Solar Photovoltaic Program Final Impact Evaluation

For

Energy Trust of Oregon

By

David Robison, P. E.

Stellar Processes

Stellar Processes, Inc.

1033 SW Yamhill Suite 405

Portland, OR 97205

(503) 827-8336

April 25, 2007

Comments, clarifications, corrections and suggestions are welcome.

Stellar Processes is a company of consulting engineers specializing in energy economics, measurement and verification. Our expertise includes monitoring and commissioning large facilities as well as diagnostic evaluation of small buildings. Projects include:

- Impact evaluation of monitored savings from improvements to small commercial HVAC systems.
- Evaluation of savings from residential weatherization and heating duct sealing for a Northwest public utility.
- Development of a simplified tool for billing analysis and simulation to diagnose problems, estimate savings and provide performance-based contracting.
- Evaluation of large building commissioning program for New England utility.
- Verification of savings and market analysis from computerized Energy Management of large buildings in the Midwest.
- Development and analysis of measurement and verification for Energy Service Company performance contracts.
- Testing of solar water heater systems

Table of Contents

List of Figures

List of Tables

Executive Summary

This study reviewed consumption records (utility bills) for participants in the Energy Trust solar photovoltaic (PV) program. The goal was to determine if the expected reduction in electricity occurred for participants.

The methodology of reviewing utility bills is subject to some uncertainty. In many cases, the home is under construction prior to the solar installation so the energy consumption is not typical. During the analysis procedure, we attempted to remove cases with atypical records but we are not certain that we have accounted for all such projects.

The solar installations also included a dedicated meter to record electricity production with reading taken by the participants. These meter readings are probably a better measure of solar production and have been normalized to reflect annual energy generation. Major conclusions are:

- PV readings taken by the participants appear to be the best measure of actual solar production. This generation averaged 3,176 kWh or 99% of expected.
- PV readings agree reasonably well with expected energy production. This suggests that the installations are generally performing as contractors promised.
- Net kWh production to the grid as estimated from billing data does not appear to differ by climate zone. This is somewhat surprising but a small difference may be masked by variability in the observations.
- Analysis of consumption records (utility bills) suggests less reduction in total energy consumption than expected. The regression estimate of consumption change averages 2,409 kWh or 73% of generation. While the reduction is statistically significant, it is not clear if it represents consumer “takeback” or inconsistent data reflecting partial occupancy prior to the solar installation.
- A comparison group showed small downward trend that was not statistically significant. However, if the regression estimate of energy change is derated for this trend, the net energy production would average 2,154 kWh or 65% of expected.

Energy Trust Staff Response to Report Findings and Conclusions

Energy Trust staff is pleased with the study's finding that the installed PV systems, on average, are performing as expected. The study does indicate the possibility that regional differences in expected and actual performance might exist. These differences may be due to the engineering model or its inputs or to the presence of systematic regional factors

Energy Trust plans on review the engineering models and their inputs for accuracy and to perform more detailed analysis of the installed system database to see if regional differences exist and if they are ongoing. Energy Trust already communicates with installers if, and when, specific system performance issues are identified.

Though the study did not conclusively identify the level of "takeback" present, it is clear to Energy Trust staff that this issue warrants further research as it may impact forecasts of electricity delivered to the gird from residential solar electric systems. Energy Trust plans to monitor this "takeback" to see if it persists over time for participants and with new participants.

To support the research into system performance and participant behavior over time Energy Trust plans on contact all participants on an annual basis to gather metered kWh production data as well as PV system and demographic data. Participants will be able to provide this information via mail or through a web-based survey.

Introduction

Program History

The Energy Trust of Oregon (ETO) in 2002 initiated incentives for on-site solar photovoltaic electricity production. Participants usually install systems that are connected to the electricity grid and "net metered". That means that when electricity is produced by the photovoltaic system, it flows to the grid and drives the participant's electric meter backward. The goal of this study was to compare participant's electricity bills pre- and post-installation to determine if the expected reduction in electricity billing occurred. In addition, we also compared results to electricity production that was separately metered with the meter read by the participant.

The essence of this program was to achieve electricity reduction by installation of solar photovoltaic panels. For purposes of this study, we reviewed participants that were listed as completed by May 12, 2006.

Energy Usage Analysis

Documentation Of Analysis Methodology

Studies of this sort frequently use the PRInceton Scorekeeping Method (PRISM) to weather normalize the consumption data and to remove 'noise' due to the weather which could influence consumption before or after treatment. The difference between consumption before and after treatment represents the energy savings due to the program. PRISM uses regression to separate the electric bill into a baseload and seasonal, weather-dependent component. This output is used in the analysis to determine whether there have been changes in the baseload or weather dependent component after treatment. The weather dependent component is assumed to represent primarily space heating. (In fact, it includes a small amount of seasonally dependent energy for increased water heating and lighting consumption during winter months.) Space heating occurs only below a balance temperature, which is unique to each home. The balance temperature depends on the thermal integrity of the house, the preferred thermostat setting of the customer and other behavioral factors. During the summer, when there is no space heating, consumption includes only a baseload component. The sum of both baseload and weather-dependent components provides the Normal Annual Consumption (NAC). NAC is the typical total annual energy consumption during a "normal" weather year.

In this analysis, we utilized PRISM as a first step, but quickly realized that a modified regression is necessary. For this program, we expect that solar electricity production will peak during summer months – which will interfere with isolating the baseload. However, the heating slope and balance temperature should not change as a result of the treatment. Thus, we established a pooled regression technique in which we used the same baseload and balance temperature but included post-installation solar irradiation as an additional variable. This model is illustrated in Figure 1.

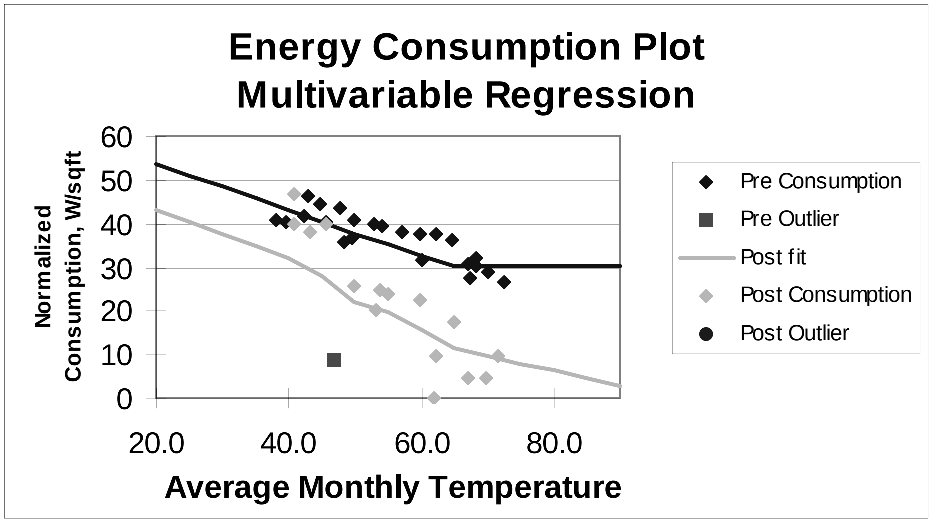

Figure 1. Example of Regression Model

One might wonder if the solar irradiation on the tilted solar panels is different than the horizontal solar irradiation used as an independent variable in the regression model. Indeed, radiation is different on a tilted surface. However, the radiation bears a linear relationship to the horizontal irradiation. Thus, the horizontal irradiation is suitable as a variable for linear regression. The “tilt effect” will be included in the regression coefficient estimated in the regression model. The same is true of the size of the solar array. Obviously a larger solar array will result in a larger solar regression coefficient.

The regression procedure identifies data points that may be outliers due to partial vacancy or other data errors. We flagged as outliers points that differ from the model by two standard errors or more. Such points may be removed from the analysis if it appears warranted or retained if they appear to be correct. Pooling of the pre- and post-retrofit data is useful because we had fewer data observations representing the post-retrofit period. With a pooled model, we are better able to utilize post-retrofit data that cover less than a full year period.

For purposes of the analysis, we do not care about the regression coefficients per se – we are only using those coefficients to compute the weather normalized Annual Consumption (NAC) for the pre- and post-installation years. It is the NAC, or rather the difference between the two NACs, that is of interest. The important criterion for analysis is that the regression model provides a relatively good fit to the individual data points. The usual metric to quantify model fit is the R2 of the regression. In this study, R2 for the analysis group ranged from 40% to 99% with an average of 72%.

Many participants, especially in South West Oregon, have noticeable cooling consumption. For these cases, we also found it necessary to develop a regression model that includes a cooling as well as heating. That is, the model includes a cooling balance temperature and cooling slope similar to that used to model heating consumption. An example of this regression model is shown in Figure 2.

(Notice that in both these figures, the line represents the modeled performance relative to the temperature variable. However, the plot does not show the full impact of solar irradiation. Thus, the fit of the regression model to the data points may be better than is suggested by the modeled line. This is particularly the case for the cold weather post-installation points in Figure 2.)

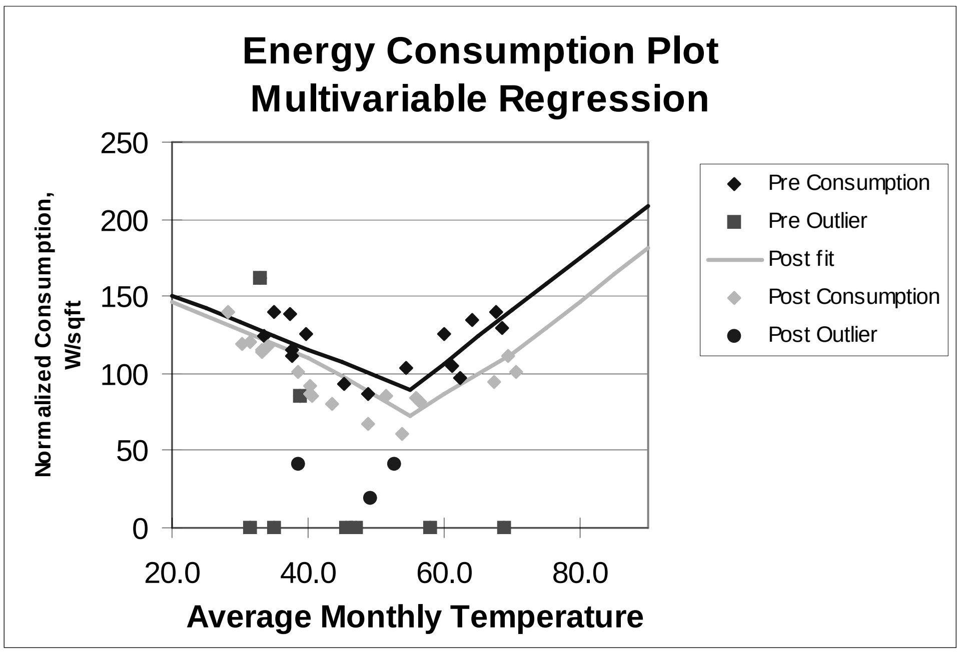

Figure 2. Example of Regression Model with Cooling

A frequent problem with regression studies is that a large portion of the study population must be removed due to the inability to form a successful regression model. Such problems might be caused by gaps in the data, vacancies or changes in the occupant's behavior. In this case, we reviewed the data carefully in order to preserve as many cases as possible. Each case was individually reviewed for outlier observations. Often these changes could be identified as periods of partial vacancy due to vacations. In these cases, the regression was computed with outliers removed.

The sample disposition describes the initial procedures used to eliminate cases from the study. Pre and post periods were labeled, and the month during which the home was treated was not included in either the pre or post periods.

Daily temperature data provided by National Oceanic and Atmospheric Administration (NOAA) were used to build the temperature file used by the regression. NOAA was also the source for long-term average weather data used to weather-normalize the data.

The University of Oregon Solar Radiation Laboratory provided solar irradiation data. Unfortunately, solar data were not always available for all sites during the study period. In particular, solar monitoring was discontinued for the Medford location, which was best suited to many sites in South West Oregon. For these sites, we utilized weather data from Klamath Falls location for both the pre- and post-installation period. It is important that the same sort of weather data be used for the pre- and post-installation periods. However, it is not so important that the data be from a close location. The regression procedure will correct for differences due to microclimate assuming that such differences are linear, which is reasonable to assume.

Programmatic impact on consumption was evaluated using a traditional quasi-experimental design. The design compares the participants to a similar but untreated group. In a true experimental design, members would be assigned randomly to either the treated or the comparison group.[1] This approach is not possible for an actual program where interested customers are allowed to participate. Hence, the design is considered “quasi” experimental. The non-participants were drawn from a pool of future program participants. This minimizes the possibility of any self-selection bias that may mark the participants as being different from typical customers. Use of future participants offers another benefit since site characteristics collected in a later year can be applied to the comparison group in an earlier year. Regression analysis of the comparison group (control group of future participants) followed the same procedures developed from the analysis of the participant group.

The analysis uses a standard pre/post cross sectional consumption (billing) analysis. The weather normalized annual consumption (NAC) before the treatment establishes a baseline, which can then be compared to weather normalized consumption after the treatment. The difference in consumption determines gross changes. That is:

Gross changes = NAC (pre) - NAC (post)

Gross changes are determined for the comparison group in the same way. The participant changes are corrected for any consumption change apparent in the comparison group. The result is net changes attributable to the power generated by the solar panels. This difference of differences approach is traditionally used in DSM evaluation to “net out” savings due only to the treatment.[2] Results are reported in terms of the average kWh reduction per participant.

---

[1] Cook, Thomas and Campbell, Donald*, Quasi-Experimentation, Design and Analysis Issues for Field Settings*, Houghton Mifflin Co., 1979. Campbell, Donald and Stanley, Julian, *Experimentation and Quasi-Experimental Design for Research*, Houghton Mifflin Co., 1963.

[2] Fels, M. *The Princeton Scorekeeping Method: An Introduction*, Princeton University, Center for Energy and Engineering, Princeton, NJ, PU/CEES 163. Fels, M., *Special Issue Devoted to*

Group Characteristics

Those cases included in the impact study matched well the location distribution of all program participants. An exception occurred with the comparison group. Due to limited consumption records, the comparison group was more likely to be located in the Willamette Valley.

Table 1. Program Participant Locations

Weather Location	Program Participant Fraction	Study Group Fraction	PV Reading Fraction	Comparison Group Fraction
Willamette Valley, Coast	33%	36%	35%	88%
South West Oregon	33%	31%	33%	8%
Central Oregon	34%	32%	33%	4%

Sample Disposition

The intent of the evaluation was to analyze consumption for all participants. Due to incomplete data and occupant changes, a full sample is never possible. For this study, we also observed some rural participants with multiple meters and inexplicable changes in the data available from the different meters. Such cases were not usable. In some other cases, the facility was obviously of recent construction and there were no comparable pre-installation data. These cases were also not usable. A serious problem with this study is that we do not have assurance that the pre- and post- installation data truly represent a similar level of occupancy by the participants.

We had sufficient consumption data to analyze approximately 40% of program participants as shown in Table 2. Not all these cases are acceptable for impact review. The data analysis procedure identified 26 cases with obviously different usage due to new construction or other significant change, 10 cases that were commercial facilities, and another 9 cases where data varied too widely for a reasonable regression fit. The final result was 118 cases for which the regression technique appears to successfully model the participant's consumption records. Of those cases, 83 also had a meter reading of total kWh produced by the PV panels that had been recorded by the participant.

Measuring Energy Savings: "The Scorekeeping Approach", Energy and Building, 9(1-2), Feb/Mar 1986.

Table 2. Summary of Sample Disposition

	Number of Participants
Total Participants, Complete by 5/12/06	277
Data Analysis	163
Acceptable Cases	118
With PV Read	83
Comparison Group	26

Participants Attrition Table

277 cases in program files (245 are residential)

⇓

163 cases with consumption data for analysis

⇓

26 cases dropped – utility records show occupancy change
10 cases dropped -- commercial facilities
9 cases dropped – unable to achieve a good regression model

⇓

Review and preservation of cases
118 cases acceptable for review

Comparison Group

Future participants were chosen from those who participated the following years, in 2005-6. In that way, the future participants would have a full year of “post” data to correspond to the participant’s year of post data. For this comparison group, we chose a “participation” date of December 2004. One would like to be confident that the comparison group is, in fact, representative of the (treated) participant population. Figure 3 compares the pre-installation energy consumption of both groups and shows reasonable agreement.

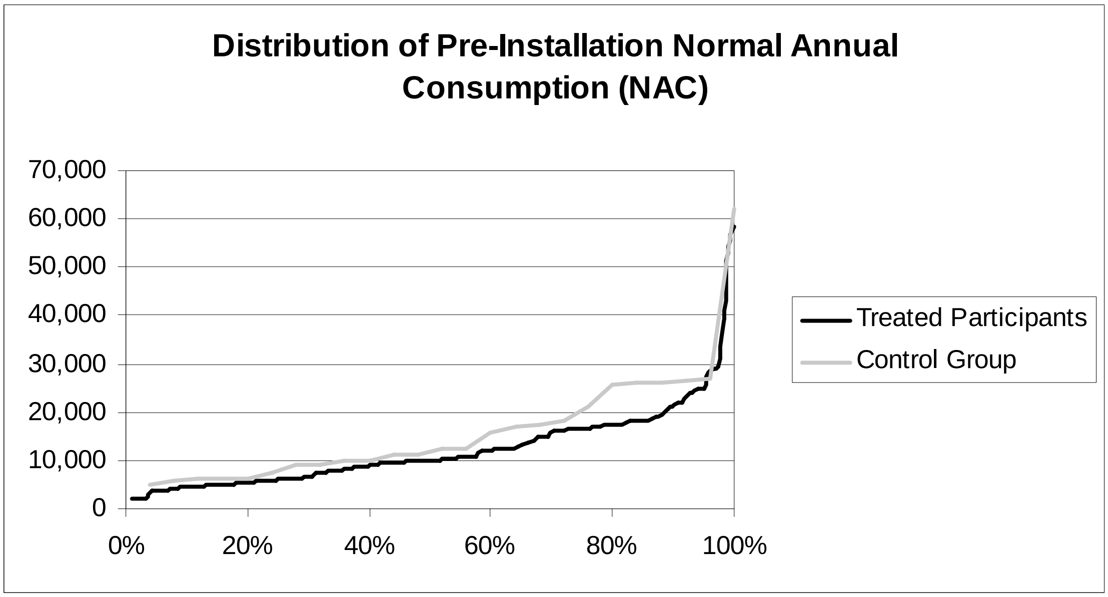

Figure 3. Comparative Distribution of Consumption

Overall Program Results And Conclusions

Energy Impacts Results

For the participant sites, we derived the energy reduction estimates in Table 3. As is apparent, the NAC reduction derived from regression analysis fall short of the expected change in energy consumption. The comparison group shows a change in consumption that appears as a small amount of reduction, so the shortfall is even larger. This leads to the question of whether participants “took back” some of their solar generation by increasing their electricity consumption. If so, the results are not necessarily a response to the solar installation. As discussed previously, there is still uncertainty about whether the pre-installation data represents homes that were fully occupied.

Table 3. Participant Analysis Group Energy Reduction Estimate (118 Cases)

Variable	Mean	Standard Deviation	Standard Error	90% CL Lower	90% CL Higher	t-test	Significance (1-tailed)
Pre-Installation NAC, kWh/year	12,380	8,966	825				
Post-Installation NAC, kWh/year	9,971	8,247	9,971				
Change in NAC, kWh/year	2,409	1,707	157	2,148	2,669	1.282	0.032
Expected Savings, kWh/year	3,318	923	85	3,177	3,459		

Table 4. Comparison Group Energy Reduction (26 Cases)

Variable	Mean	Standard Deviation	Standard Error	90% CL Lower	90% CL Higher	t-test	Significance (2-tailed)
Pre-Installation NAC, kWh/year	16,215	12,084	2,417				
Post-Installation NAC, kWh/year	15,881	11,920	2,384				
Saved NAC, kWh/year	334	1,034	207	-20	688	1.615	0.119

The reduction estimate is statistically significant for the treatment group, although it just barely fails to be significant for the comparison group.

A difficulty in viewing the reduction is that the customer’s behavior is not necessarily consistent from year to year. Changes in family size or consumption habits interfere with direct observation of the savings. One approach to minimize the effect of behavioral “noise” is to observe the distribution of pre- and post-retrofit NAC shown in Figure 4. In this graph, individual cases are sorted by consumption level. This compensates for random variation in behavioral consumption. The resulting plot shows a distinction between the pre- and post-retrofit distributions of consumption. The lower line for “post” shows that there is a reduction relative to total consumption.

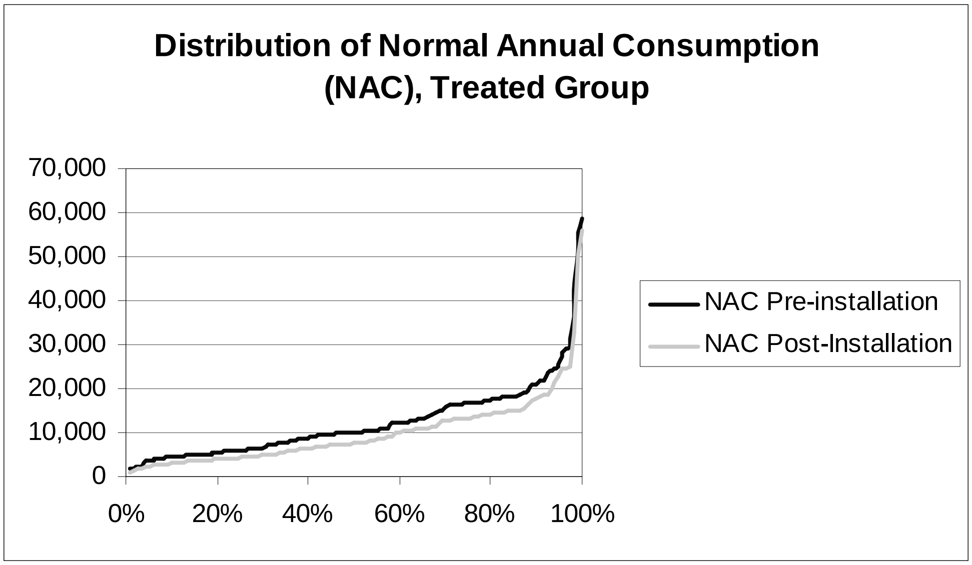

Figure 4. Pre/Post Distribution Treatment Group

A similar treatment of the comparison group is shown in Figure 5. Note that the pre- and post-retrofit consumption is almost identical for this group. That is, there is only a very small reduction for the comparison group.

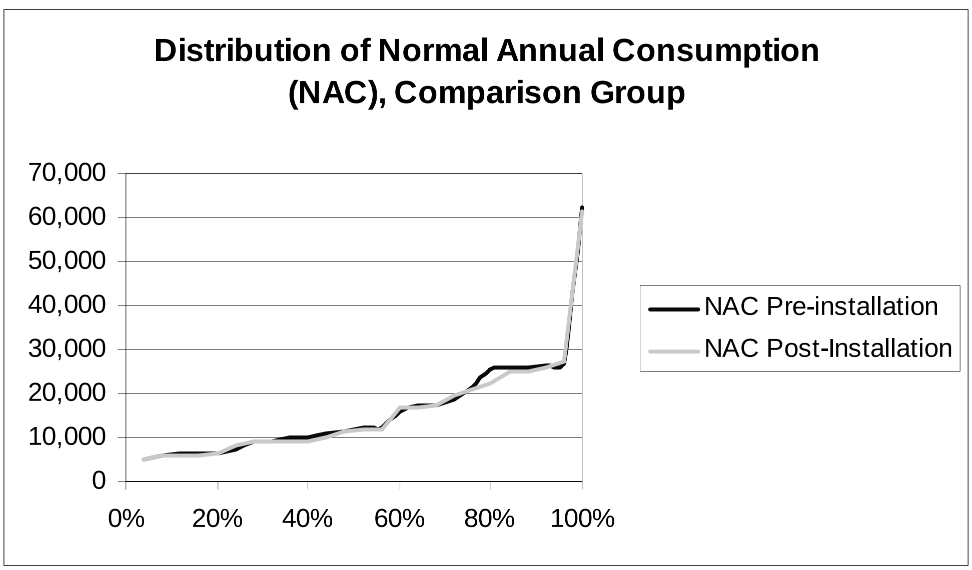

Figure 5. NAC Distribution Treatment/Comparison Groups

The comparison between these two groups is even more apparent in Figure 6. This chart shows the distribution of changes in consumption for treated and comparison groups. As expected, the comparison group has a mean saving of almost zero and the treated group is consistently larger.

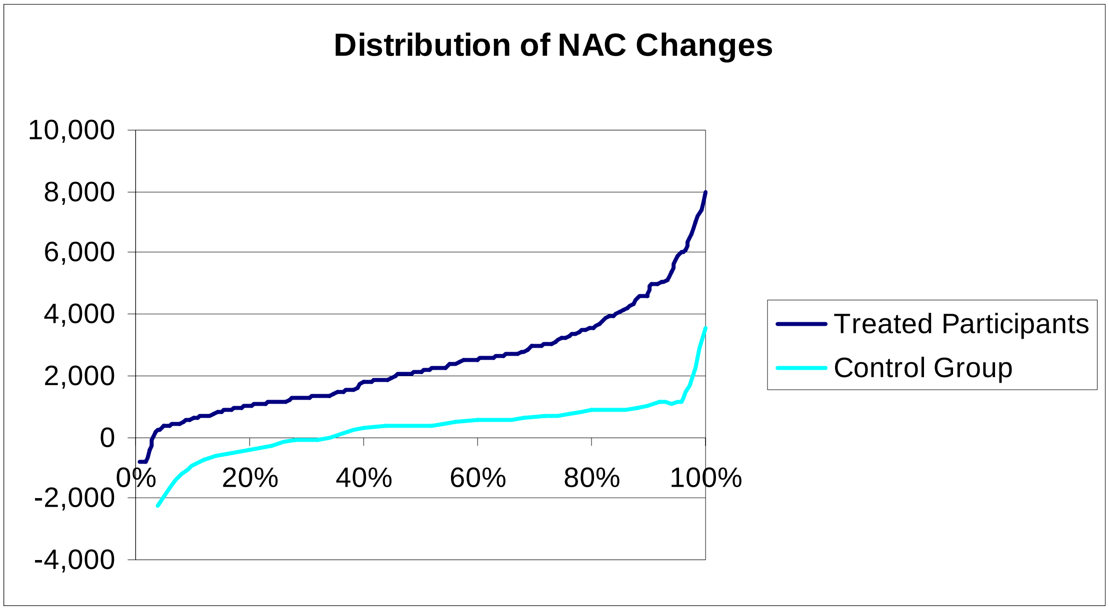

Figure 6. Distribution of Energy Reduction Compared

Since the comparison group uses slightly more energy than the program participants, one could adjust the trend by the ratio of consumption. The trend amount averages 2.1% of consumption or 255 kWh when applied to the participant group. This value was used for estimating net NAC energy reduction.

Solar Energy Production by Climate Zone

Surprisingly, there appears to be no significant difference between climate zones (ANOVA sig = .480 is not significant) . Table 5 shows similar reduction despite the fact that solar irradiation should be 20-30% better outside the Willamette Valley. One might conclude that installers have tended to compensate for less favorable solar climate by increasing the size of the solar array. However, that is not the case. Program data shows no significant difference in the installed array size – the average is about 2800 Watts in all zones. Expected energy production is significantly larger in South West Oregon which should lead to somewhat larger solar production in that zone.

Table 5. Climate Zone Comparisons

Climate Zone	Mean NAC Reduction kWh /year	Number	Standard Deviation	Standard Error	90% CL Lower	90% CL Higher
Willamette Valley, Coast	2,576	43	1,777	271	2,029	3,123
South West Oregon	2,496	37	1,644	270	1,948	3,044
Central Oregon	2,135	38	1,699	276	1,577	2,694

One can explore this observation with the sub-group that has both PV Readings and NAC Analysis. As shown in Table 6, the difference lies in the NAC analysis. The Expected Generation shows that the Willamette Valley climate should provide less opportunity for energy production. According to the PV readings, the generation in that climate was slightly better than expected. (The Realization Rate is the ratio of observed to expected.) However, the NAC Analysis Change shows little difference between the climate zones. Zone differences may be buried within the high variability of the NAC Analysis methodology.

Table 6. Production Estimates by Climate Zone, Average kWh/ Year

	Number	Expected Generation	PV Reading	PV Reading Realization Rate	NAC Analysis Change	NAC Analysis Realization Rate
Willamette Valley, Coast	27	2,766	2,935	106%	2,279	82%
South West Oregon	29	3,764	3,476	92%	2,707	72%
Central Oregon	27	3,007	3,096	103%	1,961	65%
All Zones	83	3,193	3,176	99%	2,325	73%

Comparison to PV Reading Results

As mentioned, participants installed a separate electricity meter on the solar system and independently recorded their solar production. These readings did not always occur on the basis of one calendar year. Trust staff corrected the reported readings to normalize for a one-year period. PV readings were available for only a portion of the study group but the results are interesting. Figure 7 compares expected energy production to PV readings. (Expected production is computed by the installer and reported to the customer.) The 45-degree line is a reference line, not a regression through the data points. If the two estimates agree with each other, one expects the data points to be aligned along the 45-degree line – as is the case in Figure 7. This suggests that the expected production is a good indicator of the energy reduction participants actually achieve. For comparison, Figure 8 shows less agreement between expected and NAC energy reduction.

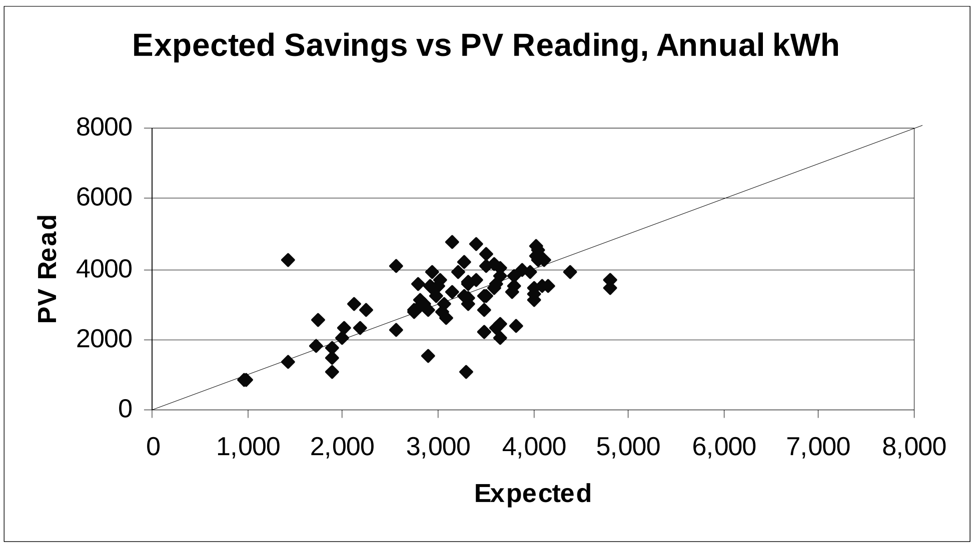

Figure 7. Expected Energy Reduction and PV Readings

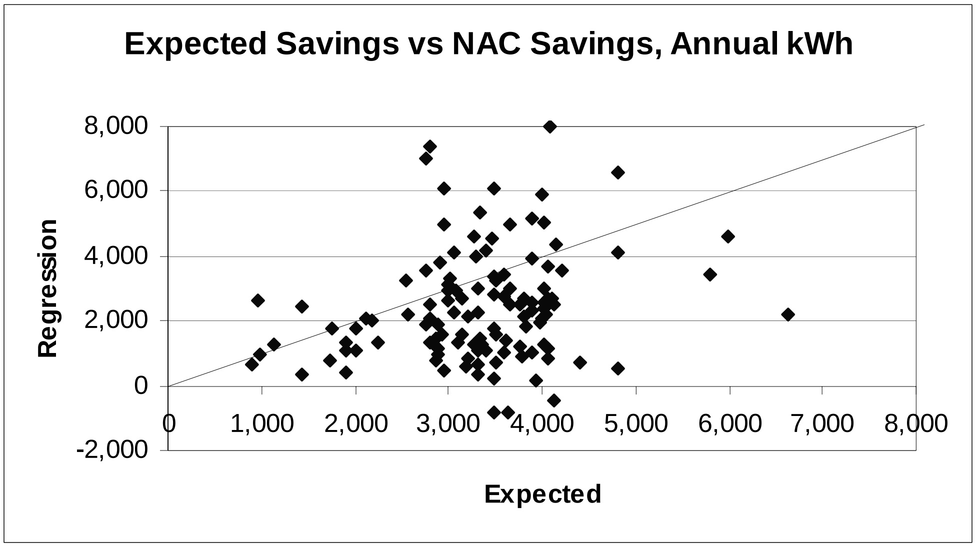

Figure 8. Expected and Regression NAC Energy Reduction

Table 7 compares NAC reduction estimates to both the expected and PV reading estimates of energy production. Net NAC reduction has been derated for the changes observed in the comparison group, even though these changes are not statistically significant. The expected production and PV readings agree closely. This suggests that the PV readings, since they were recorded directly from the solar array, are the best estimate of actual solar production. Figure 9 compares these different estimates of energy production.

Table 7. PV Reading Group Production Estimates (83 Cases)

Variable	Mean	Standard Deviation	Sample Count	90% CL Lower	90% CL Higher
Control Trend	334	1,034	25	-20	688
NAC Reduction	2,325	1,640	83	2,026	2,625
Net NAC Reduction	2,070	274	83	1,615	2,526
Expected Production	3,193	826	83	3,042	3,344
PV Read Production	3,176	927	83	3,007	3,346

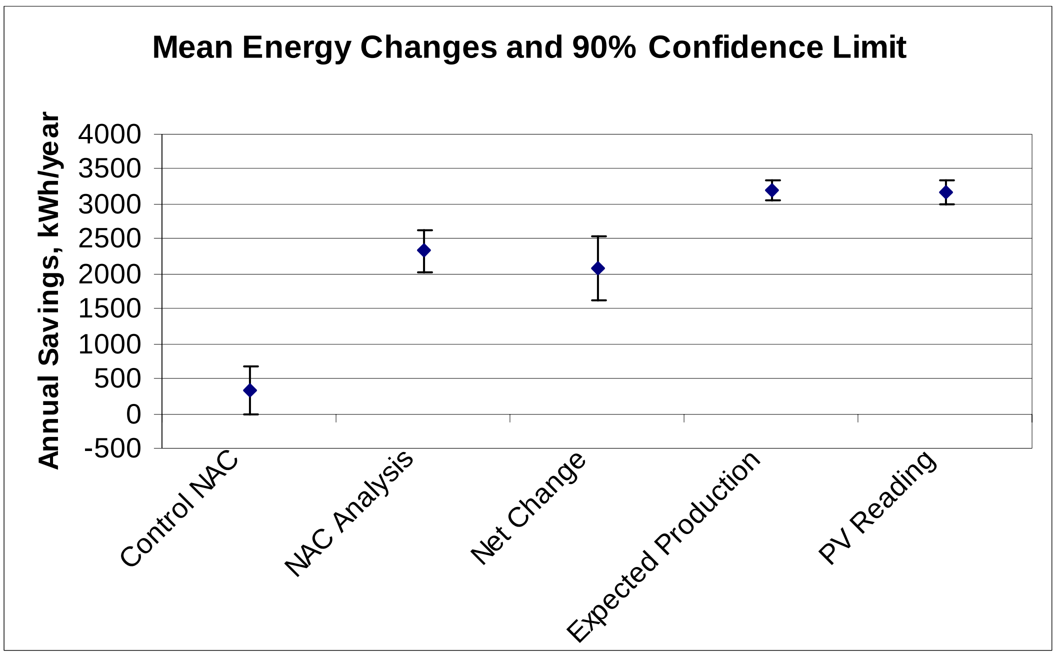

Figure 9. Different Production Estimates within PV Reading Group

The shortfall of NAC reduction compared to expected is a significant difference (z-test = 2.50, sig= 0.012 one-tailed) suggesting that “take back” may occur. Depending on whether one chooses to adjust for the comparison group change, the takeback ranges from 1185 to 851 annual kWh. The question of causality is not addressed at all by this study. Did participants increase their electricity consumption because they knew that solar power would offset their increase? Or is the apparent difference due to the fact that pre-installation consumption often reflects partial occupancy that we were unable to weed out during the analysis?

There is no clear answer without an additional survey. It is possible that customers tend to increase energy consumption, figuring that they no longer need to conserve since they are generating. If so, that change should be taken into account for planning purposes. However, given the small change and the uncertainty of full occupancy, it is not clear that a conclusion is possible.

Program Impact Results

One can apply impact results from the groups analyzed to extrapolate to the remaining cases. First, we computed a “realization rate” defined as the impact result divided by the expected energy production. These realization rates were computed for both the PV reading and NAC estimates of reduction. In neither case were there statistically significant differences by climate zone. The mean realization rates were 99% for PV readings and 73% for NAC estimates.

Table 8 shows the mean reduction for each sub-group of program participants. First, the group with PV readings has mean production computed as the actual PV readings and NAC reduction estimates. If one chooses to subtract the comparison group savings, the net NAC impact is reduced. Program energy production is then the product of the mean production times the number in each group.

The next group is the one for which regression analysis was possible. For those cases without PV readings, an estimated result was extrapolated as the average realization rate times the expected production. For those cases with PV readings, the actual reading was used. This group also has actual NAC reduction estimates.

The final group is the full set of all residential projects completed by May 12, 2006. Even though PV readings and NAC analysis were not possible for all cases, we compute an extrapolated value where data are lacking. The final result is an estimate of overall program impact.

As discussed previously, the PV Read Production is probably a good measure of actual solar production. The NAC reduction is statistically significant but could be affected by data uncertainties. The Net NAC reduction follows methodological rigor although the adjustment for the comparison group is not statistically significant.

Table 8. Extrapolated Impact Results

Program Group	Number in Sample	Mean Production			Program Production		
		PV Read Production	NAC Reduction	Net NAC Reduction	PV Read Production	NAC Reduction	Net NAC Reduction
PV Reading Group	83	3,176	2,325	2,070	263,634	193,004	171,840
Analysis Group	118	3,301	2,409	2,154	389,472	284,233	254,145
Full Residential Group	245	3,405	2,485	2,230	834,145	608,752	546,281

Conclusions

We propose the following major conclusions:

- PV readings taken by the participants appear to be the best measure of actual solar production. Energy production averaged 3,176 kWh or 99% of expected.
- PV readings agree reasonably well with expected energy production. This suggests that the installations are generally performing as contractors promised.
- Net kWh production to the grid as estimated from billing data does not appear to differ by climate zone. This is somewhat surprising but a small difference may be masked by variability in the observations.
- Analysis of consumption records (utility bills) suggests slightly reduced net kWh production to grid. The regression estimate of energy reduction averages 2,409 kWh or 73% of expected energy production. While the difference is statistically significant, it is not clear if it represents consumer “takeback” or inconsistent data reflecting partial occupancy prior to the solar installation.
- A comparison group showed small energy reduction trend that was not statistically significant. However, if the regression estimate is derated for this trend, the net energy production would average 2,154 kWh or 65% of expected.