of the weatherization installations could be made that was not clouded by the weather corresponding to the span of time matching the gas bills. A Statistically-Adjusted Engineering (SAE) model was then fit relating the weather-normalized consumption savings to the predicted annual consumption savings from tracking data to produce a ratio representing the overall proportion of the predicted savings that was observed in the sample of households in this analysis. This report discusses the methodology used to calculate these estimates and reports the results for homes in the study.

Summary of Results

The observed savings in annual gas consumption following weatherization did not reflect the savings predicted from tracking data. In all but a couple of premises annual consumption was predicted to drop following weatherization program installation. The observed annual consumption following the program installation, however, only dropped slightly overall over the year prior to installation and in close to half of the premises the consumption actually increased. This observed drop was in weather-normalized annual consumption.

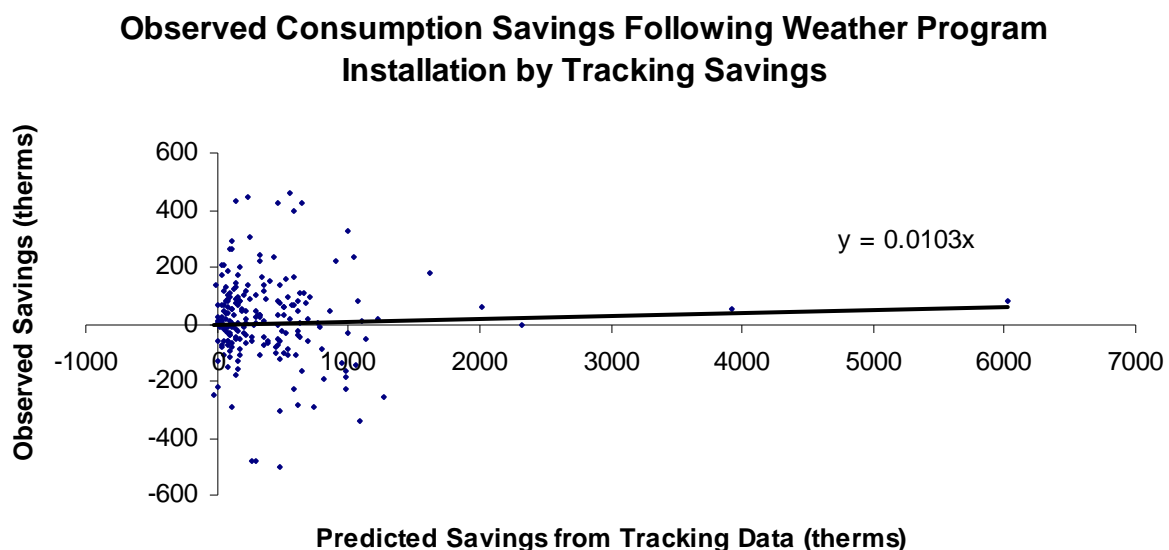


Figure 1: Drop in normalized annual consumption (NAC) following weather program installation by predicted savings from tracking data

The slightly increasing regression trend line in figure 1 indicates very weakly positive overall correlation between the tracking savings and the observed savings. Note the prevalence of

premises with negative observed savings, meaning an increase in NAC after program installation.

Total Tracking Savings from sample (therms)	Total NAC Savings from sample (therms)	Ratio Estimate of Total NAC Savings to Total Tracking Savings	Total tracking savings for program (therms)	Total Program Savings (therms)
82,605	3327	0.04	192,734	7,709

Table 1: Summary of program savings. SAE savings is total savings multiplied by SAE ratio estimate of savings.

The SAE model fit with the final 207 premise NAC differences from before to after weatherization and the tracking savings estimates produced an estimate of the ratio of observed to predicted savings of about 1 percent. This SAE model was not fit with cases from a control group of non-program participants, as the SAE model requires for best performance, due to the lack of availability of non-program participant bills. Had this control group data been available, the model could have accounted for non-weather-related trends existing in the class of natural gas users represented by this sample and the estimated savings could potentially have been substantially more than was found in this limited analysis. The straight ratio of total observed NAC savings to the total tracking savings was about 4 percent. The SAE model had to be fit without an intercept, which leads to problems interpreting parameter estimates. Because of this, it is recommended that the 4 percent savings ratio and the 3,327 therms total sample savings and 7,709 total program savings be used.

Task 3.1: Billing Data Resolution

The goal for the billing data resolution was to maximize the utility of the gas billing data so that the conclusions from the analysis tell the complete and accurate story of the effect of weatherization hidden in the data. There were a number of steps in the process of transforming the delivered billing data set into a form that could be input to the PRISM regression model and making sure that data was of sufficient quality to use.

The shape of the data sets (separate data sets were produced for the period before and after the weatherization installation) input to the model needed to be such that rows contain premise ID, weatherization installation date for that premise, billing period number for that premise, number of days for that billing period, average daily consumption for that billing period, and the average of the high and low temperature for the day within the billing period. There was to be a separate row of data for each ID/billing period/day within billing period combination. This was necessary so that heating degree-days could be computed for many different reference

temperatures. The billing data file contained consumed therms by premise and billing period as well as the read dates. It did not, however, include the weatherization dates or the temperature data. To resolve this, information from other files was linked. In many cases what appeared to be a gas bill was in fact an adjustment to a prior bill. Before any analysis was done, these adjustments were aggregated so that each record in the analysis would represent an actual billing period. This also necessitated adjusting the billing period days to match the aggregated bills.

The billing file contained a match variable to a file of temperature data by weather station. It was determined that weather stations located in airports serving three New Mexico cities – Albuquerque, Santa Fe, and Las Cruces – were sufficiently representative of the climates associated with sample for premises to be assigned to the temperature data from one of these stations. Some daily temperature averages were missing. In those cases the average from the previous day at that station was imputed for the missing values.

Linking the weatherization program installation dates to the billing data was a far more difficult process than for the weather stations because there was not a matching premise ID on the tracking data files containing the installation dates. A fuzzy matching technique was employed using a combination of the street address, last name of the resident, and the zip code. This was successful at linking in over half the cases, but significant follow-up work was required as non-standard ways of entering names and addresses, incorrect zip codes, misspellings, discrepancies in resident name associated with the premise, and other matching issues made perfect matching by algorithm impossible. Comments were recorded for each correction made to the billing file. There were 257 unique premises on the billing file for which an attempt to assign an installation date was made.

Some cases could not be matched at all, the month/year of program installation was missing, or there was resident turnover making for an insufficient span on one or both sides of the installation date. In these cases the premise was dropped from analysis. Thirteen such cases were dropped, leaving 244 premises.

Gas bills with unusual average daily consumption and bills with unusual numbers of billing days were examined. Although it was proposed that only records with a billing cycle of 25 to 35 days would be accepted, a large number of 60-day bills were found as were many 90-day bills. Instead of dropping these records, they were examined for legitimacy in terms of their average consumption over the coverage period. Highly suspect bills were dropped from analysis along with other bills before or after that billing period depending on what side of the installation date they fell.

The specified bill coverage span on each side of the program installation date was one year. Most of the sets of bills for each premise started two or more years before the program installation date, which occurred sometime in 2006 or 2007. The same cannot be said about the sets of bills following installation as no bill was dated later than October, 2008. Cases were

dropped on both sides of the due to insufficient coverage, but most of them were due to not having the full year of bills in the after (weatherization installation) set – 212 of the original 257 premises remained after this step. The means of consumed therms of gas aggregated in the before and after groups, respectively, were 687 and 719. In both groups the standard deviation for the mean was approximately 300.

Task 3.2: Weather Normalization

A refined Princeton Scorekeeping Model (PRISM) algorithm was used to produce the normalized annual gas consumption (NAC) estimates for the two sets of data. Comparing two different years of annual consumption totals for a set of premises without normalization ignores the important latent factor of temperature. Space heating makes up a bigger share of consumption than any other end-use of natural gas and is an obvious response to outside temperature.

The refined PRISM algorithm used is essentially a series of linear regression models fit with different covariates for each premise and the best of this series of models stored for each of the premises. For each day in the billing span for each premise and for each step in the 45-75 degree-day reference temperature range, the number of degree-days was computed and then averaged for each of the billing periods. The average daily therms consumed was regressed on the average number of degree-days per billing period and the square of the average degree-days per billing period per the proposed PRISM model with adjusted R-square determining model selection within each reference degree-day temperature and mean squared error determining the best reference degree-day temperature to use across the 45-75 degree range. It was determined that although the squared heating degree-day term worked well for individual billing period averages, it led to negative values for predicted consumption so this squared term was eliminated and regression was repeated with just the intercept and heating degree-day covariate.

The resulting parameter estimates and best degree-day reference temperature for each premise as computed in the PRISM model were stored. To produce the NAC for each premise, expected annual weather was set with the degree-day parameter instead of the actual temperature data corresponding to the billing data used to fit the model. The key is that for both the before and after program installation models fit for each of the premises, the same temperature data is used.

This was done as follows: for each premise, the degree-days – with reference temperature determined in the PRISM model – were computed for every day between 1/1/2005 and 12/31/2007. The resulting degree-days were then summed and divided by three to get an average annual degree-day value over the three years.

Because multiple weather stations and reference temperatures were used, the number of heating degree-days in the three-year average varied across the set of premises in the analysis with the mean about 3350 degree-days and standard deviation of about 1370.

Using the set of three-year average heating degree-days, the normalized annual consumption (NAC) was computed for both the before and after groups as follows:

$$NAC_{i,j} = 365 * b_{o(i,j)} + b_{1(i,j)} * DD(\tau)_{(i,j)}$$

Equation 1: Normalized Annual Consumption

where (i,j) represent premise i (of 212 premises) and group j (of the two groups – before and after), b_o is the estimated base gas consumption and b_1 is the estimated parameter that tells how consumption increases linearly with $DD(\tau)$, the annual heating degree-day norm with respect to reference temperature τ .

The mean resulting NAC for the before group was 681 and 667 for the after group, giving an average mean difference of 14 therms. Note that the actual therms for the before and after group were 687 and 719, respectively – so the NAC values suggest that despite the actual gas consumption going up after the program installation, after controlling for weather it went down on average.

Period	Actual Gas Consumption*	Normalized Annual Consumption (NAC)
Before Weatherization Program Installation	687	681
After Weatherization Program Installation	719	667

*Actual Gas Consumption corresponds with a span of at least one year – mostly between 365 and 400 days – while NAC is for 365 days.

Table 2: Therms of gas consumed by consumption type and period

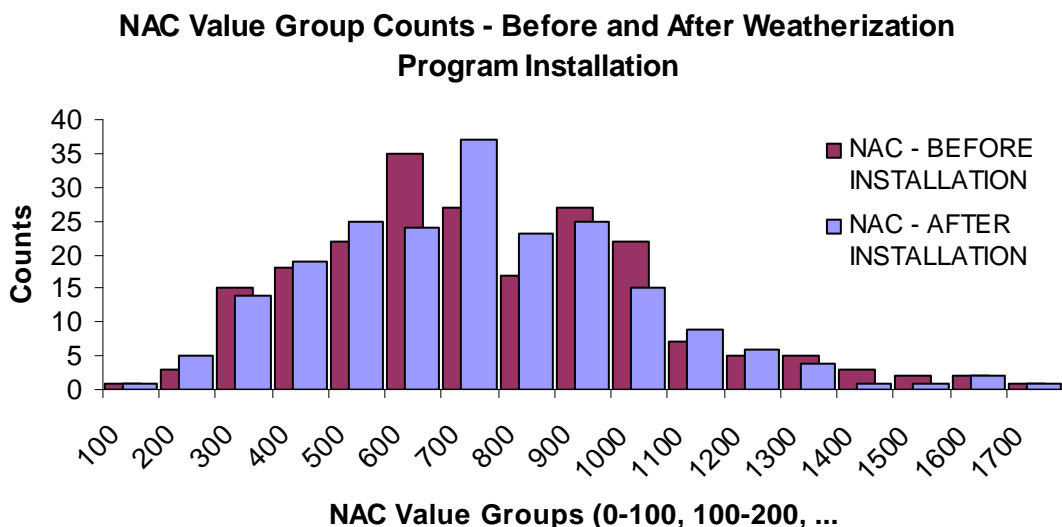


Figure 2: Histograms of NAC before and after weatherization

The estimated efficiency, however, was highly variable. While on average the normalized consumption decreased, in a large number of cases an increase was observed. In 94 of the 212 cases there was an increase in NAC.

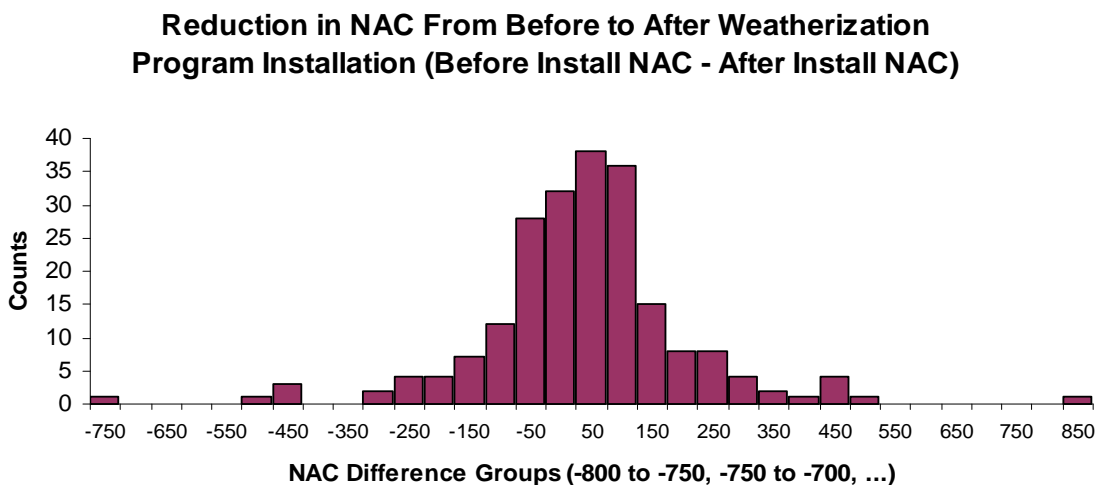


Figure 3: Histogram of reductions in NAC after weatherization

Task 3.3: Billing Analysis

The expectation was that weatherization would make more efficient use of natural gas for space heating in households, leading to a drop in gas consumption and expenditures. Annual savings estimates from tracking data were hoped to have a positive, approximately linear

relationship with the NAC differences. Using the NAC from the before and after groups, their tracking savings estimates, and before and after NAC from a control group of non-participant low income households, a statistically adjusted engineering (SAE) was to be fit as follows:

$$NAC_{diff(i,k)} = \beta_o + \beta_1 * S_{(i,k)} + \varepsilon_{(i,k)}$$

Equation 2: SAE regression model relating tracking estimates of consumption savings to observed savings

Where:

$NAC_{diff(i,k)}$ is the NAC for the year before weatherization minus the NAC for the year after weatherization for premise i in group k (participants, non-participants).

β_o is the mean observed difference in NAC for both participants and non-participants

β_1 is the estimated proportion of the tracking savings estimate that was realized for the observed sample

$S_{(i,k)}$ is the tracking savings estimate of annual consumption for premise i in group k (for the non-participant group $S=0$)

$\varepsilon_{(i,k)}$ is the error in predicting the observed drop in consumption for participant i in group k.

This model unfortunately could not be fit since non-participant bills were not available. The non-participant bills are required in order to attempt to control for other factors affecting consumption trends over the time periods of this analysis other than weather, for which heating degree-days was used as a proxy. If weather was the only factor driving consumption, the non-participant billing data would not be necessary.

Using only the participant data an SAE was fit as in equation 2 but without an intercept.

$$NAC_{diff(i)} = \beta_1 * S_{(i)} + \varepsilon_{(i)}$$

Equation 3: SAE regression model relating tracking estimates of savings to observed savings for program participants only.

The terms in equation 3 are defined as in equation 2. The model was fit on 207 cases after excluding three outliers and two cases without a tracking savings estimate. The estimate for β_1 was about 1 percent. Because this regression model was not fit with an intercept, it did not match the ratio of total observed drop in NAC to total estimated tracking savings, which is about 4 percent. Regression models forced to be fit with no intercept do not lead to the

familiar interpretations for parameter estimates and because of this it is recommended that the 4% realized savings ratio estimate be used instead of the SAE ratio estimate given.

Total Sample Tracking Savings (therms)	Total NAC Sample Savings (therms)	Ratio Estimate of Total NAC Savings to Total Tracking Savings	Total Tracking Savings for Program (therms)	Total Program Savings (therms)
82,605	3,327	0.04	192,734	7,709

As is clear from figure 1 on page 1, a positive linear relationship, even in an approximate sense, was not observed between the NAC savings and the projected annual savings. While close to half of the observed NAC values after weatherization were consumption increases, only one tracking savings estimate was negative. Some projected consumption savings values were much greater than the pre-weatherization NAC values – the observation on the far right had a savings estimate of 6020 therms but an observed drop of only 78 therms.

If, however, non-participant billing data from similar low-income New Mexico households was available to make up a control group, there would have been *potential* to see a different story after fitting the SAE model. Suppose there was an upward trend in consumption of natural gas across the time span in the analysis for non-participants. The SAE model could in that case result in a negative estimate of β_0 with a positive estimated β_1 , possibly suggesting the weatherization did lead to significant savings. Again, this is hypothetical, but it must be conceded that whatever the real story is, it will not be discovered without control group data input to the SAE model along with the program participant data.

