A Pacific Northwest Efficient Furnace Program Impact Evaluation

Ken Agnew, KEMA Inc., Madison, WI
Mimi Goldberg, KEMA Inc. Madison WI
Bobbi Wilhelm, Puget Sound Energy, Bellevue WA

ABSTRACT

The Pacific Northwest is known as a hotbed of energy efficiency activity. Since 2003, Puget Sound Energy has provided rebates for energy star qualified gas furnaces, rated with an AFUE of .90 and above, to both retrofit and new construction customers. This paper summarizes an evaluation that was conducted to confirm existing program estimates of savings.

The evaluation used a pooled time series cross-section approach to model energy savings. The approach modeled multiple households over multiple time periods, controlling for household level and time period-specific effects. Typically, billing analyses utilize bill histories for only program participants. This evaluation took advantage of monthly billing data from both program participants and non-participant furnace installers, as well as survey data for both, to obtain a direct estimate of net savings from the installation of a high efficiency furnace.

The impact evaluation generated a number of interesting results. The evaluation confirmed existing estimates of gross savings, and it also found effectively no decrease in non-participants energy usage with the installation of a new furnace, regardless of the unit efficiency. This finding indicates that this billing analysis was not fully successful at directly estimating net saving for the program. Instead secondary estimates of free ridership were used to provide the net result. We discuss this result in the context of self-selection bias in models including non-participants.

Introduction

The Puget Sound Energy (PSE) Residential Energy Efficient Natural Gas Furnace Program has promoted ENERGY STAR© qualified gas furnaces with rebates since 2004. In 2008, PSE hired KEMA to verify legacy estimates of program impacts. PSE opted for a billing analysis approach that provides the best possible results from an analysis based on billing records. The inclusion of non-participant furnace installers in a pooled time-series cross-section billing regression offers the potential for impact estimates that control for both kinds of naturally occurring savings: standard installation savings and free rider savings. The approach successfully accounts for standard installation savings, but due to specific self-selection issues, does not produce an impact estimate that is net of free ridership. Two separate free ridership estimates support this conclusion. The result is a sound estimate of both gross and net savings from the installation of an efficient gas furnace in PSE territory and a better understanding of the strengths and weakness of this billing analysis approach.

Program Background

Puget Sound Energy's (PSE) Residential Energy Efficient Natural Gas Furnace Program promotes ENERGY STAR© qualified gas furnaces with rebates of up to $250. The program is designed to serve retrofit and new construction customers using the furnaces for their primary heat source. While the program has been in effect for many years, this evaluation is based on the 2005 and 2006 program years. Since its inception, the program has used a savings estimate of 90 therms. PSE was interested in an updated estimate of savings based on actual program data. PSE opted for a billing analysis approach to take advantage of their extensive archive of billing records. KEMA proposed a billing analysis approach designed to provide the most comprehensive results available from a billing analysis.

Billing Analysis Background

In practice, the term "billing analysis" represents a wide range of evaluation approaches. Any evaluation of monthly customer consumption data can fall under this term. Within this larger scope of "billing analysis", there are two general approaches that utilize regression-based methods to account for the weather-correlated nature of most residential consumption data. In a regression context, both approaches utilize heating and cooling degree days (HDD and CDD) as explanatory variables to explain energy consumption. Electric consumption billing regressions generally include both HDD and CDD to address consumption patterns driven by both hotter and cooler weather. Because gas consumption is generally limited to heating related equipment, gas billing regressions generally include only HDD.

These billing analysis regression methods allow for weather normalization. The effects of specific weather, unique to the time and place, are used to estimate the model. The resulting regression parameters are fit to data from a chosen representative weather series (frequently, weather "normals") to provide results under representative weather conditions. PRISM, a widely recognized billing analysis program, performs weather normalization with the added feature that optimal degree day bases are chosen for each site resulting in models tailored to the unique usage characteristics of the site. In the process of normalizing consumption with respect to weather, these regressions disaggregate consumption into its constituent heating, cooling and base load parts.

Two Step vs. Pooled Approach

One approach, referred to here as the two-step approach, starts with separate PRISM-like analyses of participants' consumption in the pre- and post- installations periods. The difference between pre- and post-installation consumption, as determined by the model, represents an estimate of a site's program-related change in consumption. In a second step, these changes in consumption are regressed on site characteristics to further disaggregate the change in heating, cooling or base consumption. Site characteristics could include details gleaned from program tracking (different mixes of EE installations) or household specific details derived from survey responses (fuel-switching, increase in conditioned square footage).

The second billing analysis approach, referred to here as the pooled approach, models billing data in a single pooled time-series cross-section model. Both weather effects (HDD and CDD and site characteristics) are included in a single un-balanced panel model that includes multiple months of billing data for all the included customers. This evaluation pursued billing analysis with a pooled model and we discuss the specific advantages of the pooled approach here.

Participant Only Billing Analysis Approach

The simplest form of pooled billing analysis includes participants only. Each participant's consumption data reflects normal, ongoing usage patterns except at the time of installation. With regards to the pre-post installation change in consumption, each participant serves as a baseline for itself. In addition, those participants not in change mode inform a series of monthly dummy coefficients that control for systematic effects such as changes in consumption due to the economy, weather, etc. The model provides an estimate of average change in consumption due to the program-installation. An important advantage of this approach is that all participants with suitable billing data can be included in the regression. The inclusion of the full population bypasses the issues of sample bias. The final results from such a model are an estimate of participant change in consumption including all naturally occurring savings.

The limitations of the participant only billing model are due to the limited information that is known about each participant beyond what is captured in the billing data and the most basic tracking data. Not infrequently, the only complete and trustworthy information in tracking data are the installation dates. This is the minimum necessary to perform the model. Many tracking systems attempt to capture data on house size and vintage but this data is rarely complete or trustworthy. The same can be said about essential information on confounding situations like fuel switching, additions, etc. Without this information, the pre-post change in consumption includes, for instance, the increase in consumption associated with switching to gas heat from electric. The end goal of a billing analysis is the impact of an average, same-fuel installation.

Finally, in addition to program-related savings, the pre-post change in consumption includes what is referred to as "naturally occurring" savings. Naturally occurring savings is what would have happened in absence of the program. One part of naturally occurring savings is the savings that results from the installation of new furnaces, furnaces that meet current, improved building standards. The other part of naturally occurring savings is free drivership; that is, the installation of efficient products without a rebate. The participant-only model is ill-equipped to decompose the pre-post change in consumption into naturally occurring savings and, what remains, net savings.

Adjusting with survey results

The addition of a survey goes one step toward addressing the limitations of the participant-only billing analysis. With the use of survey data, one can make aggregate adjustments to the billing analysis estimate of change in consumption to control for non-standard installations such as fuel-switching and additional floor space. The more nuanced the change in consumption, the more complicated the adjustment. A participant who switches fuels at the installation of the new furnace goes from an effective heating consumption of zero to post-installation levels of consumption consistent with heating. Decomposing the pre-post change that includes these two kinds of participants is relatively straightforward. The participant who adds an addition to the house and increases the capacity of the furnace presents a greater challenge.

Survey data alone is still not enough to identify naturally occurring savings. There are ways to do this without including non-participants, but they require more extensive tracking data that captures unit efficiency and capacity.

Including Non-participants

We can address issues related to both non-standard installations and naturally occurring savings by including a sample of non-participant installers in the billing analysis. Unlike the previously discussed approaches where at best only a subset of participants have survey data, all sites entering into this model have associated survey data. This allows us to control for non-standard installations directly in the model rather than with after-the-fact aggregate adjustments. With the use of participant and non-participant surveys, we can control for fuel-switching, increases in capacity, etc. In addition, the surveys provide other site-specific characteristics such as square footage and number of inhabitants which can improve the fit of the model.

More importantly, by including non-participant furnace installers in the billing analysis, we can identify "naturally occurring" savings. The naturally occurring savings that participants would have realized had they not take part in the program are represented by the extent to which non-participants lower their consumption in the post-installation period. By controlling for savings in the post-installation period common to both groups, naturally occurring savings are netted out of the participant-only post-installation savings.

The inclusion of non-participants has the potential to address both kinds of naturally occurring savings. With regards to the savings generated by installing a new, up-to-code furnace, non-participants should be representative of what participants would have done had they not installed an energy efficient furnace. Non-participants provide the baseline against which program participants' gross savings should be measured.

In theory, the model also controls for non-participant post-installation savings that go beyond installing standard code furnaces. To the extent that non-participants install units that are more efficient than standard code, those savings will also be captured in the non-participant post-installation savings. If non-participants are fully representative of what participants would have done in the absence of the program, then non-participants would be purchasing energy efficient furnaces at the same rate as participants. If this were the case, then the savings of free-riders would also be contained in the shared post-installation savings and thus be netted out of the participant, post-installation estimate of savings.

Self-Selection Bias and Free Ridership

The previous sections explain how a pooled billing analysis including both participant and non-participant billing records and site-specific survey data has the potential to directly estimate participant post-installation impacts relative to standard baseline and net of free riders. To achieve this end, the sample of non-participants bills must be representative of the participant bills, had the participants not participated in the program. In practice, this proves to be difficult to achieve. One source of difficulties is self-selection.

Customers do not opt into energy efficiency programs on a random basis. Instead, they self-select based on their individual characteristics and calculation of the value of participation. Controlling for observed characteristics in the billing model to the extent feasible may solve some of this problem. However, the decision to take part in a program is still correlated with a range of unobserved customer characteristics. Both parts of the pre-post change in consumption, naturally occurring savings and net savings, the target result in the analysis, are correlated with characteristics underlying the make-up of participant and non-participant samples (Goldberg and Train 1996).

Self-selection has a number of implications but in a billing analysis with a non-participant control group the primary concern comes back to getting the appropriate baseline. For example, customers with high consumption houses might be over-represented among program participants because they stand to reap greater savings (both naturally occurring and net). In comparison, the non-participant sample, if selected randomly, would include customers with relatively lower consumption. These non-participants would provide too low a baseline consumption to properly measure the full impact of the program on the participants. The result would be a downwardly biased estimate of program impacts. An upward bias is also possible to hypothesize; take the case of a participant sample full of committed green customers who enter the program already at the low end of the consumption spectrum.

Self-selection was important to this evaluation for two reasons. First, we attempted to include techniques that address for self-selection. Second, despite the efforts to address self-selection, there was strong evidence that self-selection was still affecting the results in important ways.

Free Ridership

It's important to clarify the relationship between self-selection and free-ridership. A customer population generally includes some efficient furnace installers who would install the energy efficient option without any kind of incentive or rebate. The savings generated by these folks are naturally occurring savings, beyond those realized with a standard installation. If they opt to take part in the program then they become free riders. A billing analysis including non-participants can control for free ridership among participants to the extent that efficient furnace installers are distributed similarly among non-participants and participants. To the extent that the distribution of non-participant efficient furnace installers is different from participant distribution, the billing analysis may not be able to properly control for participant free rider savings.

In simple terms, self-selection is in evidence to the extent that this group of efficient furnace installers is more prevalent among program participants than non-participants. EE programs give efficient furnace installers cash. Furthermore, they work hard to promote widespread awareness and lower the hurdles to participation so as to attract as many participants as possible. It is not unreasonable to expect efficient furnace installers to self-select into the program.

With this in mind, it is reasonable to imagine that the energy efficiency installment portion of naturally occurring savings may have a downward bias; that is, less than the full amount of free rider savings is being removed from the impact estimate. Fortunately, this particular aspect of self selection bias is bounded within the range from net to gross savings. If all efficient furnace installers opt into the program then only savings related to standard code installation are controlled for and the participant impact estimate completely includes the savings of free riders. The challenge presented by this kind of self-selection bias is knowing the degree to which free ridership has been accounted for.

Self selection may remain an issue in the more general sense. Regardless of whether efficient furnace installers select into the program, participants and non-participants may be different in other ways as a result of the participation decision. Employing the standard double Inverse Mill Ratio self-selection correction in the pooled billing model specification may address this issue (Green 1997, Maddala 1983). The inclusion of the IMR, however, is challenging. In addition, self-selection correction techniques are limited in their ability to correct for self-selection. There are indications, in fact, that the requirements may be so limiting as to undermine their practicality.

Methodology

Heating Consumption Billing Regression

The heating consumption billing regression is the simplest form of pooled time-series, cross-section analysis. Across multiple households each with multiple time periods, this specification provides a general estimate of the relationship between average daily therm consumption and heating degree days. We start with this specification for ease of explanation as we develop the final billing regression specification.

The equation is:

$$E_{im} = \mu_i + \lambda_m + \beta_H H_{im}(\tau_H) + \varepsilon_{im}$$

where

E_{im}	=	Therms used per day during month *m* for customer *i*;
μ_i	=	Premise-specific base consumption estimate for customer i;
λ_m	=	Month-specific time period effect for month m;
$H_{im}(\tau_H)$	=	Average heating degree-days at the heating base temperature τ_H during month m, based on daily average temperatures, for customer i's meter reading period;
β_H	=	Heating coefficients, determined by the regression;

τ_H	=	Heating degree-day base temperature, determined by choice of the optimal regression; and
ε_{im}	=	Regression residual.

In this equation, gas consumption is a function of a household-specific constant (baseload μ_i) and average daily HDD $H_{mi}(\tau_H)$. Monthly bill readings divided by the number of days in the billing period provide the daily therm consumption, represented by E_{im}. Degree days for the billing period are calculated by dividing the sum of daily HDD by the number of days in the billing period. Because we obtained monthly consumption data by bill period, and not all customers are on the same bill cycle, heating degree-days for a given month vary over customers i. Finally, the time period coefficients λ_m account for systematic effects correlated with the time period but not explained by the other variables.

The household-specific intercept and month-specific time period terms are referred to as fixed effects. The fixed effects parameter estimates provide household- and month-specific information in this model where multiple households and months are combined in a single model. This approach separates out household-specific, non-heating gas consumption that occurs across all time periods. Non-heating gas consumption can include water heat, cooking and other gas appliances. The fixed effect model also separates out time-specific changes in gas consumption that occur across all households. Because of the rolling nature of meter reading, the billing regression time-period effects are inherently fuzzy. They still, however, control for time-specific changes in gas consumption. These changes can be due to economic factors, natural disasters (i.e., earthquakes), or unusual weather.

We test the full model using heating degree days calculated across a range of degree day bases. Across model runs with the same variable combinations, the R^2 provides a simple way to find the optimal combination of heating and cooling degree day bases. This effectively estimates the average outdoor temperature at which heating or cooling begins among the included households.

Pre-Post Billing Regression with Non-Participants

The basic heating consumption billing regression is enhanced to include a change point (furnace installation) and a comparison group of non-participating furnace installers. The model estimates non-participant heating consumption and separates out a participant effect, a shared participant and non-participant installation effect and the participant-only installation effect. The simplest version of the pre-post model is:

$$E_{im} = \mu_i + \lambda_m + \beta_1 P_i * G_{im} + \beta_2 H_{im}(\tau_H) + \beta_3 H_{im}(\tau_H) * P_i + \beta_4 H_{im}(\tau_H) * G_{im} + \beta_5 H_{im}(\tau_H) * P_i * G_{im} + \varepsilon_{im}$$

where

P_i	=	An indicator variable equal to zero for non-participants and one for participants
G_{im}	=	An indicator variable equal to zero prior to installation of the furnace and one after the installation of the furnace
β_1-β_5	=	Estimated coefficients

Parameters β_2 through β_5 cover the range of combinations of the participant dummy and the post-installation dummy. They function as a group and in the context of the regression model, isolate the desired estimates. All additional variables that enter the model will be included with these four combinations.

The parameter β_2 is an estimate of pre-program heating consumption for non-participants. It measures the heating consumption correlated with HDD with no other interaction terms. The parameter β_3 measures the average difference in heating consumption across all time periods for participants relative to non-participants. With this parameter, the model accounts for differences in heating consumption across the two samples.

The parameter β_4 represents average change in heating consumption from the pre- to post-installation periods across both participants and non-participants. This parameter is the estimate of naturally occurring savings; that is, the post-installment savings that are common to both participant and NP samples.

Finally, the parameter β_5 represents average change in heating consumption for participants in the post-installation period. It is an estimate of program impacts differenced from base heating consumption (β_2), net of estimate naturally occurring savings (β_4) and controlled for differences in heating consumption between Ps and NPs (β_3).

The parameter β_1 represents the one combination outside of the fix effects that is not interacted with HDD. It captures average change in consumption for participants in the post-installation period that is not correlated with HDD. This parameter is expected to be negligible and is included for completeness.

The billing regression can include additional variables that explain gas heating consumption. Because each premise has its own unique intercept, all effects that are constant across all time periods for a premise are already accounted for. HDD is the primary time-varying variable, so additional variables generally enter interacted with HDD. These additional variables will be included in the model in groups of four similar to the variables associated with β_2 through β_5 but with the addition of the additional variable in the interaction.

Not only do additional variables enter the model in groups of four but the interactions between additional variables should also enter the model in groups of four. Each combination is necessary for balance, regardless of whether the parameter reflects an estimate that is significantly different from zero. In fact, because there is multi-collinearity in combinations of dummy variables as well as the other variables that will ultimately be included in the models, the combinations are likely to have relative poor p-values on an individual basis. However, these are dynamics that need to be controlled for regardless of the ability of the regression to discern them individually in a statistically precise parameter estimate. When all the variable permutations that represent true post-installation participant consumption change are put together, the standard errors are reasonable.

The time period effects are of particular importance in the pre-post billing regression. In the pre-post model, only a small proportion of households are in change mode in any give month. All other households are in a steady state pre- or post-program consumption mode. These households that are not in change mode contribute to a measurement of time period effects across all time periods despite the fact that all households go through a structural change at some point in the evaluation period.

The pre-post billing analysis approach also relies on the time-period effects to pick up any general trends in consumption. Where the survey indicated changes in consumption, we controlled for that change explicitly. Other household-specific consumption changes add to the variation in the estimates but will not generally affect the levels of the estimated parameters. Actions that potentially change a household's consumption occur throughout the time span of the analysis and can increase or decrease consumption. To affect the savings-related parameter estimates, changes would have to affect participants' post-installation consumption but not be related to the program.

Savings for Non-Standard Conditions

The billing regressions control for a variety of non-standard program conditions. These are identified with characteristics identified in the survey. Each characteristic indicates a condition under which the change in gas consumption due to the furnace installation is likely to be different than standard conditions. The increase or decrease in gas consumption in these cases is real, but they should not be included in an estimate of net savings for a standard program participant. We assume that house additions, fuel-switching etc would have happened regardless of the program incentive. Such conditions can be accounted for, if necessary, at the program level.

The non-standard installations include:

- The furnace did not replace an existing furnace; that is, the furnace was an add-on or part of new construction;
- The furnace replaced was not a gas furnace;
- The furnace does not serve the same amount of space;
- The furnace is not the primary source of heat; or
- Furnace serves more than one dwelling.

For the purpose of this evaluation, the goal with regards to these non-standard installations was to control for their effects. The incidence for any one of these situations was no greater than five percent so a meaningful estimate of the specific effects was not possible. Instead, a single dummy variable was created that indicated a non-standard condition and that was included in the regressions to account for the non-standard installations. As a result, the savings estimate is an estimate of savings for a gas furnace replacing a gas furnace in a dwelling where the gas furnace is the primary source of heat for only that dwelling.

Self Report Free Ridership Measure

The billing analysis approach pursued for this evaluation could, in theory, control for free ridership within the billing regressions. We also pursued a simple self-report free ridership analysis using the participant survey.

There are three levels at which free ridership is assessed:

- Awareness of energy efficient alternatives: Without this awareness, it is assumed the program is responsible for motivating the choice of the energy efficient alternative.
- Likelihood of installation without the rebate: For participants aware of the energy efficient alternative, we assess the likelihood that they would have installed that energy efficient alternative without the rebate. This provides the initial level of free ridership.
- Acceleration of installation due to the rebate: Even if a participant had some level of intent to install the energy efficient alternative, the program rebate might motivate the installation to happen sooner than it otherwise would have. If there is evidence of acceleration, the free ridership level is reduced and program credit is increased.

Results

Final Model Specification and Parameter Estimates

The final billing regression specification included HDD, square footage, the non-standard (odd) dummy, the participant dummy variable and the post-installation dummy variable in all possible combinations. Table 1 illustrates the explanatory variables included in the final specification. The "X"s indicate the input variable combinations for each explanatory variable. Each explanatory variable that includes HDD (base 62° F) enters into the model in groups of four, isolating the key participant/post-installation effect. The combinations including the variable "Odd" are grayed out as they are not are not included the final results combinations below.

Table 1. Final Billing Regression Parameter Estimates and Marginal Effects, Degree Day Base = 62° F.

HDD	Sq. Ft.	Odd	Post-Install	Participant	Paramter Estimate	Standard Error	P-Value	Marginal Effect, Annual Therms
			X	X	0.009497	0.041	0.819	0.0
X					0.112310	0.033	0.001	430.9
X			X		-0.035546	0.028	0.207	-136.4
X				X	-0.065531	0.036	0.067	-251.4
X			X	X	0.024056	0.032	0.455	92.3
X	X				0.000032	0.000	0.052	272.0
X	X		X		0.000015	0.000	0.313	129.3
X	X			X	0.000030	0.000	0.095	254.3
X	X		X	X	-0.000021	0.000	0.206	-181.1
		X	X		-0.123573	0.132	0.348	
		X	X	X	0.235201	0.148	0.113	
X		X			-0.017801	0.038	0.638	
X		X	X		0.040064	0.041	0.327	
X		X		X	0.013106	0.047	0.780	
X		X	X	X	-0.063610	0.058	0.272	
X	X	X			-0.000013	0.000	0.471	
X	X	X	X		-0.000001	0.000	0.950	
X	X	X		X	0.000007	0.000	0.730	
X	X	X	X	X	0.000016	0.000	0.573	

Table 1 also provides the parameter estimates and marginal effects of the important explanatory variables. The fixed effects model was run in SAS, using PROC MIXED with the empirical option. The empirical option estimates the variance-covariance matrix with an asymptotically consistent estimator. This removes a potential downward bias on standard errors caused by the presence of heteroscedasticity. The high standard errors on many of the individual explanatory variables are expected due to the collinearity caused by the large number of combinations of just five input variables. Despite the relatively poor p-values on individual parameter estimates, we will see in the next table that the most important sets of variable combinations do have sufficient precision. The marginal effects were developed using an average participant square footage of 2,238 square feet and a TMY-based average annual HDD of 3837. With all the variable combinations, the interpretation of individual variable marginal effects is more difficult.

Impact estimates

Table 2 provides the four groups of explanatory variables that decompose consumption into non-participant, pre-installation heating consumption, participant effect, installation effect and, finally, the participant impacts net of naturally occurring savings and controlling for differences between participants and non-participants. The table also included standard errors, P-values and the magnitude of the confidence interval.

Table 2. Final Estimates of Program-Related Marginal Effects, Per-Unit Therms

Parameter Combination	Estimated Usage	Standard Error	P-Value	90% Confidence Interval (+/-)
Non-participant, pre-installation Heating Consumption (hdd)	703.0	31.3	<.0001	-51
Participant - Non-participant Heating Consumption Difference (hdd*participant)	2.8	32.2	0.930	-53
Naturally Occuring Savings -- Post-Installation Consumption Difference (hdd*post-install)	-7.1	31.2	0.821	-51
Net Impact Estimate -- Participant Post-Installation Consumption Difference (hdd*participant*post-install)	-88.8	35.4	0.012	-58

The average non-participant pre-installation annual heating consumption was 703 therms plus or minus 51 therms. Participant pre-installation annual heating consumption was not statistically different from non-participant pre-installation annual heating consumption, with an increase of only 2.8 therms. This indicates the two samples were well matched with respect to heating characteristics.

The billing regression provides an estimate of naturally occurring savings at only 7.1 therms. This is the annual heating consumption reduction common to both participants and non-participants in the post-installation period. The result is not statistically significant. It is also unexpectedly small relative to the participant savings estimate. This could be caused by take back consumption in the post-installation period. It is not uncommon for installers of new furnaces to increase their setpoints in response to perceived decrease in heating costs related to the more up-to-date furnace. Any take back effect in the non-participant sample would decrease naturally occurring savings. Any take back effect in the participant sample would appropriately reduce net impacts.

Participant savings were estimated at 88.8 therms plus or minus 58 therms. In theory, this represents the impact of the program with respect to standard baseline and net of free ridership. At face value, these last two results indicate that naturally occurring savings are less than 10 percent of savings due to the program. Even assuming no improvement in standard efficiency, (a dubious assumption) this implies a net to gross ratio of less than eight percent. This is an unrealistically low estimate of free ridership.

Free Ridership

For this evaluation we were lucky to have two different estimates of free ridership with which to compare our billing analysis results. The self-report free ridership analysis performed for the evaluation indicated a free ridership rate of 55 percent. Moreover, PSE had the opportunity to speak with contractors active in its service territory and ask questions related to free ridership among customers. The key questions established the contractors' present energy efficient furnace sales as a percentage of sales and the contractors' expected EE sales percent if the program did not exist. The answer to these questions indicated a similar level of free ridership. EE sales would drop to 44 percent of their present levels indicating a FR rate of 56 percent. These two consistent results are not consistent with the billing regression results.

Self Selection

The mixed results of the billing regression and the free ridership analyses bring us back to the issue of self-selection. The billing regression results indicate the possibility that effectively all efficient furnace installers (potential free-riders) opted into the program and became free-riders. The estimated naturally occurring savings are not big enough to account for free ridership at a level close to those estimated by the two other approaches. In fact, the estimated naturally occurring savings are relatively low even for improvements in standards alone. The improvements in standards in the PSE service territory have been modest but a one percent improvement for the installation of a new furnace appears low.

Importantly, as explained, the self selection of efficient furnace installers is an identifiable kind of self-selection. As a higher percentage of this group opt into the program, naturally occurring savings is biased down to a lower limit of the naturally occurring savings due to the installation of new up-to-code furnaces. At the same time, net savings is biased upwards to the upper limit of gross savings (relative to a standard baseline). The extreme self-selection of energy efficient installers only undermines the removal of the "net" portion of the billing regression results.

One potential drawback is that the non-participant group includes some efficient furnace installers and thus controls for some portion of free ridership. A partially net final impact would be more difficult to correct with a separate free ridership factor. There is little evidence of this being the case in this evaluation but it could be a problem for programs with higher apparent levels of free-ridership behavior among non-participants.

Aside from free ridership, all the remaining sources of self-selection still exist. The billing model controls for overall gas consumption, heating consumption, square footage and non-standard installations. In controlling for these observed characteristics it is hoped to limit self-selection correlated with unobserved characteristics. In addition, we attempted to include the double IMR self-selection correction prescribed by Train and Goldberg. This correction cannot help the efficient furnace installer self-selection issue (for the IMR to work, both outcomes must be present across the range of characteristics), but it could correct for further self-selection issues.

With the inclusion of the IMR, the model parameter estimates remained essentially unchanged, but the standard error exploded. The latter effect is not unexpected. It is a common critique of the IMR that it introduces (further) heteroscedasticity into a model. The fact that the point estimates hardly moved, however, is at least some indication that additional self-selection is not a major factor in the billing analysis results.

Conclusions

This evaluation of the PSE efficient furnace rebate program uses a billing analysis approach including non-participant installers to generate an estimate of savings for the average furnace installation. The approach was designed to directly control for both kinds of naturally occurring savings, standard installations and free riders. In actuality, the approach provided an estimate of gross impacts, impacts relative to a standard installation baseline. Due to apparent self-selection of efficient furnace installers into the program, the billing analysis did not control for free ridership. The estimated annual impact of 88.8 therms is a sound estimate of savings relative to the appropriate standard installation baseline. Two separate estimates of free ridership put the rate at 55 to 56 percent.

References

Goldberg, M. and K. Train. 1996. *Net Savings Estimation: An analysis of Regression and Discrete Choice Approaches* prepared for CADMAC Subcommittee on Base Efficiency.

Green, W. 1997. *Econometric Analysis*. 3rd. ed. Upper Saddle River, NJ: Prentice-Hall, Inc.

Madalla, G. 1983. *Limited Dependent and Quantititative Variables in Econometrics.* Cambridge, England: Cambridge University Press