How Many Mills Ratios Does it Take to Estimate Net Savings?

Katherine Van Dusen Randazzo, KVD Research
Richard Ridge, Ridge & Associates
Kenneth Train, Xenergy
Leon Clarke, Pacific Gas and Electric Company

This paper reports on an early application of a technique to correct for selectivity biases in estimating net savings using billing regression models with participants and non-participants. Traditionally, one inverse Mills ratio is used for this correction but recent investigations have shown that this method corrects only for the correlation between participation and naturally-occurring savings. This analysis demonstrates the benefit of a model that, in addition to using the ratio in the usual way, interacts the inverse Mills ratio with the participation dummy. It is argued that this has the effect of adjusting for the correlation between participation and *net* savings.

The study was conducted on a sample of participants in the PG&E 1994 Commercial HVAC Retrofit program and a comparison group of program non-participants, using variables measuring building and business-related factors, as well as macro-economic trends, weather data, and information on extra-program installations of energy-related equipment. A direct change model was estimated, using a savings delta created by subtracting post-installation consumption from pre-installation consumption, using comparable months, pre and post. Independent variables were also change-based with the exception of building type variables. Two separate models were estimated, one for energy-using equipment (air conditioners, chillers, evaporative coolers, etc.) and one for non-energy-using equipment (timeclocks, EMSs, setback thermostats, etc.). This was necessary because the two types of installations require different comparison groups. Models that were consistent across the equipment groups were estimated and the effects of the Mills ratios were consistent across the two as well.

INTRODUCTION

This paper reports on the results of a modeling effort that was part of a comprehensive evaluation of a Pacific Gas and Electric Company (PG&E) Commercial HVAC Retrofit Program. The focus of the paper, however, is not on the model, but on the correction for selectivity biases. The program, the data collection processes, and the modeling process are described briefly so that the methods and results of correcting for selectivity biases can be interpreted in context.

Selectivity bias, in the context of conservation programs, can be defined as the tendency for certain types of customers to elect to participate in the program, and, correspondingly, for other types of customers not to participate. This problem occurs in programs where it is not feasible to randomly select customers for participation in the program, as is usually the case. Generally it can be assumed that customers who elect to participate will, on average, be different in some respects than those who do not decide to participate. This becomes a problem when trying to estimate the savings that can be attributed to the program; it is difficult to assess what proportion of the savings is due to the program itself and what is due to the nature of the customer who elects to participate in the program.

Traditionally, selectivity bias has been handled by inserting an inverse Mills ratio, consisting of a transformation of a predicted probability of program participation, into a regression model predicting savings. This method is intended to compensate for a situation where, e.g., customers who are likely to take energy-efficient actions *regardless* of incentives, are also more likely to decide to participate in energy-efficiency programs. In this case, some savings that would have occurred naturally would be attributed to program participation if the correction was not made. The inverse Mills ratio correction is intended to have the effect of reducing the bias represented by the correlation of participation with naturally-occurring conservation.

Recently, Train (1994) questioned the inverse Mills ratio method based on the fact that it does *not* correct for the correlation of participation with *net* savings. Goldberg and Kademan (1995) demonstrated, using simulated data where savings are known, that models using the traditional specification produced systematically biased estimates of program savings. Goldberg and Train (1995) conducted a more extensive series of simulation experiments to test the effects of using *two* inverse Mills ratios: the traditional one applied to the entire analysis sample (participants and non-participants), and another applied only to participants, intended to correct for the correlation of participation with *net* savings. This paper takes the concept another step by applying it to actual data in an evaluation of a DSM program.

PROGRAM DESCRIPTION

This study evaluated the gross and net energy savings from commercial HVAC measures that were paid rebates in 1994 through PG&E's retrofit energy efficiency programs. PG&E offers rebates to commercial customers who adopt energy-efficient measures that reduce HVAC energy consumption and demand in existing buildings. In 1994, 1434 customer applications were approved for rebates through the Retrofit Express and Retrofit Customized Programs covered by this evaluation.

METHODS

Sample

The participant sample relevant to the net impact analysis was based on a sampling frame of 1646 rebated items, stratified by savings size and measure type. Most strata were sampled with certainty, and the remaining strata (central air cooling, reflective window film, and other) were sampled using optimal allocation procedures[1]. For interviewing and analysis purposes, some 1100 unique locations were identified from the various item strata. Identification of customer-locations was essential because interviewing was conducted at the customer-location level. These locations were defined by a process of account matching using premise number, customer name and customer address to determine the accounts that comprised a customer location. From the 1100 unique locations, 450 interviews were completed.

The non-participant sample was selected from billing records in such a way as to assure that the non-participant sample had the same percentage in each usage stratum as the participant sample. Four hundred and fifty interviews of non-participant locations were completed. Accounts were grouped to locations using the same methods applied to the participant sample.

Before settling on final participant and non-participant interview samples, billing records were inspected for presence of continuous monthly consumption for the time period elapsed to that point in the study. Billing records of inadequate quality were eliminated from the sample. When the full set of post-installation billing records were available, further checks were made to identify changes in businesses or billing record termination. This resulted in the loss of 20 of the 900 original customers for whom we have interviews.

Data

In addition to monthly billing data spanning the period one year prior to installation and at least nine months after installation, four more categories of data were collected. First, interviews were conducted with 900 location-based customers to collect information on the building and the non-rebated equipment installed or replaced during the four-year study period. Also collected within this category were the business hours and building square footage information. Changes in these factors that occurred during the study period were also coded. Second, weather data, in the form of monthly heating degree days and cooling degree days were collected. Third, several macroeconomic variables were obtained (e.g., commercial employment figures, and retail sales) by Metropolitan Statistical Area. Fourth, rebated installations of equipment occurring in 1992–1995 (in addition to the sample-related installation for the 1994 program) were taken from the program tracking system.

Analytic Decisions

The most basic decision taken in the planning of this project was to pursue a change model. That is, the dependent variable would be savings, and the independent variables would be change-related variables. There are several methods of focusing on change in a modeling effort. One is simply to subtract the mean consumption for the post-installation period from the mean consumption over the pre-installation period. This method was not pursued because simply subtracting mean post from mean pre-installation consumption would result in different time periods for pre versus post installation consumption. This situation would result from the fact that customers who installed in December of 1994 would have only nine months of post-installation consumption data and these customers would have 12 months of pre-installation consumption. A second method is to produce deltas between pre and post consumption monthly means (using only comparable months in the subtractions) and take a mean of the monthly deltas, using that as the dependent variable. A third is to predict change as a percent of the pre-installation consumption mean. This can take the form of subtracting the natural log of the post from the natural log of the pre-installation consumption, or it can take the form of subtracting post from pre and dividing by pre-installation consumption. A fourth approach is to predict the post-installation consumption with a set of independent variables that includes the pre-installation mean consumption as well as change-related variables. In this approach, the change-related variables would predict the residuals resulting when regressing post consumption on pre consumption, which is equivalent to the difference between the pre and post mean. All but the first method were employed in this study, but only the method subtracting each mean post month consumption from each corresponding mean pre month are reported here.

A second decision was to estimate separate models for energy-using equipment and for non-energy-using equipment. The appropriate comparison group for program participants who installed energy-using equipment would be a group of non-participants who also installed energy-using equipment. This way, if non-participants install equipment with a lower level of efficiency than participants do, this will be reflected in the relative kWh consumption of the two. Comparison of program installers of energy-using equipment with a general population that includes non-installers would be inappropriate. This is because a general population may include businesses with no space cooling equipment and, who would, therefore, have no opportunity to decrease consumption due to cooling equipment changes (This is a theoretical distinction as in this study customers with no air conditioning were screened out of the sample). However, the same is true for businesses that have cooling equipment but do not need to replace it. If a business is not in the market for cooling equipment, there is no opportunity to purchase efficient equipment and consumption will not go down except for reasons unrelated to cooling equipment. The issue in determining the net effect of the program is to observe the effect of the rebate on the installation decisions of customers, and on the consumption. With nonparticipant non-installers, there is no opportunity to observe the customers' decisions on what level of efficiency to purchase. Non-participant *installers* had the opportunity to choose efficient or inefficient versions of this equipment category; thus, they serve as the appropriate point of comparison for program participants who have installed energy-using equipment.

Non-energy-using technologies present different issues, and require a different type of comparison group. The appropriate comparison group for this type of installer is all customers who have current equipment and situations that make it feasible to install a non-energy-using technology. Ideally, one would have a comparison group for each technology type, and that comparison group would consist of customers who don't have that technology but have a situation that would make it feasible. Of course, it is not practical to assess the feasibility of each technology for each site in the potential comparison group. It was, therefore, necessary to assume feasibility for each. The main assumption underlying comparison group decisions in this area is that essentially all customers have the option of deciding to install one or more non-energy-using technologies. It is always possible to add window film (unless it is already installed), and having installed that, it is possible to install a setback thermostat. Even with those items in place it may well be feasible to install an EMS, and so on. Because of the nearly constant possibility of these types of installations, and because they do not vary in efficiency level, a general population of program non-participants, all of whom have air conditioning, are reasonable points of comparison. It could be argued that the ideal comparison group for these participants would be customers who had not installed any of these technologies over the past several years. However, screening for such customers is not practical. Of course it would cost a great deal to do it, but more than that, customers often do not know if such equipment has been installed over the years. This is especially true for EMS systems and HVAC motors. Therefore, the general population of non-participants may include some businesses that installed some non-energy-using equipment without rebates. This actually is not a problem since it represents customer behavior outside of the program. Of course, to the extent that such installations are spillover effects of prior program exposure or reflect market transformation, they contribute to unfair minimization of program effects. However, this could not be avoided in this study.

On the other hand, it would be inappropriate use non-energy-using equipment *installers* as a comparison group since there is no efficiency variation in these types of equipment. If you install the equipment, you get the fixed effect (fixed based on building, climate, and usage characteristics). To make the point more concrete: using a group of non-participant installers of, e.g., setback thermostats, as a comparison group for a group of participant installers of the same measure would yield an apparent net impact of zero. The logic applies under the usual circumstance of combining many measure types together such as thermostats, window film, insulation, or EMSs. Using non-participant installers of such equipment as a comparison group would reduce the apparent net savings to something close to zero. Therefore, the appropriate comparison group for this type of equipment is a general population of customers who have air conditioning and, therefore, could conceivably add a non-energy-using, energy conserving measure.

It is also important to note that the decision to estimate models on these two separate groups has implications for sample size and variances. This is true because the sample on which each model is estimated will be divided into two smaller parts. Because the two parts will, by definition, be smaller than the sample as a whole, variances will be increased and it will be more difficult to find statistically significant t's for the regression coefficients.

Another analysis decision concerns how to define the date that divides the pre-installation period from the post, i.e., the installation date. There was ambiguity in the records as to the date of the installation, including the fact that month and year were available, but the day was not. Beyond that, however, it was often not clear in which month the actual installation took place. The date most consistently available was that defining the month in which authorization for payment occurred. Since this implied that the installation occurred before that date, the authorization month was included as the first month in the post-installation period. One or more months were not eliminated from the analysis because this would mean that it would be impossible to have 12 consecutive pre-installation months (required by the protocols) in the analysis. Thus, the authorization month was defined as the pivot month that marks the beginning of the post-installation consumption period. Therefore, we almost certainly identified some post-installation billing data as pre-installation. Consequently we will have under-estimated savings.

The categories of variables measured/collected to explain savings were all (with one exception) change-based and all would be expected to have an impact on the change in consumption over time. Six groups of variables were involved:

(1) Weather differences

(2) Macro-economic trends

(3) Changes in business hours

(4) Changes in square footage

(5) Equipment installed without rebates (self reported)

(6) Equipment rebated by other programs (1992,93/95 HVAC, 1992–95 Lighting, data taken from the tracking system)

Building types were also considered important, in spite of being constant over time, because they can represent unobserved trends in some segments of the economy that may not apply to all.

The Basic Modeling Approach

The change-based regression model that was employed has the following general form:

$$\Delta kWh_i = \alpha + \delta_1 Part_i + \delta_2 Mills_i + \delta_3 Miles_i * Part_i + \sum \delta_k \Delta x_i + \varepsilon_i$$

where

ΔkWh_i = change in kWh consumption from before the program to after the program for the i^{th} customer;

α = a constant that captures the energy consumed by the unspecified equipment;

δ_1 = a coefficient that reflects the energy change associated with participation

$Part_i$ = a binary indicator representing participation status;

δ_2 = a coefficient that reflects the energy change associated with the selectivity correction factor

δ_3 = a coefficient that reflects the energy change associated with the selectivity correction factor for participants who installed a rebated measure;

$Mills_i$ = the selectivity correction factor for the i^{th} customer

$Mills_i * Part_i$ = an interaction term that captures self-selection for participants only

δ_k = a vector of k coefficients that reflect the energy change associated with a one-unit change in the k^{th} explanatory variable;

ΔX_i = a vector of other explanatory variables, such as changes in square feet, operating hours, equipment stock, and the rate of inflation from before the program to after the program for the i^{th} customer; and

ε_i = the differences in energy consumption that are not explained by the model.

As indicated by the subscript of the change variable, the analysis was carried out on a dataset with one observation per customer location.

Mills Ratios

Selectivity biases are likely to be present in conservation programs. Selectivity bias can be defined as the correlation of naturally-occurring savings with the decision to participate. Consider the following example: customers who tend to conserve energy may also be more likely to participate in conservation programs. Traditionally, this source of bias is corrected by inserting an inverse Mills ratio into the model. This method was intended to correct for the correlation between participation and naturally-occurring savings. Without that correction some naturally-occurring savings among participants would be interpreted as net savings and would appear in the parameter connected with the participation variable.

In this analysis, we have added a second inverse Mills ratio, interacted with the participation variable, which has the effect of accounting for the correlation between *net* savings and participation. This will be explained further in the results section.

The process of correcting for selectivity biases was implemented in the following steps: (1) estimate a logit model for participation, i.e., find the determinants of participation in the program, (2) calculate an inverse Mills ratio from the probability of participation, and (3) estimate the regression model of consumption change integrating the inverse Mills ratios. The result of the logit model (a predicted probability of participation) was used here to calculate the inverse Mills ratio.

For participants, it was calculated as:

$$\text{Mills} = -\left[\frac{(1-P) \times 1\text{n}(1-P)}{P} + 1\text{n}P\right]$$

For non-participants:

$$\text{Mills} = \frac{(P) \times 1\text{n}(P)}{1-P} + 1\text{n}(1-P)$$

where P = the probability of participation.

Modeling Goals

The goals in the search for the best model were to find a set of predictors:

(1) with coefficients stable enough to apply to both energy-using and non-energy-using models;

(2) with plausible coefficients both in terms of size and sign (plausibility was gauged against knowledge of the tracking system reported savings and by taking the estimated savings as a proportion of pre-installation consumption);

(3) that pass diagnostic tests for influential observations, heteroskedasticity (i.e., the variances around estimates are different for different levels or values of the predicted variable), and collinearity (two or more independent variables in a model are highly intercorrelated); and

(4) that include all categories of variables that are expected, *a priori*, to belong in the model.

Results

The variables ultimately included in the model, the process by which they were selected, and the interpretation of their behavior in the model are not discussed in this paper. This decision was necessary in order to focus more on the central issue of the paper: the addition of a second inverse Mills ratio to the regression model.

The two models (energy-using and non-energy-using models), are shown without correction for selectivity in Table 1. The same variables are included for both equipment groups. The final set of variables (applied to both equipment groups) was selected *before* applying Mills ratios for correction of selectivity biases. The reasoning behind this approach was that the variables that best predict consumption behavior in the most stable and consistent manner should be established first before addressing higher-order issues. Then the measures of self-selection biases should be applied to the model and their effects on the model determined. This model has the following characteristics. First, the dependent variable is the direct mean of the monthly pre/post deltas using comparable months. This dependent variable, in combination with the explanatory variables chosen, produces a model that allows the same specification to be used for both equipment groups. That is, the variables remain the same and, with one exception (SUMHRS in the non-energy-using model), the signs of the coefficients are appropriate and remain the same (ignoring the building type signs, for which we have no directional expectations). Further, the magnitudes of the coefficients are plausible and it is possible to understand why they are operating as they do.

The use of the direct monthly delta as the dependent variable results in some complexities in interpretation. The interpretations are complicated by the fact that the explanatory variables are also change oriented, so the interpretation has to take into account the effect of the *change* in the independent variable on the *change* in the dependent variable. It is important to remember that the post period kWh value is always subtracted from the pre-period value. Thus, a decrease in consumption will result in a *positive* delta, or a positive savings. Both weather and economic data were treated the same way: post values were subtracted from pre, with the pivot point defining the two time periods being the same as that used for consumption. Therefore, a decrease in employment between the pre-installation and post-installation periods would result in a positive delta just as it would for the consumption data. Since we would expect less consumption during periods of less employment, we would expect a positive correlation between employment change and consumption change. This is intuitively apparent. It is more complex, however, for the other change variables. To help the reader make interpretations more easily, plus and minus signs are included in the tables to illustrate the expected directions of the relations of the independent variables with the dependent variable of consumption change.

Results of Corrections for Selectivity Bias

Without correcting for selectivity bias, the model shows that the average monthly net savings for premises installing

Table 1. *Final Models Before Correction for Selectivity*

Variable	Variable Description	Energy-Using Expected Sign	Energy-Using Parameter Estimate	Energy-Using T Prob	Non-Energy-Using Parameter Estimate	Non-Energy-Using T Prob
INTERCEPT			980	0.82	385	0.8
PART	Program Participation	−	2608	0.56	3909	0.03
DIFHDDMN	Pre minus post heating degree days	−	58	0.66	48	0.64
SUMHRS	Change in business hours	−	6.6	0.98	42	0.87
SUMSQ	Change in square footage	−	−2.1	0.02	−1.3	0.0001
INSTOUT	Some installation outside of program	+/−	−8490	0.06	−5859	0.28
DIFEMPMN	Pre minus post employment rates	+	410	0.62	875	0.06
OFFICE	Office building	+/−	−1549	0.56	3434	0.04
REST	Restaurant	+/−	−1497	0.74	15	0.99
RETAIL	Lg. or sm. retail	+/−	2017	0.62	1594	0.51
FOOD	Food store	+/−	724	0.91	3618	0.5
WARE	Warehouse	+/−	−2868	0.59	−2134	0.54
SCHOOL	Primary/secondary schools	+/−	−211	0.95	−348	0.93
CLINIC	Clinics	+/−	368	0.94	−1019	0.82
HOTLMOTL	Hotel or Motel	+/−	926	0.9	−1740	0.71
PERSREP	Personal repair services	+/−	−2289	0.75	107	0.97
COMSERV	Community Services	+/−	510	0.9	−1844	0.47
N			169		740	
R^2			0.07		0.17	

energy-using equipment would be 2609 kWh. The analogous savings estimate for premises that installed non-energy-using equipment would be 3909 kWh. However, it is important to take account of the fact that participants and non-participants are not randomly assigned to those categories. Different businesses may be differently exposed to the program and may be differently inclined to participate. Traditionally, an inverse Mills ratio is included in the model to account for the relation between naturally occurring savings and participation. Now a second inverse Mills ratio is entered for participants to account for the relation between *net* savings and participation. When the two inverse Mills ratios are added to the models for energy-using and non-energy-using equipment, the same pattern emerges in each model (Table 2). In particular, the first inverse Mills ratio enters with a positive sign in each model. This indicates that naturally occurring savings is positively correlated with participation. Stated equivalently, a customer who would have a

Table 2. *Models After Correction for Selectivity*

Variable	Variable Description	Energy-Using Expected Sign	Energy-Using Parameter Estimate	Energy-Using T Prob	Non-Energy-Using Parameter Estimate	Non-Energy-Using T Prob
INTERCEPT		+/−	2154	0.9	845	.86
PART	Program Participation	+	8438	0.62	7618	.19
DIFHDDMN	Pre minus post heating degree days	+	64	0.63	59	0.56
SUMHRS	Change in business hours	−	−74	0.78	43	0.87
SUMSQ	Change in square footage	−	−2.7	0.007	−1.3	0.0001
INSTOUT	Some installation outside of program	−	−12924	0.014	−7014	0.2
DIFEMPMN	Pre minus post employment rates	+	344	0.68	845	0.08
OFFICE	Office building	+/−	−2604	0.34	3359	0.04
REST	Restaurant	+/−	−830	0.85	191	0.95
RETAIL	Lg. or sm. retail	+/−	881	0.83	1561	0.52
FOOD	Food store	+/−	774	0.9	3787	0.49
WARE	Warehouse	+/−	−2798	0.6	−1745	0.62
SCHOOL	Primary/secondary schools	+/−	−1853	0.61	−882	0.82
CLINIC	Clinics	+/−	178	0.97	−1355	0.76
HOTLMOTL	Hotel or Motel	+/−	1767	0.81	−1293	0.78
PERSREP	Personal repair services	+/−	−1069	0.88	68	0.98
COMWERV	Community Services	+/−	2106	0.62	−2132	0.41
MILLS	Inverse Mills ratio-all premises		744	0.95	330	0.92
MILLS2	Inv Mills ratio for participants		−5884	0.63	−3762	0.38
N			169		740	
R^2			0.09		0.18	

tendency to take measures even without the program (that is, customers who are naturally inclined toward conservation) tend to participate in conservation programs.

The second inverse Mills ratio enters with a negative sign in each model. This result implies that net savings is negatively correlated with participation. In a sense, the estimated relation between net savings and participation is the expected consequence of the estimated relation between naturally occurring savings and participation. Customers who would have taken the measures even without the program have high naturally occurring savings; the positive coefficient of the first inverse Mills ratio implies that these customers tend to participate in the program more readily. Since these customers would have taken the measures without the program, their net savings are generally low[2] the fact that they participate in the program more readily means that the net savings is negatively correlated with participation—as implied by the negative coefficient on the second inverse Mills ratio.[3] The inverse Mills ratio therefore enters the two models in a consistent and highly plausible manner.

Net KWh Savings

Having settled on the model shown in Table 1, and on the method of correcting for selectivity as shown in Table 2, the next step in the process was to calculate average, premise-level net savings. This is calculated by adding the product of the MILLS2 parameter and the mean of the MILLS2 for participants to the parameter estimate for the participation dummy (PART). For the energy-using model, this translates to: $8438+(-5884*1.136)$, and for the non-energy-using model, it means: $7618+(-3762*1.182)$.

With the two inverse Mills ratios included, the model for energy-using equipment provides an estimated net savings of 1754 kWh per month, as shown in Table 2. The model for non-energy-using equipment provides an estimated savings of 3171 kWh per month when the two inverse Mills ratios are included. The corrected estimate for energy-using equipment savings is 854 kWh or 33 percent lower than the uncorrected estimate. The corrected estimate for non-energy-using equipment is 738 kWh or 19 percent lower than the uncorrected estimate. These figures are plausible and within a reasonable range of each other, lending credence to this adjustment method. Further indication that the double Mills method is working appropriately is the fact that the model with the ratios fits better than without, as measured by the R^2 (although the R^2 and the increment in R^2 are quite small).

CONCLUSION

This study has been an early application of the method of adding a second inverse Mills ratio to a model that estimates program-related savings. The purpose of the method is to avoid the biased estimate of program savings that occurs with the traditional use of a single inverse Mills ratio. The bias comes from the fact that, while attempting to correct for naturally-occurring savings, the application of the inverse Mills ratio to the entire sample of participants and nonparticipants neglects the correlation of potential *net* savings with the participation decision. The addition of a second inverse Mills ratio only for participants has the effect of adjusting for *net* savings, which is correlated with the probability of participation as represented by the inverse Mills ratio. Since this component of savings that is related to the inverse Mills is interpreted as net savings, it is used to adjust the net savings represented by the participation variable parameter. The interpretation of the second inverse Mills as the net savings that is correlated with participation is supported by the entry of its parameters in both equipment models with the appropriate negative sign, in the context of the positive sign found with the original inverse Mills.

END NOTES

1. Specifically, the Dalenius and Hodges method of setting strata boundaries to minimize variance was used, using a Neyman allocation. The method is described in: Dalenius, T., and Hodges, J.L., Jr. (1959). Minimum variance stratification. Journal of the American Statistical Association, 54, 88–101.

2. Consider two customers who have installed equipment that would save 1000 kWh. For the first customer, 80 percent of this, or 800 kWh was naturally-occurring savings, implying that 200 kWh was net savings. For the second customer 30 percent of the 1000 kWh gross savings is naturally occurring, implying 700 kWh as net savings. This examples demonstrates the idea that the higher the naturally-occurring savings, the lower the net savings must be.

3. More precisely, the negative sign indicates that the unobserved factors associated with the participation decision are negatively correlated with the unobserved factors related to net savings.

REFERENCES

Goldberg, M.L. And K. Train, (1995) *Net Savings Estimation: An Analysis of Regression and Discrete Choice Approaches.* Submitted by Xenergy Inc. To CADMAC Subcommittee on Base Efficiency.

Heckman, J., (1979). *Sample Selection Bias as a Specification Error*. Econometrica, 47(1), pp. 153–162.

Train, K., 1994, ''Estimation of Net Savings from Energy-Conservation Programs,'' *Energy*, Vol. 19, No. 4, pp. 423–441.

Goldberg, M. And E. Kademan, 1995, ''Is it Net or Not? A simulation Study of Two Methods,'' *Proceedings* of the Seventh International Energy Program Evaluation Conference, Chicago.