Standard Approach to Non-Standard Industrial Projects

Kevin Warren, Warren Energy Engineering, West Grove, PA
Carter Membrino, Warren Energy Engineering, West Grove, PA

ABSTRACT

Energy savings from custom large-scale industrial projects can be difficult to quantify. The savings may depend on the level of production or other factors, such as ambient air temperature. Savings can also be influenced by hard-to-quantify variables such as product mix, number of operating lines, or the number of personnel shifts. This paper discusses a systematic, albeit customizable, approach to evaluating industrial custom projects. The method was first developed for evaluating compressed air projects but it also applies to process cooling and other types of measures. In many cases, ex-ante savings for such projects were calculated using methods that either did not account for the impact of production, or did so in ways that were imprecise.

The method involves collecting data on the power use of the affected the system, the output of the affected system, and the facility production. To accurately estimate savings, the authors have found that four relationships should be quantified: the dependence of pre-installation supply side efficiency and post-installation supply side efficiency on load and of pre-installation demand side efficiency and post-installation demand side efficiency on production or other independent variables. Impact evaluators commonly have access to short-term post-installation data and sometimes have access to short-term pre-installation data. Short-term data of a week or two is typically sufficient to establish the required relationships.

The method is an alternative to the common approach of tying the dependence of pre-installation and post-installation energy use directly to production or other independent variables and the related concepts of "energy performance" and "energy use intensity." If you tie energy use directly to production, it impossible to determine if the change in energy use was due to a change in the efficiency of the affected system or a drop in system load. For projects that impact only the supply or demand side of the system, using the proposed method enables the evaluator to control for the unaffected parameter.

The Challenge

Evaluation of industrial energy efficiency projects are among the most difficult challenges for impact evaluators. Though there are some common measure categories which may be seen again and again, nearly every facility and every project is unique. Similar measures at different sites may need to be evaluated using different methods because the availability and quality of data are highly variable. The drivers of energy use may be difficult to determine. Data may be available only with a low degree of granularity, and customers can be reluctant to provide production information they consider proprietary. Evaluators can apply the standard four IPMVP approach Options to these types of projects, with varying degrees of success.

Evaluators first look to apply either Option A: Key Parameter Measurement or Option B: All Parameter Measurement. For the evaluator wanting to use data to determine the savings achieved by a project, the only viable method is typically to use some form of relatively short-term data collection and then apply some rationale for annualizing the data. Production level commonly has a significant impact on industrial system energy use, so methods need to account for this dependence. The paper focuses on one recommended Option A/B method that we use as our default starting point for all industrial projects. It relies on characterizing the supply side efficiency (SSE) and demand side efficiency (DSE) of the system separately and annualizing the results based on typical production levels. The method is sometimes criticized as being overly complex, so we delve into some of the common issues with the alternative methods.

IPMVP Option C: Billing Analysis as it is conventionally applied is typically not a useful method for evaluating energy efficiency projects in industrial facilities. Conventional Option C involves the comparison of whole-facility monthly energy consumption versus one or more independent variables, such as ambient temperature, degree days, or production output. The conventional rule of thumb is that the technique should be applied only when energy savings are expected to be at least 10% of total metered energy use. In our experience, it is very rare for incentivized projects in industrial facilities to achieve savings that are a high percentage of total facility consumption, particularly if there is only one meter for the site. Where the customer has installed sub-metering it is more likely that the measure savings will be greater than 10% of the annual sub-metered usage, and the conventional Option C approach may be applied. Newer applications of Option C utilizing daily or hourly interval data are able to reliably estimate savings that are below the conventional 10% threshold, but they have primarily been tested on commercial building projects. Production is rarely available on an interval that would allow for comparison to interval utility data. We have successfully used the method for industrial projects where the independent variable was weather rather than production, such as refrigerated warehouses.

IPMVP Option D involves use of calibrated simulations models. This technique is also rarely of use for industrial projects. Simulating all end uses in an industrial facility is rarely possible, and the data collected to calibrate the model should instead be used more directly in an Option A or B approach.

Energy use intensity (EUI) is a concept that is commonly applied to industrial projects. The energy use per some production metric is calculated and the pre- and post-installation EUIs are applied to typical production levels to estimate savings. The methods are applied at the facility level for the Superior Energy Performance program[1] (where it can be thought of as a form of Option C) and can be applied at the system level by evaluators (where is can be thought of as a form of Option B). Some pitfalls of the approach are described later.

Part of what makes industrial projects challenging for evaluators is the different information that is available from site to site. For a compressed air project, continuous measurements of compressed air plant output (CFM) and input (kW) are ideal. With these data we can apply a rigorous IPMVP Option B approach. However, these data are often unavailable. Similarly, for a process cooling system, our job is simple when we have access to continuous measurements of cooling plant output (tons) and input (kW). Unfortunately for the evaluator, continuous load information is unavailable in the majority of facilities. The only parameters for which long term data typically exists are weather and production.

Types of Measures

Proprietary and complex systems are at the heart of many industrial facilities. Fortunately for evaluators, most energy efficiency projects do not affect actual manufacturing processes, but rather they affect the secondary support systems. These include process cooling, compressed air, lighting, and space conditioning. There is much less variability in the design and operation of these support systems between facilities. Experienced M&V engineers are more easily able to understand measures and develop expertise with these types of systems. Unfortunately, the information available in each situation is often very different.

There is a common pattern to many industrial secondary support systems: 1) the efficiency of the system is dependent on the output of the system, and 2) the output of the system is related to the production output of the facility. This is true of compressed air systems and process cooling systems,

[1] The SEP program and its M&V protocols are described later in the paper. http://www.energy.gov/eere/amo/downloads/superior-energy-performance-measurement-and-verification-protocol-industry

where the amount of chilled water or air produced is typically dependent on production. When production levels increase or decrease, so does the load on the secondary systems. The variable with the largest effect on the efficiency with which the air or cooling is produced is typically the load on the system. Many systems get dramatically more efficient as loads increase. As stated earlier, it is useful to divide industrial measures into 2 groups: supply side efficiency measures (SSE) and demand side efficiency measures (DSE).

Table 1. Typical Measures for Supply and Demand Efficiency

Efficiency Type	*Typical Efficiency Units*	
	Compressed Air Measures	*Process Cooling Measures*
Supply Side (SSE)	kW/CFM	kW/Tons of Cooling
Demand Side (DSE)	CFM/Production	Tons of Cooling/Production

The proposed method was first developed for evaluating compressed air projects, some of which improved the efficiency with which air was supplied, some of which decreased the demand for air, and some that did both. We later discovered that the same technique worked equally well for process cooling measures. The method could also be applicable to conveying systems, dust collection systems, and other production dependent industrial systems.

The Method

To accurately estimate savings, the authors have found the following relationships must be quantified: the dependence of both pre-installation and post-installation SSE on load, and the dependence of both pre-installation and post-installation DSE on production and other independent variables. Impact evaluators commonly only have access to short-term pre and post-installation data. Fortunately, short-term data of a week or two is typically adequate for quantifying the required relationships. Since production is typically the driver of energy use, the duration of pre- and post-installation metering should ideally be long enough to capture at least one production cycle, commonly one week.

The SSE is defined as the energy consumption required to produce a unit of system output. Common measurements are kW/CFM and kW/ton. It can be characterized in several ways, but it is generally possible to create a plot of SSE versus load such as a plot of kW/ton against cooling tons.

The DSE is defined as the support system output (e.g. tons of cooling or CFM of compressed air) required to produce a unit of plant output. For a compressed air project, the method requires that the demand for compressed air CFM be characterized in terms of plant production or a proxy for production, such as production mode (full, part, off line, etc.). Annualizing the facility production load profile is necessary for annualizing the results of the short-term measurement and verification.

The data required are as follows for a typical compressed air project:

1. Customer/vendor will measure compressor CFM and input power to all pre-retrofit compressors prior to project implementation for at least one week.
2. Customer will record units of production during the same pre-retrofit period (weight or quantity of product, production hours, number of employees per shift, etc.). A sensitive and careful interview with the customer is often required to determine what data are available. This may involve contacting customer personnel at different levels of the organization.

3. Customer/vendor will measure compressor CFM and input power to the post-retrofit compressor for at least one week after measure installation and commissioning is complete.
4. Customer will record units of productivity for the same post-installation period.
5. Customer will provide annual information on units of activity (weight or quantity of product, production hours, number of employees per shift, etc.).

Following upon the data collection, the SSE and DSE must be characterized. SSE can be characterized as a plot of system kW versus load or as the system efficiency versus load. Either variable can be used similarly in calculations. The choice should be based on which characterization presents the most accurate narrative of the project. We have found that the plot of efficiency rather than kW can provide a more intuitive demonstration. For example, the plot below shows the same data, but the presentation of kW/CFM is more intuitive since it shows a clearly constant efficiency. The "story" of the kW plot is perhaps less clear to an inexperienced consumer of the report.

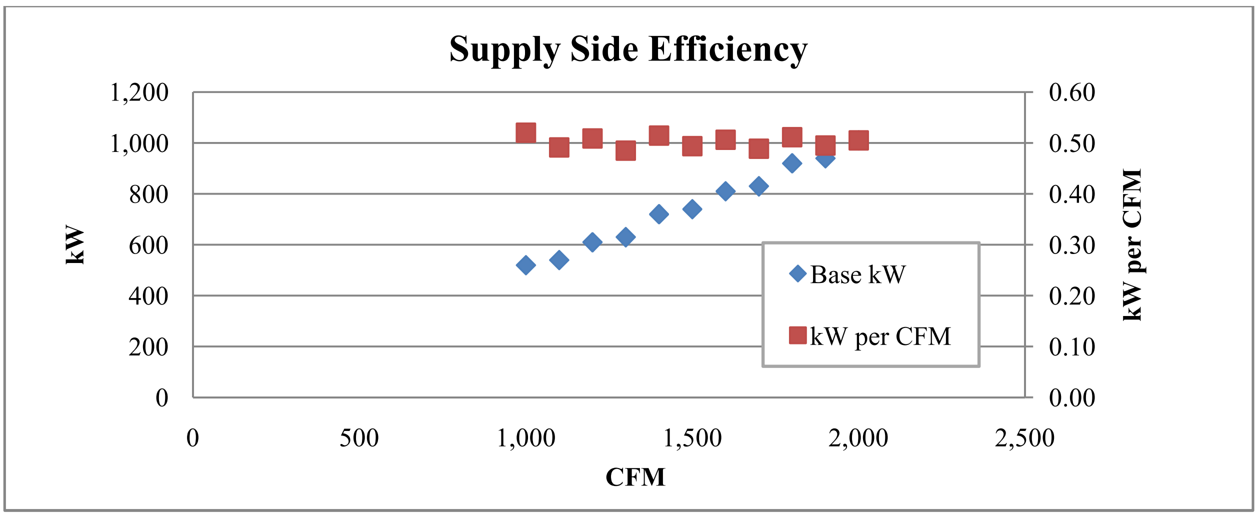

Figure 1. Alternative Presentations of SSE

Characterizing the DSE can be more challenging because of the wide variety of production data available. Production data are typically available on longer time intervals than consumption data. For example, having kW data on an hourly interval and production data per shift gives both needed metrics, albeit on different intervals. In such cases, you typically need to analyze the data on the larger of the two time intervals. We would need to compare the average kW used over each shift. It is important to note that even though the time interval is relatively long (several hours to a day) there is still great utility in calculating the DSE for each interval rather than for the entire metering period as a whole. Breaking the data collection period into as small intervals as possible allows a more accurate view of the relationship between DSE and production.

Regression analysis is typically used to characterize the SSE but may not be useful for the DSE. The production can be characterized as one of a few modes (weekdays and weekends, high and low production, regular shifts vs. maintenance shifts, etc.). The average value of the dependent variable (e.g. tons of cooling) for each scenario, or bin, will be determined and the result will be a matrix of production scenarios with an associated tons. If the production levels are found to vary more significantly, there may have to be more scenarios/bins. The bin strategy will be based on both an understanding of the production and system demand characteristics of the facility as well as analysis of the metering data. The system load should be plotted against time (histogram). The metered data are examined to determine if the load varies in a continuous manner or if the data can reasonably be divided into several scenarios/bins. If these bins or strata emerge, they can be compared to any available production metric during the observed periods.

Finally, savings are calculated using one of the algorithms presented below. An important concept in the calculation approach is that the unaffected parameter must be held constant[2]. This is a key advantage of the method over EUI methods. It is generally not realistic to think that the measure is causing either the SSE or the DSE to get worse. If the measure includes supply side measures only, the measure has no impact on the demand for compressed air, and only the SSE is expected to change. The DSE can often show significant, apparently random, variation that cannot be fully explained by the available production data. If the measure improves the SSE by 10% but the DSE fluctuates by a similar amount, the calculated savings will be extremely inaccurate if the calculation does not control for DSE. Conversely, if the measure involves repairing leaks, solenoid valves, or more efficient air using equipment, it will affect the demand for air (DSE) but not the SSE efficiency curve In that case, only the demand on the system (tons or CFM) is expected to change and the efficiency curve of the supply side is unchanged. Note that the actual SSE value may change, because the system will be operating at a new point on the SSE curve, but the SSE curve itself is unchanged.

Savings Algorithm

The general form of the algorithm applies where measures include both supply side and demand side measures. The efficiency of supply system is expected to improve and the overall load on the system (CFM or tons of cooling) is also expected to decrease. The savings are calculated according to the following equation:

Annual kWhsave = kWhbase - kWhpost

Where:

$\text{kWhbase} = \sum \text{i}(\text{SSE}_{\text{pre,i}} \text{ x } \text{DSE}_{\text{pre}}) \text{ x production,i x hours,i}$

$\text{kWhpost} = \sum \text{i}(\text{SSE}_{\text{post,i}} \text{ x } \text{DSE}_{\text{post}}) \text{ x production,i x hours,i}$

Equation 1. Savings Algorithms where Supply and Demand Side Improvements Installed

Where "i" is each production mode. The annual hours at each load increment will be estimated using available production data from the customer. This formulation applies to situations where production is the main independent variable affecting both the DSE and SSE.

Since there are demand side measures, the load for each production mode is expected to decrease; therefore, DSE_{pos} and DSE_{pre} are not expected to be equal at each typical production level. Should DSE_{post} be greater than DSE_{pre} for any mode, then DSE_{post} in the equation will be set equal to DSE_{pre}. This adjustment would have the effect of capturing the savings from the supply side aspects of the measure but treating the demand side measures as having no savings (but not negative savings). The same logic applies to the SSE, and SSE_{pos} will not be allowed to be worse than SSE_{pre}.

The equations assume that production modes will be used rather than a direct continuous measure of production output. If more production data is available to make the use of production modes or bins unnecessary, the pre-installation and post-installation energy use can be more accurately determined.

[2] Note that the DSE and SSE and generally not constants, but instead are characterized as described above, so by holding them constant we mean applying the same characterization to both pre and post periods in the analysis.

Alternative - End Use Energy Use Intensity

Both CFM and Tons are relatively difficult to meter on a short-term basis if the meters are not already installed by the customer. As a result, program participants and program implementers can be reluctant to install or require the installation of monitoring equipment. It is hard for the evaluator to justify the expense and the customer impact of performing the metering after an incentive has been dispersed. Furthermore, the evaluator is typically unable to perform this metering for the baseline because they were not aware of the project until after it has been installed (new "real-time" evaluation approaches are changing this traditional impediment). Though it can be difficult to obtain the output of the system (tons or CFM), it is relatively easy to perform short-term metering of the input power (kW or amps). As a result, there is often a preference for comparing the dependence of pre-project energy use and post-project energy use directly to production or other independent variables.

The alternative method looks at the system kW directly rather than separately analyzing the supply side efficiency and demand side efficiency terms. The easiest short-term data to collect is input kW or amps. Say the implementer or the customer's contractor monitored amps on (4) existing air compressors prior to the retrofit and then the evaluator monitors kW on (3) existing and (1) new compressor for 2 weeks post-installation. Implementers and evaluators are accustomed to creating a method for annualizing the short-term results.

Annual $kWh_{saved} = kWh_{base} - kWh_{post}$

Where:

$kWh_{base} = f_{base}$ (typical annual production)

$kWh_{post} = f_{post}$ (typical annual production)

f_{base} is a mathematical model (such as regression, bin, binary) of the baseline system (such as compressed air) energy use, as a function of production level or production mode, developed from pre-installation metering and production records.

f_{post} is a mathematical model (such as regression, bin, binary) of the post-installation system energy use, as a function of production level or production mode, developed from post-installation metering and production records.

Equation 2. Savings Algorithms where Total System Improvement is Measured Together

The Superior Energy Performance (SEP) Program is a DOE program that encourages and assists facilities in achieving ISO 50001 certification and improving their energy performance. The SEP M&V protocol focuses on measures of plant level energy performance but some of the concepts could be applied to system level measurements. The allowable model forms described in this program's materials are "ratio of energy consumption to a single production level", simple regression model, and complex regression model.

1. **Simple average.** The SEP program refers to this as "ratio of energy consumption to a single production level". The average kW of the system in the pre-installation metering period less the average kW of the system in the pre-installation metering period. In the most simplistic analysis, the functions are a simple average demand value. You can simply compare the pre-installation kW and the post-installation kW. This theoretically can be acceptable as long as you first test the regression options and can find no relevant variable. One step more sophisticated is to determine a single average kWh/production for the baseline period and apply that to the post project productions levels to determine baseline energy use and hence savings.
2. **Simple regression model.** Linear regression is used to determine the dependence of the pre-installation system kW on an independent variable of production. The same is done for the post-

installation short-term data. The pre- and post-installation regressions are then applied to the same set of annual conditions to determine normalized savings.

3. **Complex regression model.** Higher order forms are used to characterize the dependence of system kW on an independent variable of production. Since the method typically involves a degree of extrapolation, higher order forms should be used with extreme caution. They are prone to deliver unrealistic results at values of the independent variable even slightly outside the range of data used to develop the relationship.

The following example illustrates the potential pitfalls of the EUI method. Let us consider a case where we obtain energy usage data (kW) and production data, for a period prior to and following installation of a measure. The energy usage information might be of the total facility or could be the usage of a system (e.g. a process cooling plant). From the data in the table, it is tempting to estimate savings but we actually have insufficient information to determine the savings.

Table 2. Example of insufficient EUI data

Period	*Production*	*Energy Use (average kW)*	*EUI*
Pre-installation	**1,500**	**750**	**0.5**
Post-installation	2,000	900	0.45

If we take the simple change in average demand, the project has no savings and created an energy penalty of 150 kW. Knowing that production is not the same in both periods, this conclusion would be correct only if we assume there is no dependence on production. If we assume, or know, that the metered post-installation production levels are typical of typical annual levels, standard practice is to estimate the savings as the difference between baseline energy use at the post-production levels and the post-installation energy use. If the average baseline EUI is used, the baseline energy consumption would be: 0.5 x 2000 = 1000 kW, resulting in savings of 1000 - 900 = 100 kW. We have seen this approach applied, though it is often incorrect. The method only works when the baseline and post EUIs are constant across all production levels.

The following plot demonstrates the required conditions for the use of a simple average EUI to be correct. The system demand is linearly dependent with production and the intercept is zero. As production approaches zero, so does the system power. As a result, the EUI is constant.

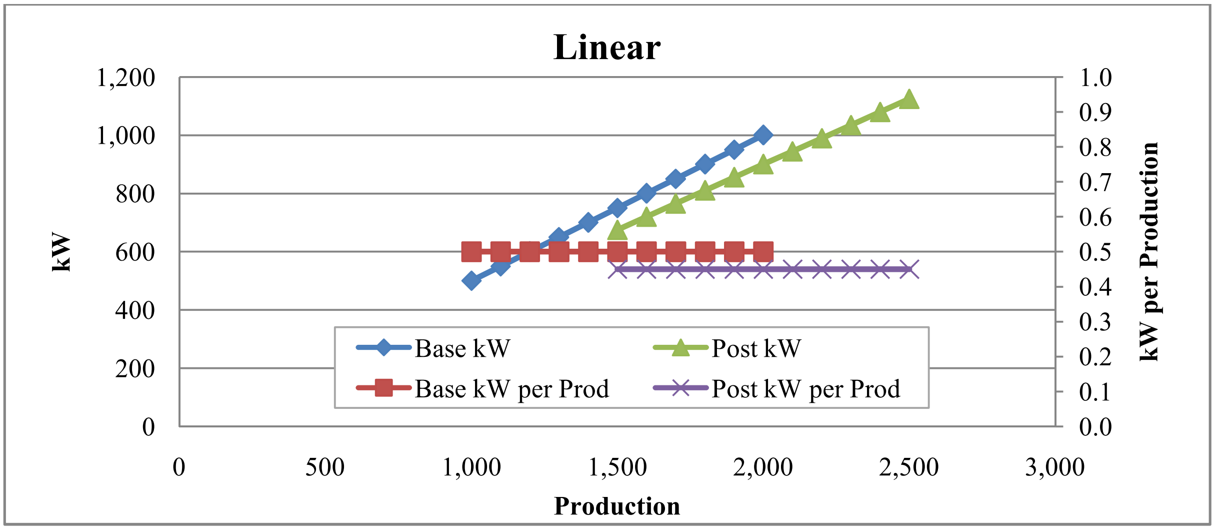

Figure 2. Constant EUI

What happens when the EUI is not constant? If you apply the same logic as above, the savings estimate can be dramatically incorrect. Energy use is often dependent on production but does not tend toward zero as production decreases as is demonstrated below. There is a dependence of demand on production but demand does not trend to zero. As a result, the EUI increases as load decreased.

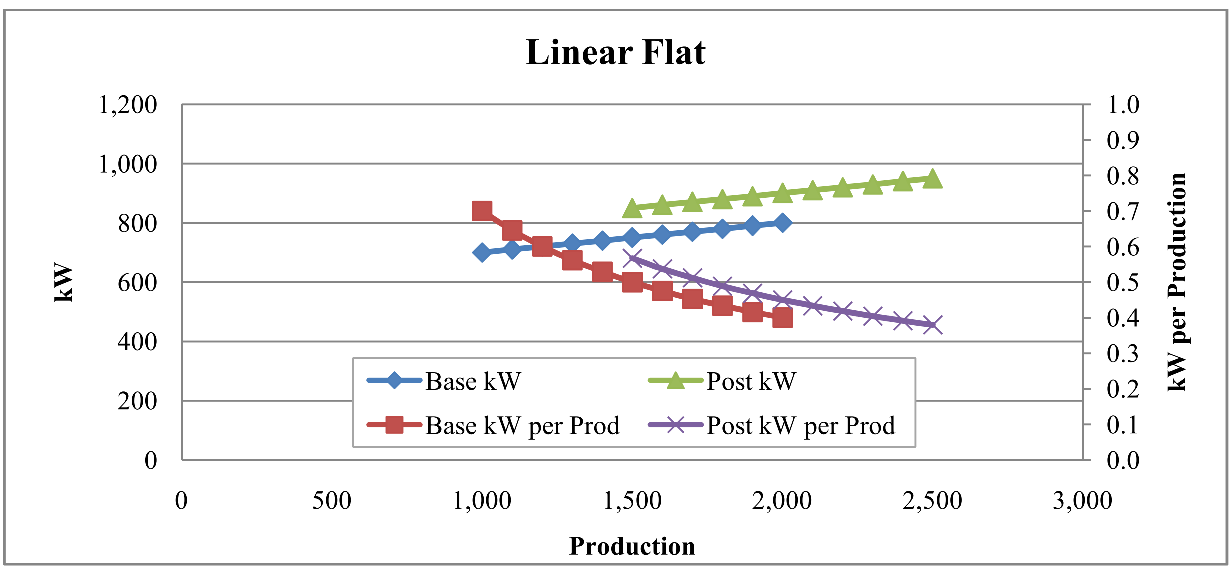

Figure 3. Variable EUI

For the facility shown in the second plot, the baseline energy use at the post-production level is 800 kW, resulting in negative savings of 100 kW. The average pre- and post-installation kW and production levels are exactly the same in the two examples. Only the demand vs production slope is different. The table tells us nothing about the shape of the EUI curve, but the examples show that different shapes can yield savings between positive 100 kW and negative 100 kW.. Clearly, the information in the table was insufficient for determining the savings.

Why Is the Proposed Method Generally Better?

A key issue with the EUI method is that it can be easily misused. Having an incomplete understanding of the dependence of EUI on production can lead to inaccurate or incorrect results. Pre-installation and post-installation production levels are rarely exactly the same, and the system energy use intensity can vary wildly across different operating points for the facility.

However, even with an understanding of the EUI dependence, this method does not give us any insight into the SSE and DSE individually, making evaluation of savings due to different measure types less certain.

It is impossible to say if the observed change was due to the efficiency of the affected system or a drop in the system load. A small variation in CFM or cooling load, unrelated to the project, can mask or overstate the savings from the efficiency upgrade. Particularly for projects that impact only the supply or demand side of the system (but not both), using separate demand-side and supply-side efficiency parameters enables the evaluator to control for the unaffected parameter.

We saw this recently on a large pharmaceutical chiller plant optimization project (SSE measure). The customer plots of chiller plant kW against ambient temperature showed a compelling decrease in energy use. However, upon examining the DSE (tons as a function of weather) and SSE (kW/ton), it was determined that the major cause of the reduction in energy use was an improvement in the DSE curve. The project was not expected to have had an impact on the DSE curve, so the focus of the evaluation became determining if such an improvement could reasonably be attributed to the measure.

This is a common issue with compressed air projects as well. Many measures affect only the SSE or the DSE, while some measures can affect both. Leak detection programs, replacement of "open blows," and installation of solenoid control valves are examples of projects that improve the DSE. New air compressors, VFD driven compressors, improved storage, and control strategies all improve the SSE. Many custom projects involve the installation of multiple measures and will thus improve both SSE and DSE.

Consistency with Existing Protocols

The BPA Verification by Equipment or End Use Metering Protocol suggests the following equations for measures with variable load and timed schedule, for which the efficiency is changed by the measure but the hours of use are not changed.

$$kWh = \Sigma\ (kW_{base,i} \times hrs_{base,i} - kW_{post,i} \times hrs_{post,i}) \text{ or}$$

$$kWh = \Sigma\ [\ (Eff_{base,i} \times hrs_{post,i} - Eff_{post,i} \times hrs_{post,i}) \times Load_{post,i}\ ]$$

Equation 3: BPA Usage Savings Algorithms

The first equation compares direct measurements of pre and post-installation kW. This can be a continuous measurement or bins could be created. In an industrial setting, this would generally be presented as a table of pre-installation and post-installation kW at various operating conditions or modes. This method does not allow for the supply side and demand side impacts to be separately assessed.

The second equation assumes that the load is unchanged by the project but the system efficiency is improved. In our example, the demand side efficiency is unchanged but the supply side efficiency is improved. The protocol does not provide guidance in the cases where both DSE and SSE are affected.

The Uniform Methods Project (UMP) Protocol for Compressed Air only applies to two compressed air measures: 1) installation of a VFD compressor rather than a single speed compressor, and 2) leak repairs. Larger compressed air projects may not fall into these categories. Typical larger projects involve various changes to multiple-compressor plants. Added storage, improved sequencing, pressure reduction, and leak repair are frequently implemented together.

Our proposed approach is consistent with some of the data collection suggested by the UMP for compressed air but is more widely applicable. For the VFD compressor measure, the UMP protocol calls for development of a "CFM demand profile. A demand profile must be developed to provide accurate estimates of annual energy consumption. A demand profile typically consists of a CFM-bin hour table summarizing hours of usage under all common loading conditions throughout a given year . . . The annual CFM profile is used to determine base case and proposed case energy use. For both, compressor electricity demand for each CFM bin should be determined from actual metering data, spot power measurements, or CFM-to-kW lookup tables."

What If it is Infeasible to Measure Load?

If it is infeasible to directly measure the system load (CFM or tons), is the evaluator compelled to use an EUI method? No, it is often better to calculate the CFM from performance curves (CAGI data sheets). The calculation of CFM introduces uncertainty, but it is our view that this is more than offset by the ability to correct for the nonlinear form of the demand side and supply side efficiencies. The UMP states;

> "One common method is to measure compressor power. The percent power can be correlated to percent flow using the appropriate compressor curve for the given type of control type. In this way, a load profile can be developed that can be used to compare the baseline and post systems at equivalent flow.
>
> Measured or trended airflow (SCFM) data can be quite advantageous when evaluating compressed-air ECMs; however, this information can be difficult to obtain and is not generally collected unless the existing compressed-air system controls already have the capability. In the absence of measured or trended CFM data, the evaluator must develop parameters..."

Another Benefit: Engineering Insight into the Measures

A true real-time impact evaluation approach was applied to a recent custom compressed air project. The evaluation was conducted in parallel with the program implementation, so the evaluator was able to guide the data collection and M&V approach to the project. The measure involved new air knives, solenoid valves to shut off air to open blowing applications, and adding a new VFD-driven compressor to a multi-compressor plant. It was expected to improve both the SSE and the DSE of the compressed air plant.

The customer monitors CFM and production but not compressed air plant kW. In order to obtain the SSE, it was thus necessary to perform some power metering. The customer performed a pre-installation metering study of kW and CFM for a period of 20 days to establish the baseline. The metering was repeated for an additional 18 day period following installation of the measure. The kW was actually higher in the post period. It initially looked as if no incentive would be available. However, because both kW and CFM were metered, we were able to determine if the cause of the higher kW was either the SSE or the DSE. In this case, the SSE had improved but the situation with the DSE was more complex.

The post CFM was higher but so was production. The average DSE had improved slightly, but when the actual curve fit of the baseline DSE was applied to the higher post-production levels, it was found that the baseline DSE at these conditions was actually better than the post case. The higher CFM values were brought to the customer's attention and it was determined that the wrong air knives had been purchased. The new knives actually used more air than the old knives and led to an increase in energy use[3]. In addition, the customer discovered that a number of the solenoid valves had failed. The purchased valves were not suitable to the harsh environment in which they were placed. According to the customer, these errors would almost certainly never have been found absent the program-required metering, and the savings would not have materialized. The customer did not have budget to buy the correct knives, but they were able to put the original knives back in service. At this point, the demand side of the project was essentially uninstalled, so the "baseline" metering was repeated by locking out the new compressor for a week. The metering was left in place and the new compressor energized in order to collect nearly 4 more weeks of post-installation kW data.

The results were still much lower than expected. The evidence for reduced CFM demand was very weak. When faced with a regression with a bad[4] R-squared, you get more data if you can. We were able to do this because the site had continuous CFM monitoring and were willing to extend the post-installation data collection period. Another two months of CFM data was collected. Though the R2 of the post-installation DSE is relatively low, the shape is what was expected - with savings occurring mostly at low production rates when the valves would be expected to be closed more often.

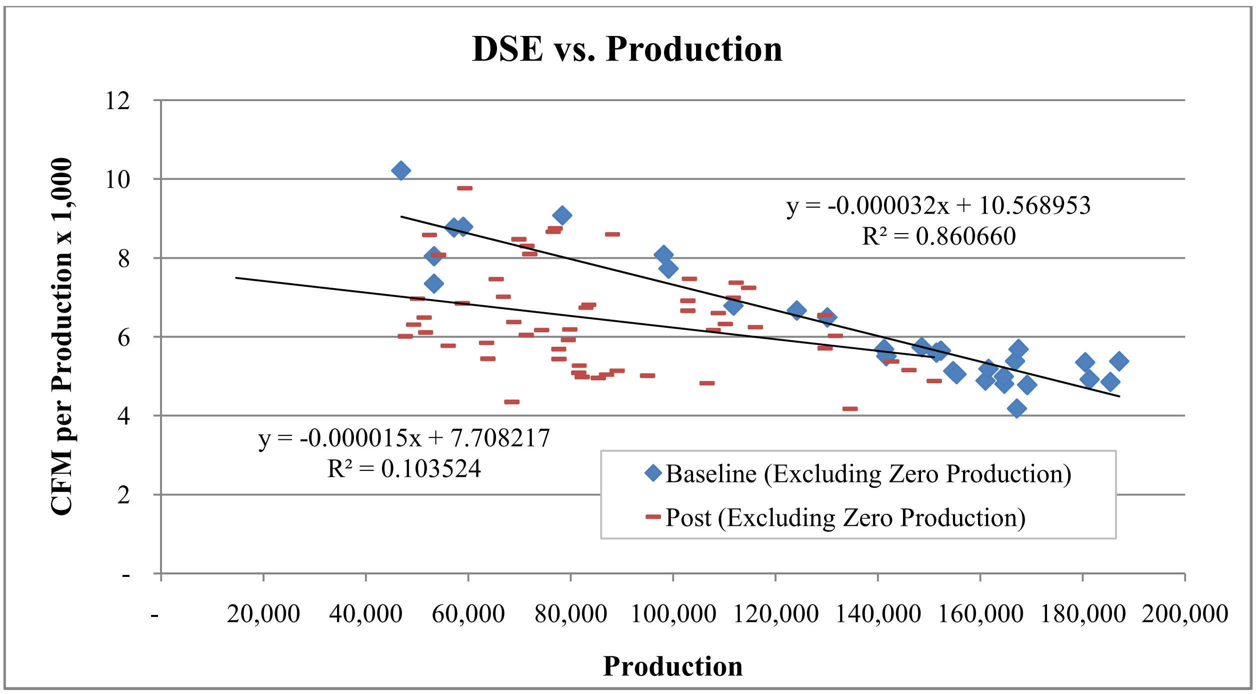

Figure 4. Example Demand-Side Efficiency vs. Production

The data for SSE is extremely "good" (excellent R2), but was also not what was expected. We noticed a discontinuity in the data and separately analyzed low flow and high flow conditions. The penalty at low flows was brought to the attention of the customer, for which they were grateful, and it was discovered that the site's oldest compressors were being used in these conditions as a result of a maintenance concern with running the new larger units at low flows. The customer was unwilling to change the sequencing, and accepted that the savings were significantly less than they had originally projected.

[3] Note that this project was an exception to our earlier contention that measures apparent DSE or SSE penalties should generally be reduced to zero in the analysis.
[4] We hesitate to specify a minimum acceptable R-squared value. Several sources point to 0.75 as a threshold, but we frequently find that this level is impossible to achieve when production is the independent variable. Even with r-squared values, the question to ask is if the dependence on the independent parameter is clear enough that it should be used. We typically feel that yes, it is better to use the variable than to use a simple uncorrected average.

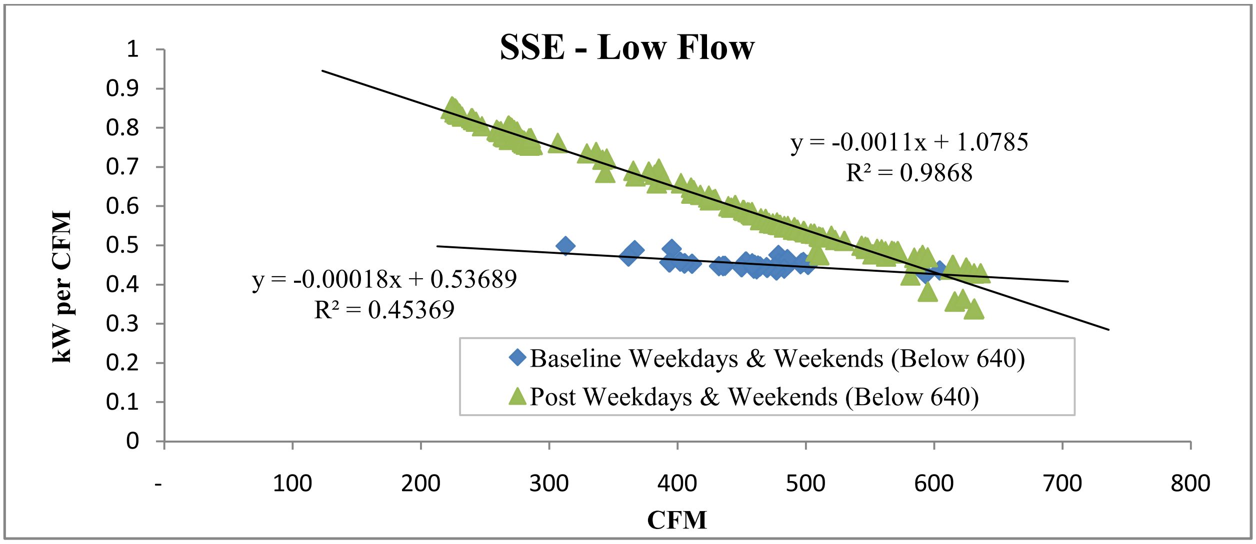

Figure 5. Example Supply Side Efficiency at Low Flow

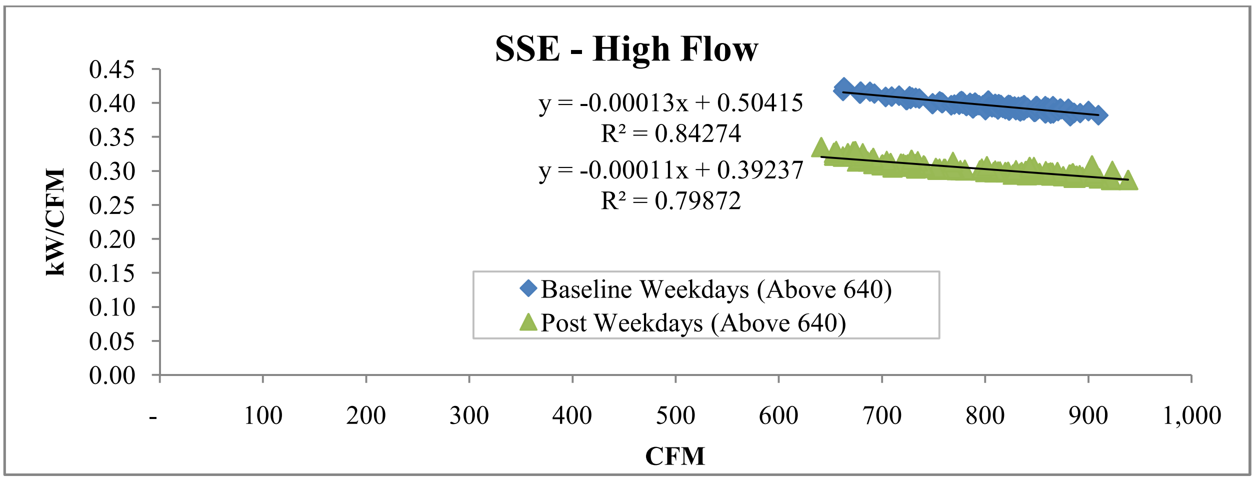

Figure 6. Example Supply Side Efficiency at High Flow

Though metering was performed an extra time and the incentive was somewhat delayed, the utility wound up with a very happy customer. If only short-term metering of power had been performed (without separate analysis of SSE and DSE), there would simply have been a poorly performing project, a small incentive, and an unsatisfied customer.

Conclusion

Many industrial energy efficiency projects involve the systems that support the core industrial process. Whether the process is making pet food, soda, or bottles, there is often a need for the same types of support systems such as compressed air, process cooling, and conveyance. It is often useful to think in terms of a "Demand Side Efficiency" and a "Supply Side Efficiency" when analyzing industrial energy efficiency projects. Characterizing systems such as compressed and process cooling in these terms has several advantages over approaches which consider the energy use of the system as a function of production level. A key advantage is that the unaffected parameter can be held constant in the savings calculations, leading to more rigorous results.

While it is commonplace to consider the part load efficiency of a chiller or fan, it is equally important to realize that the efficiency of industrial systems are also typically not constant. This is true of the energy use intensity of an entire facility and also of the DSE and SSE of individual systems. It is typically inaccurate to "correct" for production level by applying the DSE or SSE determined at one production level to a new level. Instead, the dependence of DSE and SSE on production should be determined so that the correct adjustment can be made.