Estimation of Net Savings for Rebate Programs: A Three-Option Nested Logit Approach

Kenneth Train, University of California, Berkeley
Susan Buller, Pacific Gas and Electric Company
Bruce Mast and Kirtida Parikh, Pacific Consulting Services
Eric Paquette, Cambridge Systematic, Inc.

We introduce a method to estimate net-to-gross ratios for programs, such as rebate programs, in which a customer can only participate in the program (i.e., receive a rebate) by implementing an energy-efficient measure. A nested logit model describes the customer's choice of whether to implement the measure and whether to apply for and receive a rebate. The impact of the program is determined by simulating the implementation choice of customers when they do not have the option of receiving a rebate. The method is applied to Pacific Gas and Electric's incentives program and Southern California Edison's hardware rebate program for commercial, industrial and agricultural customers.

Introduction

To estimate the net savings of a demand-side management program, it is necessary to disentangle the patterns of causation between program participation and measure implementation: did customers participate in the program because they had already decided to implement energy-efficient measures and simply wanted the program benefits, or did the program actually induce them to implement the measures? In an earlier paper, Train (1994) described econometric methods for identifying the two directions of causation in certain kinds of programs, namely, those in which a customer can participate without necessarily implementing any measures, such as an audit program. These methods have been applied to data from Southern California Edison's (SCE's) audit program for commercial and industrial customers, with reasonable results (Pacific Consulting Services 1993.) For rebate programs, however, these methods are not applicable, since participation in the program (i.e., receiving the rebate) requires that the customer implement a targeted measure.

The current paper describes and applies econometric methods to estimate net savings for rebate and other incentive programs where participation requires that a targeted measure be implemented. Under a typical rebate program, one or more measures (such as the installation of high efficiency lighting) are eligible for rebates. Each customer has a choice among three options regarding each eligible measure; in particular, the customer can (1) implement the measure and receive the rebate, (2) implement the measure but not apply for or receive the rebate, and (3) not implement the measure. A customer who chooses the first option is considered a participant in the program, and the energy savings from the measure are counted in the gross savings of the program. To determine net savings, a discrete choice model is estimated that describes customers' choices among these three options, using data on the actual choices that customers made during the program period. This model is specified as nested logit, though multinominal probit could be used instead. Once the three-option model is estimated, it is used to simulate the behavior of customers with the first option removed (that is, to "forecast" what customers would have done if they had not had the option of implementing the measure with a rebate.) This simulation indicates the extent to which customers would have implemented the measures without the program; the energy savings under this simulation are the estimate of naturally occurring savings. The net savings of the program are then calculated as the difference between (i) the savings that occurred with the program (i.e., when all three options are available), and (ii) the naturally occurring savings. Details of the method are provided in the next section below.

In the "results" section we describe the application of this method to the commercial, industrial and agricultural

(CIA) incentives programs of Pacific Gas and Electric (PG&E) and SCE, respectively. Different types of data were used to estimated the models for the two programs, as well as somewhat different estimation algorithms. The resulting estimates of net savings are reasonable and conform to estimates obtained with other methods (such as self-report surveys) for the same programs. The last section concludes with a general discussion of the difficulties and uncertainties entailed in applying the method.

Methodology

During the program period, each customer faces three options, denoted A-C, regarding each rebate-eligible measure:

A. Implement the measure and receive a rebate
B. Implement the measure and not obtain a rebate
C. Not implement the measure.

The customer chooses the option that provides it with the greatest profit, net benefits, or, more generally, "utility." The utility that the customer obtains from each option depends on the investment cost, energy consumption costs, and other factors associated with the option. For example, for installation of occupancy sensors, the investment cost for option A is the cost of installing the sensors minus the rebate, for option B is the installation cost (not net of rebate), and for option C is zero. Savings from the sensors enter options A and B, since these savings are the same whether or not the customer receives a rebate.

Some factors that affect the utility of each option are observable (such as the installation cost and expected savings). However, other factors are not. For example, the non-monetary "hassle" of making changes, which cannot meaningfully be measured, might affect the utility to the customer of options A and B. The customer's uncertainty about the cost and especially about the savings can also be expected to affect its utility from options A and B but are not generally observed. Some unobserved factors affect the utility of option A that do not affect option B; for example: the hassle of applying for a rebate, and the cost and difficulty of documenting the installation and its cost (which is usually needed to receive a rebate). In fact, it is because of these factors that a customer might choose option B over option A (i.e., implement the measure but not apply for a rebate.)

Denote the utility associated with each option as consisting of observed factors and an error term that captures the effect of unobserved factors:

$U_i = x_i + \epsilon_i$ for i = A, B, and C.

The distribution of the unobserved factors ϵ_i determines the probability of each option being chosen. If the unobserved factors were independent over options, then a standard log it model could describe the probabilities. However, the unobserved factors are clearly not independent: unobserved factors relating to the installation of the measure enter the utility for both options A and B, since both of these options entail implementing the measure. A model is required, therefore, that recognizes the correlation between ϵ_A and ϵ_B.

Recognition of this correlation is critical to the estimation of net savings. An assumption that the errors are *not* correlated (as for a simple logit model) is equivalent to assuming that the implementation rate among customers who did not receive a rebate (i.e., who chose options B or C) is the implementation rate that rebated customers (i.e., customers who chose option A) would have had in the absence of rebates. This is essentially the same as saying that participants would have behaved like non-participants if the program had not existed, which Train (1994) indicates is an inappropriate assumption for net savings analysis.

Nested logit explicitly recognizes the correlation in unobserved factors over options. Let the errors be jointly distributed generalized extreme value, with correlation between ϵ_A and ϵ_B. McFadden (1978) shows that under this distribution, the probability that the customer chooses option A is:

$$P(A/A,B,C) = \frac{e^{\beta x_A/(1-\lambda)}\left[e^{\beta x_A/(1-\lambda)} + e^{\beta x_B/(1-\lambda)}\right]^{-\lambda}}{\left[e^{\beta x_A/(1-\lambda)} + e^{\beta x_B/(1-\lambda)}\right]^{1-\lambda} + e^{\beta x_C}} \qquad (1)$$

where: X is a measure of the similarity, or correlation, between 6A and ~B, and the notation A/A,B,C means that A is chosen when options A,B, and C are available. The probability of option B is a similar formula, with xB replacing xA in the first term in the numerator. The probability function for option C is different, since the error associated with option C is not correlated with the other errors:

$$P(C/A,B,C) = \frac{e^{\beta x_c}}{\left[e^{\beta x_A/(1-\lambda)} + e^{\beta x_B/(1-\lambda)}\right]^{1-\lambda} + e^{\beta x_C}} \qquad (2)$$

The parameters of the model are estimated on a sample of customers with information for each customer on whether it implemented the measure during the program period and whether it received a rebate (which together identify which of the three options the customer chose) plus data on the cost and savings associated with the measure. Note that data are required on the cost and savings of each measure that *could* have been taken by the customer during the program period, whether or not the customer actually implemented the measure or received a rebate.

Once the parameters are estimated, net savings are estimated through a series of steps. First, the model is used to estimate the gross savings from the program associated with this measure. Gross savings are the savings from all rebated measures. Let the subscript n denote a sampled customer, w_n be the sampling weight for customer n (equal to the inverse of the customer's chance of having been sampled), and S_n. denotes the energy savings that customer n would obtain from implementing the measure. The estimated gross savings from the program for this measure (e. g., from occupancy sensors) is:

$$GS = \sum_n w_n \cdot P_n(A/A,B,C) \cdot S_n \ , \qquad (3)$$

where $P_n(A/A, B, C)$ is the probability that customer n chose option A when all three options were available.

Second, the model is used to estimate the total savings that were obtained from implemented measures, whether the customer received a rebate or not:

$$TS = \sum_n w_n \cdot \left[P_n(A/A,B,C) + P_n(B/A,B,C)\right] \cdot S_n \ . \qquad (4)$$

Note that total savings is equal to gross savings plus the savings from measures that were implemented without the customer receiving a rebate.

Third, naturally occurring savings is estimated. Naturally occurring savings is the savings that would have occurred without the program. If the program had not existed, then customers would not have had option A available: they would not have been able to implement the measure and receive a rebate. Their choice would have been to either implement the measure without a rebate (option B) or not implement the measure (option C). The probability of the customer choosing to implement the measure without the program is:

$$P(B/B,C) = \frac{e^{\beta x_B}}{e^{\beta x_B} + e^{\beta x_C}} \qquad (5)$$

The estimated naturally occurring savings is then:

$$OS = \sum_n w_n \cdot P_n(B/B,C) \cdot S_n \ . \qquad (6)$$

Finally, net savings is the difference between the savings that occurred with the program and the naturally occurring savings, which is estimated as: NS = TS - OS. The estimated net-to-gross ratio is NTGR = NS/GS. [1]

Results

Application to PG&E's CIA Incentives Program

PG&E has an on-going program that provides rebates to commercial, industrial, and agricultural customers for the installation of energy efficient measures. For regulatory purposes and internal planning, estimates of the net savings from the program were required. The method described above was applied, as well as three other procedures. In this section, we describe the data that were collected for this application, the estimated models, and the resulting net-to-gross ratios. We then compare the results to those obtained with the other three procedures.

Data were obtained on a sample of 926 customer premises. The sample was specially designed to include a large share of participants (i.e., customers who chose option A for at least one measure) as well as a large share of customers who installed those measures that accounted for most of the program's gross savings (e.g., lighting and HVAC measures.) A weight for each customer was derived from the sampling scheme, such that weighted sample is representative of the population. On-site surveys were conducted for each sampled customer. The customer's premises were divided into "inventory groups," representing distinct portions of the building. This breakdown by inventory groups reflects the fact that customers can make separate decisions relating to each inventory group, such as installing high efficiency lighting in one inventory group but not another. For each inventory group, the surveyor determined which rebate-eligible measures could have been performed since 1990. For example, the surveyor determined whether it would have been possible to install high efficiency lighting after the beginning of 1990; if such lighting had already been installed prior to 1990, then installation during the relevant program period was considered not feasible. For each measure that could have been installed in each inventory space, the surveyor estimated the cost of implementing the measure and the energy savings that could be expected from the measure. The surveyor also determined whether the customer actually implemented the measure since 1990.

For each measure that could have been taken in each inventory group, the customer chose one of the three options, A,B, or C, given above. The 926 sampled customers provided information on 6734 inventory group implementation decisions, which are observations in the estimation of the nested logit model of customer's choices among the three options. Since different factors can enter the choice for different types of measures, separate models were estimated for each type of measure. The measure groups are:

- Lighting Conversions (from incandescent to fluorescent)
- Lighting Controls (such as sensors and time clocks)
- Lighting Upgrades (energy-efficient lamps, electronic ballasts, delamping)
- HVAC Controls (such as set-back thermostats and time clocks)
- HVAC Maintenance (mainly coil cleaning)
- New HVAC Installations

These groups of measures account for 58% of the gross savings for the 1990 program; the remaining 42% came from measures that were either too diverse to model or for which on-site surveyors could not meaningfully determine whether the measures could have been taken in 1990.

The nested logit model for lighting conversions is given in Table 1. The parameters are estimated by maximum likelihood. The first column gives the variables that enter the utility for each alternative (that is, the elements of x.) The second and third columns give the estimated parameters (i.e., *ß)* and t-statistics. Option C represents no change (that is, not implementing the measure), and its utility is normalized to zero.

The first variable is the cost of the measure divided by the annual savings, with the cost being net of the rebate for option A and not net of the rebate for option B. The variable enters with a negative sign, as expected, meaning: (i) a rise in the cost of implementing the measure reduces the probability that the customer implements it (i.e., chooses either option A or B) if the rebate and savings from the measure remain the same, (ii) an increase in the rebate offered on a measure increases the probability that the customer will implement the measure and receive the rebate (i.e., choose option A), and (iii) a rise in the savings from a measure will increase the probability that the customer will implement the measure (i. e., will choose either option A or B).

Variables 2- 5 enter the utility of options A and B the same. Therefore, these variables affect the customer's probability of implementing the measure (i.e., choosing options A or B) but not whether the customer receives a rebate. The estimates indicate that restaurants and retail stores are less likely to implement lighting conversions than other types of businesses, holding all else constant. Businesses that consider themselves to be environmentally oriented are more likely to implement these measures. And measures that cost less than $100 are less likely to be implemented than measures that cost more but have the same payback period; this result is presumably due to the fact that a small investment provides a small savings (when the payback period is the same as for a larger investment) and businesses do not consider the effort required to implement the measure to be worth the small savings.

Variables 6-8 enter the utility of option A and not option B. These variables therefore affect the probability that the customer will receive a rebate. As expected, customers who are aware of the rebate program or had the measure recommended to them during the course of an audit are more likely to implement the measure and receive a rebate than other customers. [2]

The estimate of λ is given at the bottom of the table. As described above, λ captures the similarity between options A and B, or more precisely, in the unobserved factors affecting the utility of these two options. The parameter is 0.39, indicating a moderate correlation. This parameter is the most important factor in the estimation of net savings. If $\lambda = 0$, then the nested logit model is equivalent to a simple logit model, in which case P(B/B,C) = P(B/A,B,C)/[P(B/A,B,C)+P(C/A,B,C)]. That is, the predicted implementation rate without the program is the same as the actual implementation rate that occurred among customers who did not receive a rebate for this measure. If the program targets only one measure, this is equivalent to saying that the participants, without the program, would have implemented the measure at the same rate as non-participants did with the program. [3] As Train (1994) has argued, this is generally an unreasonable assumption to make for net savings analysis. The nested logit model, by allowing λ to take values other than zero, does not make this assumption. However, since λ is estimated, the estimated value of λ can be zero, if reality actually conforms to this assumption.

Generally, a larger estimated value of λ translates into a larger estimate of naturally occurring savings and a smaller estimated net-to-gross ratio. The underlying logic is straightforward: a high λ means that implementing with the rebate is seen by customers to be very similar to implementing without the rebate, such that the rebate has little effect on their decision to implement.

Table 2 gives the estimated net-to-gross ratio for each of the measures for which models were estimated. The

Table 1. Nested Logit Model for Lighting Conversion Measures for PG&E's CIA Incentive Program

Variable (options that variable enters are given in parentheses)	Estimated coefficient	T-statistics
1. Paycheck period (A, B)	-0.059	0.1
2. Dummy for: restaurants (A,B)	-1.736	1.3
3. Dummy for: retail stores (A,B)	-1.527	1.3
4. Dummy for: cost of measure is below $100 (A,B)	-2.954	1.4
5. Dummy for: customer describes self as environmentally oriented (A,B)	0.449	0.8
6. Dummy for: measure was recommended in an audit (A)	1.031	2.3
7. Dummy for: customer was aware of rebate program (A)	2.149	3.8
8. Size of rebate times dummy for customers who were aware of rebate program (A)	0.0000755	0.5
9. Constant (A)	-2.117	1.5
10. Constant (B)	-5.269	3.3
λ	0.3935	0.91

net-to-gross ratio for lighting conversions is estimated to be 0.53. This estimate is based on the model of Table 1, with its estimated λ of 0.39. The net-to-gross ratios for the other measures are based on models of similar specification with their own estimated parameters. These models are given in Cambridge Systematic (1993a) and, for brevity, are not repeated here. In addition, that volume provides many alternative models for each measure with the resulting net-to-gross ratios. Generally, these alternative ratio estimates are similar to, or below, those reported in Table 2. The net-to-gross ratio for the program as a whole is estimated as a weighted average of the measure-specific ratios, with the weights being proportional to gross savings from each measure. The estimated ratio for the program as a whole is 0.73.

Three other methods have been used to estimate the net-to-gross ratio for PG&E's CIA Incentives Programs and the results are similar to those show in Table 2 (see Cambridge Systematic 1993b.) These methods are the following. (1) Self-report surveys were utilized, in which customers were asked, in various ways, whether they would have implemented the measures without the rebates. These surveys produced estimated net-to-gross ratios that ranged from 0.39 to 0.60, depending on how inconsistencies and other issues were handled. The most reasonable estimate was considered to be 0.58. (2) A survey of dealers and contractors was conducted to determine how much the sales of energy-efficient products had changed over the last several years (when the program was operating.) The responses of dealers and contractors in PG&E's territory were compared with those in areas not covered by PG&E's program (i.e, "control" areas). The difference-in sales growth was 'attributed to the program. The net-to-gross ratios implied by this study ranged from 0.72 to 0.86. (3) The implementation rates of customers in areas covered by PG&E's program were compared with the implementation rates of customers in areas that did not have such a program. After accounting for other differences (such as building types and size), the remaining differences were attributed to the program. The net-to-gross ratio obtained with this study is 0.75.

The estimate of 0.73 using the nested logit approach is within the range of results found in these other studies (0.58 for self-report, 0.72-0.86 for the dealer/contractor survey, and 0.75 for the comparison of implementation

Table 2. Estimated Net-to-Gross Ratios for PG&E's CIA Incentives Programs

Lighting conversions	53%
Lighting controls	80%
Lighting upgrade	84%
HVAC controls	75%
HVAC maintenance	84%
HVAC new installations	43%
Weighted average	73%

rates.) It is important to note, however, that these studies do not include the same factors in their estimates, and so the results are not directly comparable. In particular, the dealer/contractor survey and the analysis of customer implementation rates compare behavior in the program area with behavior in a control area that did not have the program. As such, these studies include all effects of the rebate program, including spillover.[4] In contrast, the nested logit approach does not include spillover effects in its calculations. The nested log it method is therefore most comparable with the self-report survey, which also does not consider spillover. In this comparison, the nested logit model provides a higher estimate: 0.73 compared to 0.58 from the self-report. Of course, it is not clear whether this difference is due to the nested logit providing too high estimates, the self-report providing too low estimates, or simply randomness.

Application to SCE's CIA Hardware Rebate Program

SCE's program was examined in a manner analogous to that described above for PG&E, with two differences. First, data were obtained from a telephone survey rather than on-site visits. That is, sampled customers were called and asked what measures they installed and what measures they could have installed during 1990. Second, the nested logit models were estimated sequentially rather than simultaneously. One convenient property of a nested logit model is that it can be decomposed into two sub-models, each of which is a simple logit (see Train 1986.) In the case at hand, the nested logit model of choice between the three options A, B, and C, can be expressed as two simple logit models, as follows. First, the customer has a choice of whether to implement the measure, that is, whether or not to choose option C. If the customer chooses to implement (i.e., not choose option C), then the customer has a choice of whether to apply for and receive a rebate; this is a choice between option A and option B, made by customers who choose to implement (i.e., do not choose option C.) Sequential estimation exploits this decomposition by estimating each sub-model separately. This procedure is easier computationally and can be performed on standard software for logit estimation (unlike simultaneous estimation which requires special software.) However, sequential estimation is less efficient than simultaneous estimation, since only part of the full information is being used in estimation of each sub-model.

Separate models were estimated for groups of measures, with different groups for commercial and industrial customers. For commercial customers, three models were estimated, one each for: HVAC measures, lighting measures, and other hardware measures. For industrial customers, three models were estimated: for lighting measures, process measures, and other hardware measures. Table 3 gives the model for lighting measures for commercial customers. The notation in the tables is the same as that for PG&E's model for lighting conversions. In particular, variables that enter the utility of options A and B affect the probability that the customer implements the measure, and variables that enter only option A affect whether the customer applies for and receives a rebate. The variables and their estimated coefficients are self-explanatory from the table. Models for the other measures are provided by Pacific Consulting Services (1993a,b).

Table 4 presents the estimated net-to-gross ratios for each measure group. The estimates seem reasonable; however, there are no other measures of net-to-gross ratios for the SCE program against which to compare.

Conclusions

The method presented in this paper can be used to estimate net savings from programs, like rebate programs, in which a participant necessarily implemented a targeted measure. The method has the advantage of utilizing data on customers from the energy utility's territory; that is, it does not require the use of a control area for comparison. There are serious limitations of the method, however, that need to be considered in any application. First, the method does not estimate spillover effects. For some programs, especially large programs that have been operating for a long time, spillover can be very important. Second, the method requires data that are somewhat difficult to obtain. In particular, data are required for a sample of customers on what measures each customer could have implemented, the cost and savings of each of these measures, and which of the measures the customer did implement. Third, the estimates of net-to-gross ratios are not as precise as one would perhaps want. Confidence intervals were calculated for the estimated ratio for PG&E's program, using the method described in the appendix. The 80% confidence interval is 0.43 to 0.98. Furthermore, this confidence interval does not indicate the full degree of uncertainty about the estimates, since the confidence interval reflects the variance due to sampling only. As is true in any econometric analysis, changes in model specification can change the resulting estimates.

Despite the uncertainty surrounding the estimates, it is not clear that they are less precise than those from other methods that appropriately account for the fact that customers choose to participate in programs and that this choice is related to their decision of whether or not to implement energy-efficient measures. (Methods that do not account for the fact that participation is voluntary can appear to be more accurate. However, these methods give a false sense of precision. In particular, by erroneously

Table 3. Nested Logit Model for Lighting Measures for SCE's Commercial Sector Hardware Rebate Program

Submodel for choice of whether to implement (option C versus options A and B):

Variable (entering the utility of options A and B)	Estimated coefficient	T-statistics
1. Dummy for: offices	0.19	1.36
2. Dummy for: retail stores	0.45	2.50
3. Dummy for: universities	1.07	2.97
4. Dummy for: medical facilities	0.77	2.03
5. Dummy for: hotels	0.44	1.57
6. Dummy for: misc. services	0.38	2.24
7. Square footage group	0.04	2.00
8. Dummy for: customer expanded in 1990	0.74	3.08
9. Measure of whether customer plans to relocate	0.05	2.50
10. Customer's ranking of SCE's helpfulness	0.04	1.33
11. Dummy for: customer received an audit in 1989-91	0.54	4.50
12. Dummy for: customer reports that SCE is its prime source of information about energy efficiency	0.33	2.75
13. Dummy for: customer located in Valley region	0.45	1.73
14. Dummy for: customer located in Coastal region	0.46	2.88
15. Dummy for: customer located in Inland region	0.18	0.90
16. Constant	-2.99	11.9
$1-\lambda$	1.03	3.12

Submodel for choice of whether to receive rebate, given that measure is implemented:

Variable (entering the utility of option A)	Estimated coefficient	T-statistics
17. Dummy for: customer operates 24 hrs/day, 7 days/week	0.24	0.73
18. Dummy for: offices	0.55	1.57
19. Dummy for: retail stores	0.71	1.82
20. Dummy for: warehouses	0.41	0.87
21. Dummy for: schools	1.16	2.58
22. Dummy for: universities	0.23	0.34
23. Dummy for: hospitals	1.25	1.76
24. Dummy for: hotels	1.38	2.82
25. Dummy for: misc. services	0.99	2.83
26. Customer's ranking of its familiarity with SCE's rebate program	0.11	2.20
27. Dummy for: building built after 1983	-0.31	1.55
28. Dummy for: annual usage between 60,000 and 334,000 kwh	-0.06	0.22
29. Dummy for: annual usage between 334,000 and 1.674 million kwh	-0.27	0.96
30. Dummy for: annual usage over 1.674 million kwh	-0.70	2.12
31. Dummy for: customer located in Valley region	0.12	0.26
32. Dummy for: customer located in Coastal region	0.16	0.55
33. Dummy for: customer located in Inland region	0.69	2.09
34. Constant	-1.39	3.02

assuming that participation is exogenous to the implementation decision, the methods appear more precise while actually being less precise. Train, 1994, discusses this issue.) The estimation of net savings is an inherently difficult task, for which highly precise estimates cannot reasonably be expected. The method in this paper provides another modeling strategy than analysts can consider to use, either alone or in parallel with other methods, in trying to tackle this difficult issue.

Appendix: Calculation of Approximate Confidence Intervals

Recall that the net-to-gross ratio for any measure is calculated as the ratio of sums of estimated probabilities, and the net-to-gross ratio for the program as a whole is the weighted average of these measure-specific ratios. The distribution of the estimated ratio for each measure

Table 4. Estimated Net-to-Gross Ratios for SCE's Hardware Rebate Program

Commercial customers	
HVAC measures	83%
Lighting measures	72%
Other hardware measures	80%
Industrial customers	
Lighting measures	48%
Process measures	29%
Other hardware measures	53%
Weighted average	77%

(i) depends on the covariance matrix for all the parameters in the model for that measure, since the probabilities depend on all these parameters, (ii) depends on this covariance matrix in a highly nonlinear way, since the net-to-gross ratio is the ratio two sums of probabilities, and (iii) is not normal, even when the sample is large enough for the parameters to be normally distributed. The distribution of the estimated ratio for the program as a whole depends on the distribution for each of the measure-specific ratios. To simplify the calculations, an approximate distribution was calculated for the program ratio. The method is based on the fact that, for each measure, λ is the critical parameter in determining the net-to-gross ratio for that measure, and that Monte Carlo methods can approximate a true distribution. The confidence intervals were calculated in three steps. First, an approximate distribution for each measure's estimated ratio was calculated. The distribution of λ around its estimated value was taken to be normal (since the sample was large) with variance equal to the square of the standard error. Random draws were taken from this distribution, with each draw providing a value of λ from its distribution. For each draw (i.e., for each value of λ), the net-to-gross ratio for the measure was calculated. This process provided a distribution of net-to-gross ratios arising from the distribution of λ. Second, the approximate distribution of the program ratio was calculated. The program ratio is the weighted sum of the measure ratios. One random draw was taken from the distribution of each measure ratio. The weighted average of these draws was taken to obtain a value of the program ratio. This process was repeated for many sets of draws, producing many values of the program ratio. These values constitute the approximate distribution of the program ratio. Third, confidence intervals for the ratios were tabulated on the basis of this approximate distribution. In particular, the 80% confidence interval was obtained by removing the highest 10% and the lowest 10% of the values; the confidence interval is from the lowest of the remaining values to the highest of the remaining values. (Stated more succinctly, the 80% confidence interval is the 10% quantile break up to the 90% quantile break.)

These approximate confidence intervals differ from the true intervals because (i) only the variance in λ is utilized rather than the variance-covariance for all the parameters in the choice model for each measure, and (ii) the Monte Carlo method is not exact for any finite number of draws. It is important to note, however, that the approximate confidence intervals are not necessarily smaller than the true ones simply because they are based on the variance of λ only. Since there is a covariance of λ with each other parameter, using the entire covariance matrix can result in either larger or smaller confidence intervals.

Endnotes

1. Often the gross savings of the program are known (or have been estimated through other procedures that are considered reliable.) Inevitably, the gross savings estimated from the nested logit model (i.e., GS) will not equal this known gross savings, which we denote GS*. The ratio GS*/GS gives the proportion by which the actual gross savings differ from the model's estimate. This information can be used to correct the model's estimates of total savings, naturally occurring savings, and net savings by multiplying TS, OS and NS by this ratio. This correction will not, of course, affect the estimated net-to-gross ratio.

2. These variables are probably endogenous to the decision of whether to implement and receive a rebate. In particular, customers who are likely to implement a measure are probably more likely to become aware of the program. Inclusion of these variables was necessary, however, to obtain convergence when estimating λ. The difficulty of estimating λ, and its instability with respect to different specifications, is perhaps the most important limitation of this approach to estimating net savings.

   Customers who are not aware of the program do not, in some sense, actually have the option of receiving a rebate. Alternative models were estimated in which these customers were considered to face a choice between options B and C. The resulting net-to-gross ratio does not change substantially with this change in specification. See Cambridge Systematic (1993a) for details.

3. With multiple methods, a customer is a participant if it chooses option A for at least one measure. A

customer that chooses option B or C for one measure might still be a participant if it chooses option A for another measure. Because of this, P(B/A,B,C)/[P(B/A,B,C)+P(C/A<B<C)] is not the share of nonparticipants that choose option B over option C, since the customers choosing option B or C for this measure might be a participant due to another measure.

4. Spillover s the impact of the program on customers' decisions to take measures without a rebate. For example, publicity about a rebate program might induce a customer to install an energy-efficient measure, even though the customer does not bother to apply for the rebate. Or, the rebate program might induce dealers to stock high-efficiency equipment; a customer might purchase this equipment without knowing that a rebate is available, or without bothering to apply. In the context of the nested logit model, spillover means that a customer who chose option B would, without the program, have chosen option C. For the nested logit models described above, the impact of the program is estimated by eliminating option A; therefore any customer who chose option B with the program is necessarily predicted to choose option B without the program. While the model sin this paper do not capture spillover, nested logit model of this form can capture some types of spillover, depending on the variables that enter the model. Train and Paquette (1994) provide an example estimated on data from eight New York state utilities.

References

Cambridge Systematic, Inc., 1993a, *Net-to-Gross Ratios for PG&E's CIA Rebate Program: Study D–Decision Analysis Model*, report to Pacific Gas and Electric Company.

Cambridge Systematic, Inc., 1993b, *Net-to-Gross Ratios for PG&E's CIA Rebate Program*, report to Pacific Gas and Electric Company.

McFadden, D., 1978, "Modeling the Choice of Residential Location, " in A. Karquist, et al. (eds.), *Spatial Interaction Theory and Planning Models*, North-Holland Publishing Company.

Pacific Consulting Services, 1993a, 1990 *Southern California Edison Energy Management Services and Hardware Rebate Program Evaluation, Vol. 6, Estimating Net KWH and KW Reductions in the Commercial Sector*, Chapter 3.

Pacific Consulting Services, 1993b, 1990 *Southern California Edison Energy Management Services and Hardware Rebate Program Evaluation, Vol. 7: Estimating Gross and Net kWh Savings and kW Reductions in the Industrial and Agricultural Sectors Using Conditional Demand Analysis*, Chapter 7.

Train, K., 1994, "Estimation of Net Savings from Energy Conservation Programs," *Energy: The International Journal*, Vol. 19, No. 4, pp. 423-441.

Train, K., 1986, *Qualitative Choice Analysis*, MIT Press.

Train, K., and E. Paquette, 1994, "Rebate Programs: Discrete Choice Methods to Estimate Free-Ridership, Net-to-Gross Ratios, and the Effect of Program Advertising, forthcoming, *Journal of the Association of Demand Side Management Professionals.*