

Estimating Energy Savings Resulting from Strategic Energy Management Programs: Methodology Comparison

Andrew Bernath, The Cadmus Group, Portland, OR
Maggie Buffum, The Cadmus Group, Portland, OR

ABSTRACT

Strategic energy management (SEM) programs seek to implement long-lasting, comprehensive changes in energy management and consumption at industrial facilities.¹ Evaluators typically estimate energy savings associated with SEM programs using regression analysis to model facility energy consumption as a function of weather and other variables. They then use the regression models to compare energy consumption before and after program implementation. Evaluators can choose among several regression model frameworks, with the accuracy of results depending on the model selected. Therefore, selecting the right framework proves integral to completing an accurate, robust evaluation. Testing frameworks to determine those best suited to evaluate savings for a particular facility (and under a specific set of circumstances) presents a challenge as true energy savings value at a facility remain unknown. Using all available frameworks, evaluators can estimate savings and compare results, conjecturing about how well the frameworks estimate savings, but they cannot compare the estimates to a true energy savings value to determine, with certainty, which works best.²

By using simulated data, evaluators can begin answering questions about which regression models perform best and under what circumstances. If simulating facility energy consumption data using a function of weather and other known variables (and because they generate it), evaluators can determine “true” savings and savings drivers underlying the model. In this study, the authors use simulated energy consumption, based on models representing industrial facilities that participate in SEM, to answer two important questions: Which regression analysis framework should evaluators use to calculate accurate, robust savings estimates? How should the facility regression model be specified within that framework?

The simulation revealed several important findings:

- The forecast and fully specified pre/post model frameworks produce unbiased estimates and capture true savings at the nominal 80% confidence level under most scenarios— for both simple and complex facility specifications.
- For facilities with simple model specifications, the simple pre/post model framework performs comparably to the forecast and fully specified pre/post models, with slightly elevated mean absolute percentage errors. For facilities with complex model specifications, the simple pre/post framework proves unreliable, consistently producing biased savings and failing to reach nominal capture rates.
- All model specifications tend to produce biased estimates with poor capture rates when omitting variables from the model.
- For a scenario where an event affecting energy consumption occurs during the post period (where an indicator variable can be included in a pre/post model, but no estimate is available to adjust savings in a forecast model) the forecast model produces biased savings with low capture rates, while fully specified pre/post produces unbiased savings estimates and high capture rates.
- On average, savings are slightly biased for a scenario where the regression model specification does not account for autocorrelation, but this produces poor capture rates for true savings.

¹ Though commercial facilities also implement SEM, this research focused on industrial SEM evaluation.

² Evaluators with the true savings estimate clearly do not need to apply a regression model.

Introduction

A rapidly growing approach to energy efficiency, strategic energy management (SEM) programs realize savings by using capital in concert with operations and management improvements to facilities. By using various frameworks, evaluators of these programs can estimate reductions in energy consumption before and after program implementation. These frameworks typically involve facility-level regression models, but they use different model specifications and underlying assumptions. Industrial facility models pose challenges due to large numbers of variables that effect energy consumption and relatively little data available to capture those relationships. Typically, they include complex interactions between energy consumption drivers and sometimes include nonroutine adjustments before or after an SEM program begins, which affect energy consumption but are not related to the program. Nonroutine adjustments may or may not include engineering estimates of energy savings. Furthermore, determining which regression framework and model specifications to use in an evaluation can be challenging as no industry-standard produces results with proven accuracy when including a variety of possible scenarios.

In this study, the authors used a simulation approach to test different frameworks and model specifications. We sought to shed light on those that produce accurate, robust results (and under which conditions), providing guidance for future evaluations. A simulation approach proved particularly effective for this purpose as we determined “true” savings and model facility energy consumption, based on this savings value and on other variables driving energy consumption. As we knew savings underlying the simulated data, we could compare true savings to estimates resulting from different regression model frameworks, thus determining—on average—the estimates’ accuracy. We assessed accuracy and robustness for the following three regression frameworks, described in detail below:

- Forecast
- Simple Pre/Post
- Fully-Specified Pre/Post

Further, in our evaluations of actual facility data, we observed that results tended to vary greatly, depending on weather and facility conditions, including the following:

- Weather (e.g., heating degree day [HDD], cooling degree day [CDD], mean temperature)
- Production and occupancy
- Weekday versus weekend activity
- Shutdown or closure periods
- Nonroutine adjustments

Finally, due to differing data collection protocols and procedures, evaluators may not have access to data for certain variables, and some data may exhibit autocorrelation or non-normality that the regressions does not include. Therefore, we considered evaluations that did or did not include the following model specifications:

- Extraneous variables included or known drivers of energy consumption omitted
- Correctly or incorrectly specified relationships between variables and energy consumption
- Nonroutine adjustments included or not included in the regression model
- Accounted for or did not account for time-series autocorrelation or non-normality in the regression analysis

The remainder of this paper provides details on the regression frameworks, weather and facility conditions, and additional considerations that we examined using the simulation study. By summarizing the

findings, we identify trends regarding how accurately each framework estimates energy savings under different circumstances. Ultimately, we provide answers to two important questions:

- Which regression analysis framework should evaluators use to calculate accurate savings estimates?
- How should the facility regression model be specified within that framework?

Methodology

Evaluators typically estimate SEM program energy savings by comparing energy consumption prior to implementing the SEM program (i.e., the baseline) to energy consumption after implementing the SEM program. Building baseline models only requires use of pre-program period data and represents a facility's energy consumption in the program's absence. Still, in modeling the baseline, evaluators typically require data on several possible energy drivers:

- Weather (e.g., HDD, CDD, mean temperature)
- Production and occupancy
- Weekday versus weekend activities
- Shutdown or closure periods
- Timing of nonroutine adjustments

In many facilities, some or all of these variables correlate with energy consumption. In others, however, they may not. Using simulated data, we built several baseline models to test misspecification's effects.

Upon establishing the baseline model specification, evaluators model post-program period energy consumption to find the difference before and after program implementation, and then calculate savings. We explored three common frameworks for doing so. In one, we used pre-program energy consumption models to forecast energy consumption in the post-program period, absent the SEM program's effects. We calculated savings as the difference between predicted baseline usage and the metered usage. The second framework used the baseline model specification, in addition to a post-program indicator. This allowed us to estimate savings based on the coefficient corresponding to the indicator, representing average energy savings per time interval. The third also used the baseline model, but it included effects from other variables estimated to capture pre-program and post-program effects (e.g., HDD effects can vary with changes in production or production may interact with program effects). These frameworks can be summarized as follows:

1. **Forecast:** Use a baseline regression model to predict energy consumption in the post-program period, absent the program. Sum the differences between predicted usage and metered usage to estimate total savings during the post-program period.
2. **Simple Pre/Post:** Use a baseline regression model, based on pre-program energy consumption and predictor variables. Estimate the model using both pre- and post-program period data, with an additional indicator signaling the post-program period's beginning. Use the post-program period indicator's coefficient to estimate average energy savings per time interval (e.g., day, week, month; one or more of these may be an option, depending on the data frequency). Multiply the average by the number of time intervals in the post-program period to estimate total energy savings during the post-program period.
3. **Fully-specified Pre/Post:** Similar to the simple pre/post framework, use a baseline regression model with both pre- and post-program data; in this framework, however, include interaction terms with the post-program indicator and all predictor variables. The interactions allow predictors to produce different effects on energy consumption in the pre- and post-program periods. Estimate the model, take the post-program period indicator's (i.e., main effect) coefficient, multiplied by the number of post-program time periods, add each coefficient of the post-program period interactions, multiplied

by the sums of their respective variable values during the post-program period, and estimate total energy savings during the post-program period.

Simulated Data

The authors simulated two data sets to represent facilities similar to those observed in SEM program evaluations. One data set represented a facility with energy consumption as a function of one production process, an indicator of production interruptions, and weather (specifically HDD). We called this the “simple facility,” with energy consumption driven by two variables without interactions. The second data set represented a facility with energy consumption as a function of two different production processes, an indicator of production interruptions, weather (HDD), the interaction of weather and production, and a nonroutine adjustment (i.e., an “event”) occurring in the pre-program period, resulting in reduced energy consumption. This reflected a “complex facility,” where a number of variables drove energy consumption and energy drivers interacted. Both models included a nonroutine event in the post-program period, with an engineering estimate associated with energy consumption changes. Such events occur frequently in industrial facilities. Examples include installing new equipment at a facility, closing part of a facility temporarily or permanently, or staffing changes.

Error! Reference source not found. provides “true” baseline models for simple and complex facilities (pre/post models also include a post-period indicator and interactions between the post-period indicator and each other model variable):

Simple Facility: $kWh_t = \beta_0 + \beta_1 Production_t + \beta_2 Production\ Interruptions_t + \beta_3 HDD_t + \varepsilon_t$

Complex Facility: $kWh_t = \beta_0 + \beta_1 Production1_t + \beta_2 Production2_t + \beta_3 Production\ Interruptions_t + \beta_4 HDD_t + \beta_5 HDD_t \times Production1_t + \beta_6 Pre-Program\ Event_t + \varepsilon_t$

Where:

- β_i = The coefficient of the i^{th} variable in the model ($i = 0$ represents the model intercept)
- kWh_t = Energy consumption at the facility at time t
- $Production_t$ = Production at the facility at time t
- $Production\ Interruptions_t$ = An indicator of production interruptions at time t
- HDD_t = HDDs at time t
- $HDD_t \times Production_t$ = The interaction between HDD and production at time t
- $Pre-Program\ Event_t$ = An indicator representing a nonprogram-related change in a facility’s energy consumption at time t

We specified the error term using random draws from a normal probability distribution, with variance specified in a number of ways. This allowed us to study how accurately each regression framework estimated savings in the face of small and large variations, in addition to autocorrelation and heteroscedasticity, and it allowed us to generate numerous simulated datasets for each facility type, thus testing how the framework performed on average, given random data variations.

The production and weather variables used in this simulation were constructed so that they were uncorrelated. There is a slight correlation between production 1 and the non-production indicator. Figure 1 visualizes the relationships between energy drivers and energy consumption in the baseline period for a complex facility.




Figure 1. Scatterplot Matrix of Variable Correlations for a Simple Facility.

The simple facility consumes energy in a similar manner, but correlates only with a single production variable, a non-production variable, and HDD.

Testing Models

After simulating the facility data, the authors applied the above-outlined regression frameworks to determine those producing accurate results and performing best for each facility. We considered scenarios where data were unavailable for one or more production variables, data were missing for the non-routine adjustment, data were available for an extraneous variable that did not drive energy consumption, and models where the error terms did or did not account for heteroscedasticity and serial correlation. In summary, we tested each regression framework using regression model specifications with the following characteristics or cases:

1. Correctly specified
2. Missing post-program period event data
3. Omitted weather variable(s)
4. Omitted production2 variable (complex facility only)
5. Extraneous variable (CDD)
6. Heteroscedastic error
7. Serial correlation in errors
8. Known post-program period event without engineering estimates (i.e., changes to facility consumption resulting from the event cannot be separated from consumption changes due to the program in the forecast framework). This scenario violates compliance with American Society of

Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) Guideline 14-2002³ and International Performance Measurement and Verification Protocol (IPMVP) Option C⁴ whole building modeling standards, it occurs repeatedly in practice, and therefore proves relevant to the model framework comparison.

Table 1 summarizes each case by indicating predictor variables included in the evaluation regression model: black dots indicate that the regression model included the variable. Some variables (and case four) do not apply for a simple facility (cells greyed out in the table).

Table 1. Energy Drivers Used for Each Test Scenario

Case Number	Energy Drivers Included in Evaluation Regression Model								
	Production 1	Production 2	Interruptions to Production	HDD	Extraneous Variable (CDD)	HDD x Production 1	Event in Baseline	Event In Post-Period	
								Event Savings Indicated (Pre/Post)	Event Savings Subtracted (Forecast)
Simple Facility									
1	•		•	•					
2	•		•	•				•	•
3	•		•						
4									
5	•		•	•					
6	•		•	•	•				
7	•		•	•					
8	•		•	•				•	
Complex Facility									
1	•	•	•	•		•	•		
2	•	•	•	•		•	•		•
3	•	•	•				•		
4	•		•	•		•	•		
5	•	•	•	•	•	•	•		
6	•	•	•	•		•	•		
7	•	•	•	•		•	•		
8	•	•	•	•		•	•		•

Using these 15 total cases to estimate savings in each regression framework produced 45 sets of results, each of which included savings estimations and measurements for testing the model’s accuracy for 10,000 simulated datasets, generated for each facility type and error specification. A findings summary follows.

Findings

The authors fit regression models according to each of the 15 cases described above and in all three regression frameworks. This produced 45 sets of results for each of 10,000 simulated data sets. The results corresponding to each data set included an estimate of energy savings and an 80% confidence interval. We

³ ASHRAE Guideline 14-2002. *Measurement of Energy and Demand Savings*. American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Inc. 2002. www.ashrae.org

⁴ International Performance Measurement and Verification Protocol: Concepts and Options for Determining Energy and Water Savings. 2012. Available online: <http://www.coned.com/energyefficiency/PDF/EVO%20-%20IPMVP%202012.pdf>

compared the confidence interval to the “true” savings value to determine whether it included the true savings value, and then we estimated a coverage rate by tallying the number of datasets where this proved true.

For example, if the confidence interval around the savings estimate included true savings in 8,000 of 10,000 datasets, we concluded the method had a 80% coverage rate. Figure 2 illustrates this concept using 100 savings estimates and their associated confidence intervals for the correctly specified complex facility forecast regression framework. The horizontal axis represents true savings, each dot represents one savings estimate, and the lines extending from the dots represent 80% confidence intervals for the respective savings estimate.

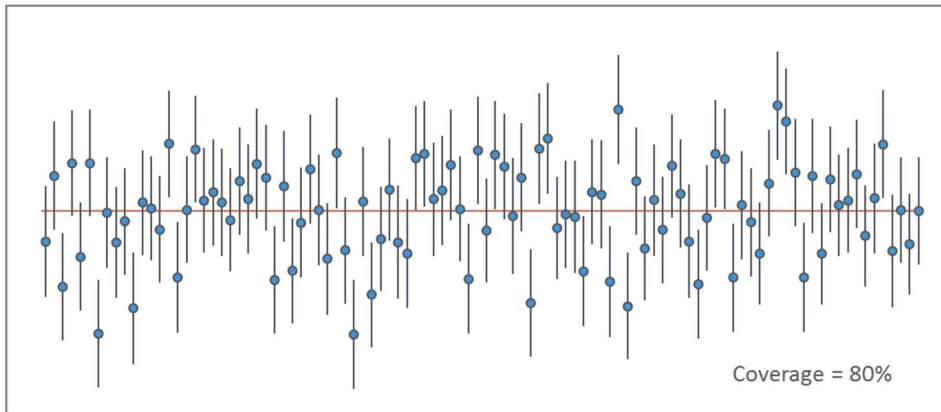


Figure 2. Coverage Plot for 100 Observed Complex Facility Savings Estimates Using the Forecast Regression Framework

With 80% confidence intervals, we expected coverage rates close to 80%, but note how high or low the rates are under the various modeling specifications or misspecifications defined by cases. Coverage provides an important indication of the accuracy of results from each framework: the closer the framework comes to 80% coverage, the more accurate its estimates.

Table 2 summarizes findings that provide the true savings value, average estimated savings from each framework, and 80% confidence interval coverage.

Table 2. Estimated Savings and Confidence Interval Coverage

Case Number	True Savings (MWh)	Estimated Savings (MWh)			80% Confidence Interval Coverage		
		Forecast	Simple Pre/Post	Fully-Specified Pre/Post	Forecast	Simple Pre/Post	Fully-Specified Pre/Post
Simple Facility							
1	6,849	6,847	6,697	6,847	79%	63%	80%
2	6,849	7,814	7,683	7,815	0%	0%	0%
3	6,849	6,738	6,601	6,735	71%	39%	70%
4	N/A	N/A	N/A	N/A	N/A	N/A	N/A
5	6,849	6,851	6,700	6,851	80%	62%	80%
6	6,849	6,852	6,702	6,852	79%	70%	79%
7	6,849	6,848	6,698	6,848	41%	42%	41%
8	6,849	7,814	6,919	6,850	0%	81%	80%
Complex Facility							
1	7,615	7,617	5,874	7,617	79.8%	0.0%	79.9%
2	7,615	9,279	7,590	9,277	0.0%	96.3%	0.0%
3	7,615	7,598	7,547	7,598	79.9%	98.9%	99.2%
4	7,615	6,742	5,047	6,745	66.6%	0.0%	51.2%
5	7,615	7,616	5,858	7,616	79.3%	0.0%	79.6%
6	7,615	7,620	5,876	7,620	79.1%	0.1%	79.8%
7	7,615	7,612	5,866	7,612	42.1%	5.0%	42.3%
8	7,615	9,279	6,123	7,618	0.0%	0.1%	80.4%

These results indicate that the fully specified pre/post framework provides coverage fairly close to 80% more frequently than the other two frameworks. None of the frameworks produces good coverage in Cases 2 or 7 for the facilities, and the forecast and simple pre/post models do not produce good coverage for case 4 for the complex facility. Case 2 represents a scenario with missing event data (and hence not included in the regression model). Case 7 represents a scenario where serial correlation exists but is unaccounted for. Case 4 represents the scenario (in the complex facility) with one production variable missing or omitted from the model. Except for the simple pre/post model in the simple facility, coverage rates are at or above the nominal level for Case 3, where the HDD variable is missing or omitted from the model. Based on these findings, we conclude that none of the frameworks that the study examined reliably produced accurate savings estimates when omitting variables or not accounting for serial correlation.

We also investigated the bias in savings estimation (i.e., the difference between the estimated and “true” savings) by computing the mean absolute percentage error (MAPE) and the median percentage error. The MAPE told us the average magnitude of the estimation bias, and the median percentage error told us whether the model tended to overpredict or underpredict savings. Table 3 summarizes these results.

Table 3. Mean Absolute Percentage Error and Median Percentage Error

Case Number	MAPE			Median Percentage Error		
	Forecast	Simple Pre/Post	Fully-Specified Pre/Post	Forecast	Simple Pre/Post	Fully-Specified Pre/Post
Simple Facility						
1	1.6%	2.5%	1.6%	0.0%	-2.2%	0.0%
2	14.1%	12.2%	14.1%	14.1%	12.2%	14.1%
3	2.1%	3.7%	2.1%	-1.6%	-3.6%	-1.7%
4	N/A	N/A	N/A	N/A	N/A	N/A
5	1.6%	2.5%	1.6%	0.0%	-2.2%	0.0%
6	2.3%	2.9%	2.3%	0.0%	-2.2%	0.0%
7	4.5%	4.9%	4.5%	0.1%	-2.1%	0.0%
8	14.1%	2.1%	2.0%	14.1%	1.0%	0.0%
Complex Facility						
1	2.7%	22.9%	2.7%	0.0%	-22.9%	0.0%
2	21.9%	2.7%	21.8%	21.9%	-0.3%	21.8%
3	2.6%	2.7%	2.6%	-0.2%	-0.9%	-0.2%
4	11.5%	33.7%	11.4%	-11.5%	-33.7%	-11.4%
5	2.7%	23.1%	2.7%	0.0%	-23.1%	0.0%
6	3.5%	22.8%	3.5%	0.1%	-22.8%	0.1%
7	7.5%	23.0%	7.5%	-0.1%	-23.1%	-0.1%
8	21.9%	19.6%	3.0%	21.9%	-19.6%	0.1%

These results indicate, on average, absolute percentage error is typically within 5% of the true savings for most scenarios. Similar to the confidence interval coverage results, missing or omitted variables and autocorrelation not accounted for tends to elevate the absolute percentage error. In most cases, the error did not display a tendency to be biased consistently in the same direction. Missing or omitted production variables led to overestimation in the complex facility, both for the forecast and the fully specified pre/post regression frameworks. Figure 3 provides the MAPE for the simple and complex facilities, and compares values for the three savings estimation methods.

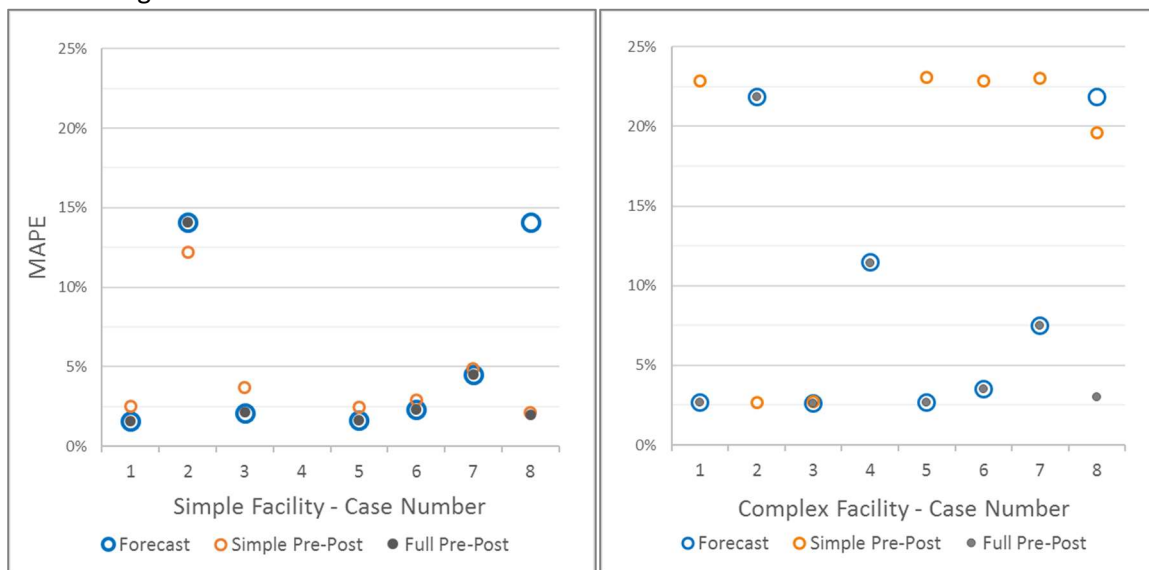


Figure 3. MAPE for Each Facility and Regression Framework

Consistent with the capture rates, the simple pre/post method typically produced larger MAPE values around estimated savings for the complex facility. Additionally, median percentage errors in the complex facility

tended to underestimate savings in all but the omitted post-program period and weather variable scenarios for the simple pre/post framework. This provides further evidence that the simple pre/post regression framework produces unreliable and biased savings estimates.

Two scenarios served to violate the typical regression assumption that model error would be distributed normally. The first included a heteroscedastic error (i.e., variance increased with production). The second included serial correlation in the model error as autoregressive order 2. Figure 4 visualizes the heteroscedastic model error in the complex facility, and Figure 5 depicts the model error for the first 365 days in the complex facility data.

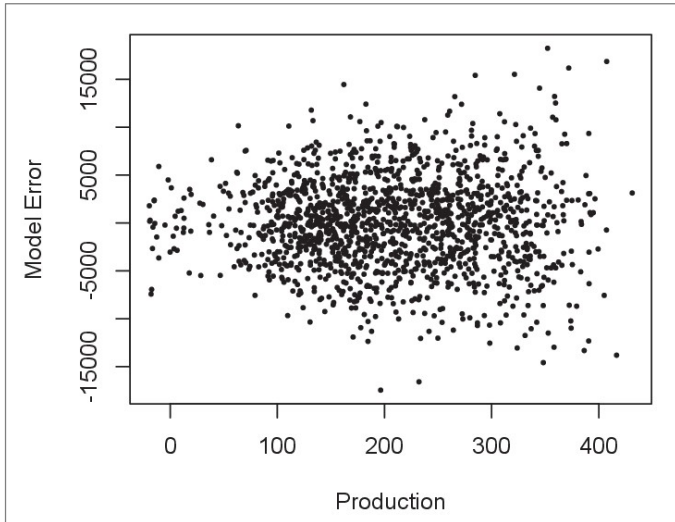


Figure 4. Heteroscedastic Model Error Plotted Versus Production

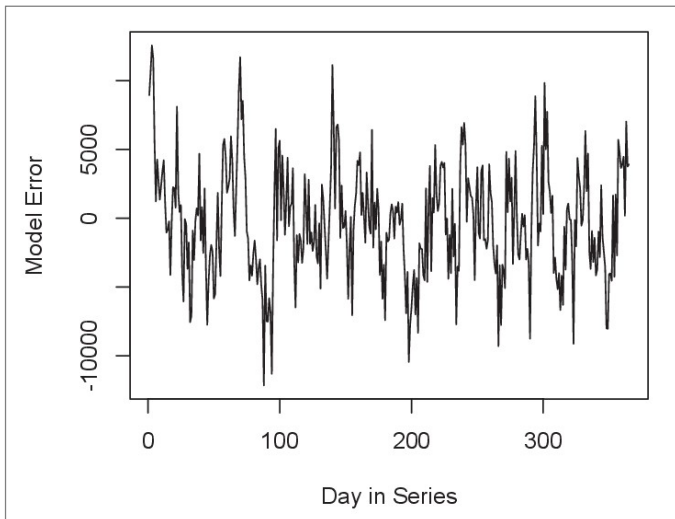


Figure 5. Serial Correlation in the Model Errors for the First 365 Days in the Complex Facility

Despite violation of the normality assumption, the three regression frameworks produced unbiased estimates of energy savings. In both scenarios, standard errors were larger for models with non-normally distributed errors. Additionally, the scenario with a missing or omitted production variable in the complex facility resulted in increased CV's Table 4 shows the coefficient of variation (CV) for each energy savings estimate, calculated as the margin of error for an estimate (at 80% confidence) divided by the estimate itself. A larger CV implies increased error for a particular estimate. All CV's are within the tolerance levels recommended by ASHRAE (< 20%).

Table 1. CV for Simple and Complex Facilities

Case Number	CV—Simple Facility			CV—Complex Facility		
	Forecast	Simple Pre/Post	Fully-Specified Pre/Post	Forecast	Simple Pre/Post	Fully-Specified Pre/Post
1	2.0%	2.3%	2.0%	3.3%	6.9%	3.4%
2	1.8%	2.0%	1.8%	2.7%	5.4%	2.9%
3	2.2%	2.5%	2.2%	3.3%	6.9%	6.8%
4	N/A	N/A	N/A	11.3%	14.9%	10.2%
5	2.0%	2.3%	2.0%	3.4%	6.9%	3.4%
6	2.8%	3.1%	2.9%	4.3%	7.8%	4.4%
7	2.4%	2.7%	2.4%	4.1%	7.6%	4.1%
8	1.8%	2.7%	2.4%	2.7%	7.5%	3.8%

Conclusions

Regardless of the framework used to estimate savings, we find omitting significant energy drivers from regression models (i.e., post-program period events, production, and weather) tends to result in inaccurate program savings estimates and 80% confidence intervals that do not actually include the true savings value 80% of the time. The fully specified pre/post framework provides coverage fairly close to 80% in more cases than the other two frameworks. The simple pre/post framework typically fails to capture true savings at least 80% of the time for all model specifications using both simple and complex facility data, even when correctly specifying the model.

On average, savings estimates for scenarios where the regression model specification omits important energy drivers are biased compared to cases where models include the necessary (or even extraneous) energy drivers. In these scenarios, the bias does not consistently over- or underpredict. Rather, the bias direction depends on the omitted variable. Additionally, some variable omissions resulted in increased standard errors for savings estimates.

When the model's error is not distributed normally, the savings estimates are unbiased with only slight increases to standard errors. When the model's error exhibits unaccounted for serial correlation, the savings estimates are slightly biased with increased standard errors. Savings estimates result in low coverage rates where regression model specifications do not account for serial correlation.

We draw several important conclusions from the simulation results.

Evaluators should ensure they include all key energy drivers in their regression models. For most scenarios and regression frameworks, missing or omitted key energy drivers produced biased estimates, increased standard errors, and failed to capture true savings with target confidence levels. In addition, consequences did not appear for including extraneous parameters; so overspecification of the energy model may not be a concern. The primary takeaway from this is that evaluators should attempt to identify all key energy drivers for SEM evaluations.

Forecast and fully specified pre/post models result in accurate, unbiased estimates under most scenarios. For both simple and complex facilities, the forecast and fully specified pre/post methods appear to produce very similar estimates and capture rates. These frameworks produced unbiased estimates (i.e., less than 5% of the "true" savings) for four of seven scenarios in the complex facility data and for five of the six scenarios in the simple facility data, outperforming the simple pre/post regression frameworks. These frameworks even produced an unbiased result in one of the three variable omission scenarios for both facilities. Note that for this comparison to be fair, we are excluding the case where a post-program event is captured by pre/post models and not the forecast model. This case will be discussed separately. Consequently, when non-routine adjustments are unnecessary, we recommend evaluators use the forecast or fully specified pre/post as robust models for savings estimation.

The unreliable simple pre/post model should not be used to estimate facility energy savings. The simple pre/post model performed consistently with the forecast and the fully specified pre/post model for simple facility data, but failed to produce accurate, unbiased estimates with the complex facility data. Additionally, the simple pre/post almost never captured true savings at the target capture rate, even when the model was correctly specified. In practice, industrial facilities tend to be complex, with many potential energy drivers and interactions. Hence, an actual facility probably will not behave like the simple facility modeled in this study, and we do not recommend use of the simple pre/post regression framework.

Evaluators should investigate serial correlation and attempt to model it in their chosen regression framework. As many industrial facilities have seasonal or cyclic production, data from these facilities likely contains some degree of serial correlation. This particularly holds true with higher-frequency data (e.g., daily intervals). Our results suggest that failing to account for this autocorrelation can greatly reduce the chance of capturing true savings parameters in the estimated confidence interval.

Only the fully specified pre/post regression framework should be used when an event or nonroutine adjustment occurs in the post-program period and when an estimate of energy reduction or increase does not exist. As evaluators, we have encountered several cases where a facility experienced a significant change during the post-program period. In some cases (e.g., new equipment installations or upgrades), an engineering estimate will be available that quantifies the energy consumption change resulting from the installation. For many situations, however, no estimate will be available for the expected increase or reduction in energy consumption resulting from a change to facility energy consumption. Examples encountered include facility layoffs, temporary building closures, management changes, or equipment breakdowns. In this simulated scenario, the forecast framework overestimated savings by including the event's reduction in consumption in total savings attributed to the SEM program. Conversely, the fully specified pre/post framework accounted for the consumption reduction through an added indicator variable, resulting in unbiased savings estimates.