The Advanced M&V Proving Ground: It's Working and It's Yielding Great Results

Abigail Hotaling and Robert Stephenson, VEIC, Winooski, VT

ABSTRACT

The ever-finer granularity of utility meter data makes it increasingly attractive for energy efficiency programs to adopt advanced measurement and verification (M&V 2.0) methods. Even though M&V 2.0's value to an efficiency program's evaluation process is clear, there are barriers to widespread adoption: (1) evaluators' incomplete understanding of which project types are appropriate for these advanced methods; (2) the inability of some methods to meet M&V requirements; (3) insufficient, relevant case studies that can inform program staff, regulators, and evaluators; and (4) insufficient data management to support M&V 2.0 at scale. This paper addresses these barriers, presenting an investigation of statewide energy efficiency programs' projects containing the type of data necessary to support the appropriate use of M&V 2.0 for assessing more real-world benefits of energy efficiency measures. The study team investigated the feasibility, applicability, and potential for cost and staff-time savings from replacing custom site-metering evaluations with M&V 2.0 in 142 commercial, industrial, and multifamily projects. The team also compared evaluation methods from M&V 2.0 model results with *ex ante* savings claims for those projects. The team then applied three open-source M&V 2.0 models to assess uncertainty, and precision requirements for savings claims within the portfolio, segmented by building end use and savings level. The team finally tested procedures for pre-screening projects for their applicability to M&V 2.0 methods. M&V 2.0 pre-screening reduces staff hours in analysis scoping and yields non-subjective quantitative results. The research provides the information regulators and evaluators need to accept M&V 2.0 methods with confidence.

Introduction

The ongoing enhancements in computer technology and improved computational methods in the last several decades have significantly advanced the scope and quality of energy efficiency evaluation. A next-generation method, advanced measurement and verification (M&V 2.0), uses higher-frequency energy meter data to feed statistical models,[1] and processes large amounts of data through automated analytics. These capabilities correspondingly allow energy efficiency savings estimation to be conducted more quickly, accurately, and potentially at lower cost, compared to other M&V options.[2]

The primary use cases for applying M&V 2.0 are building energy performance tracking, pay-for-performance programs, strategic energy management programs, flexible load management projects and programs, and impact evaluation. Although the benefits from applying M&V 2.0 to more evaluation processes are obvious, hurdles to widespread adoption still exist:

1. Efficiency program staff, regulators, and their evaluators have an incomplete understanding of which project types are appropriate for these advanced methods.
2. Methods do not meet M&V requirements in some cases.
3. Program staff, regulators, and evaluators do not have access to sufficient M&V 2.0 case study literature to inform regular M&V assessments.

[1] The ideal frequency of data collection is sub-hourly. In any event, although monthly billing data might be sufficient for regular M&V methods, they are insufficient for M&V 2.0.
[2] That is, International Performance Measurement and Verification Protocol (IPMVP) Options A, B, and D.

4. Staff, regulators, and evaluators need data management systems that are adequate to support M&V 2.0, prior to adoption at scale.

Background

This investigation offers a case study that examines the use of M&V 2.0 methods by a regulated statewide energy efficiency utility, Efficiency Vermont, which had claimed energy savings for historical commercial, industrial, and multifamily retrofit projects. Efficiency Vermont's current M&V work informs savings claims that are reported to ISO New England's Forward Capacity Market (FCM). The FCM market participation prompted this study, which sought to test the feasibility of an IPMVP Option C whole-facility M&V analysis strategy to replace or augment other M&V methods. Even though the study examined FCM impact evaluation as the primary use case, the team has concluded that the findings can apply to other efficiency program activity such as building energy performance tracking, strategic energy management, and flexible load management. The team considered the potential for Option C (specifically, M&V 2.0) to generate cost and staff-time savings, refine current savings estimation methods, and allow for long-term monitoring of project performance.

The study design also sought to achieve results that would allow program administrators and evaluators to assess the feasibility of applying M&V 2.0 methods to advanced metering infrastructure (AMI) energy data, particularly short-interval data. In this way, they could have the information necessary to decide whether to augment or replace on-site metering, where such methods meet required levels of confidence and precision.

Each year, Efficiency Vermont submits peak demand reduction claims to the FCM. A third-party evaluator conducts the M&V to verify demand reduction claims and determine if the confidence and precision of the portfolio meet the grid operator's M&V requirements. Typically, these M&V requirements are satisfied via stipulated load shapes or on-site metering of specific equipment to verify savings claims. Evaluators rely on M&V Options A or B, but these methods involve costs to deploy sub-meters at building sites and further staff costs to analyze the results. Significant cost reductions are possible via Option C, which requires no meter deployment and uses utility data, instead. Because Option C can be applied on a large scale, automated data collection and analysis are more cost-effective. This study has investigated the applicability of using a whole-building facility analysis to verify project savings and satisfy M&V requirements. The working hypothesis for this project was that a reduction in meter deployment, analysis planning, and data collection would provide significant administrative savings for Efficiency Vermont. This paper:

1. Illustrates how well-designed data infrastructure can support broad application of M&V 2.0
2. Identifies the parts of custom commercial and industrial (C&I) portfolios where M&V 2.0 is effective
3. Helps evaluators understand the applicability of M&V 2.0
4. Evaluates the use of M&V 2.0 for pre-screening program administrators' projects

Data Infrastructure

The success of this study and the success of scaling AM&V are contingent on well-designed data infrastructure. The infrastructure makes it possible for high levels of automation to be added to the analysis without adding significant staff time for data collection and preparation. VEIC, the organization that operates Efficiency Vermont, has undertaken several infrastructure projects that store external and internal data sources so that staff can access them through user interfaces—or, in this case, via data access software and analytics tools.

The relevant data management and analytics components of VEIC's infrastructure are:

- **Project tracking warehouse.** VEIC's project tracking database is the source for all customer data necessary for managing energy efficiency projects. The warehoused components for this analysis are information about measures, installation dates, and claimed savings. Each project is also associated with a site with its own attributes, such as mapping to utility identifier, facility address, business end use, and records of all past projects.
- **AMI warehouse.** VEIC has built its AMI infrastructure on an extract, transform, load (ETL) framework that takes utility AMI data and ingests them into a VEIC-maintained warehouse. The framework makes it possible for data from utilities to be loaded and transformed into an accessible, standard format, ready for analysis. VEIC ingests the data on a schedule so that the warehouse is regularly and automatically updated. This framework also makes it possible to streamline the onboarding of new data from utilities added to the portfolio.
- **Curated weather.** VEIC also maintains curated weather data for all weather stations in the service locations in which the organization operates an energy efficiency utility. A semi-automated process updates the weather files in consistent directories and file formats. This data source also maps local ZIP Codes to corresponding weather stations. The system makes weather data accessible VEIC-wide and enables automated data retrieval for any building.
- **Analytics framework.** VEIC built a general software framework for conducting and collaborating on analysis projects. This framework involves modular architecture for building analysis pipelines, parallelized computing, and a metric database schema applicable to a wide variety of energy efficiency applications. VEIC has also created a code library that wraps open-source code such as EEmeter to perform energy modeling, savings estimates, and model diagnostics (EEmeter 2019). Using this framework, VEIC built an M&V 2.0 analysis pipeline, which queries data from the organization's other infrastructure and uses this code library to perform M&V 2.0 on a large scale. A database stores the metrics computed in this pipeline, and can be accessed for reporting and evaluation.

Figure 1 shows how the interconnected pieces of VEIC infrastructure enable an automated Option C or M&V 2.0 analysis pipeline, without having to do significant data preparation up front.

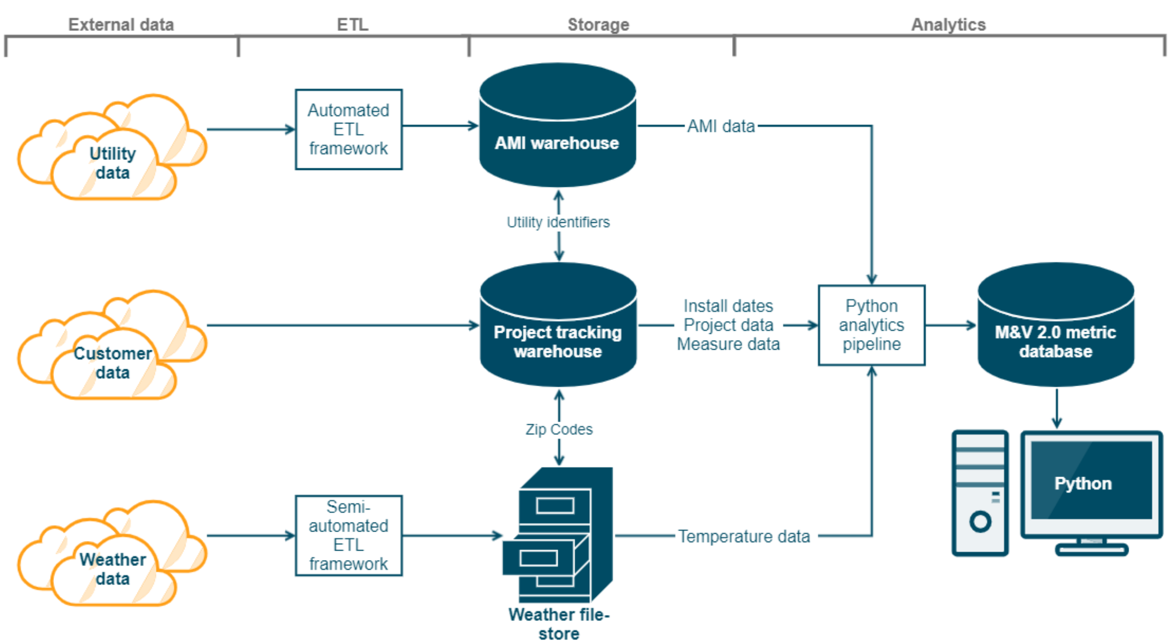

Figure 1. Overview of data infrastructure used for M&V 2.0 analysis.

From the project tracking warehouse, the pipeline can create baseline and reporting periods from measure installation dates, retrieve Efficiency Vermont savings claims, and obtain utility identifiers and building ZIP Codes. Using the utility identifiers, the analysis pipeline can query the AMI data directly from the warehouse. The final piece of the puzzle is weather data. The site's ZIP Code is the only identifier needed to extract specific temperature data from the network file structure.

Other case studies have shown that data access, management, and quality can be a substantial barrier to the broader deployment of these methods (Franconi, Gee, Goldberg, Granderson, Guiterman, Li, and Smith 2017; Crowe, Granderson, and Fernandes 2019). If more resources are needed for data preparation and management this can increase the per project analysis costs for M&V 2.0. The investment in this infrastructure makes it possible for M&V 2.0 to be integrated into many more programs and use cases. Data sources can be updated in near real time, also setting up opportunities for more real-time monitoring of energy efficiency programs.

Case Study Data

The data for this case study come from energy performance of the custom C&I retrofit projects that were claimed in the FCM during program year 2018. Study staff narrowed the dataset to look at sites for which AMI data are available in VEIC's core infrastructure. Of the custom C&I retrofit projects claimed in 2018, 70 percent had AMI data available. Finally, the staff filtered the projects for consistent data availability. This analysis required at least 12 months of baseline and reporting period data. Table 1 summarizes the sample sizes set after determining each requirement for the study population. Nearly 40 percent of projects claimed in 2018 met data requirements needed for a whole-facility analysis. The team expects this factor to increase in future years, since the AMI data stored by Efficiency Vermont starts in 2017, limiting the amount of available baseline data.

Table 1. Summary of project data requirements

		Amount remaining after requirement is met			
Requirement	Description	Number of sites	Percent of sites	Summer demand savings (kW)	Winter demand savings (kW)
FCM custom C&I project claims	FCM custom C&I project claims that exhibit either winter or summer demand savings	373	100%	2,764	2,709
AMI data available	Projects with AMI data in VEIC core infrastructure	260	70%	2,459	2,378
Data requirements	At least 12 months of baseline and reporting period data are available	142	38%	1,273	1,196

Analysis

The study team based the M&V 2.0 methods in this project on Option C normalized energy savings (EVO 2016). This approach models the actual energy use of a project facility, post-installation, against a baseline modeled on pre-installation consumption data. The team established baseline and reporting periods for each project from the installation dates recorded in the project tracking warehouse. The study required 12 months of baseline and reporting data. The team excluded a total of 4 months of data around the installation to ensure that the installation process did not bias energy savings.

The baseline models included independent variables such as outdoor air temperature, time-of-day, and time-of-week indicators to model the relationship with building energy load. The baselining methods were:

- **Change-point.** The study used the daily heating and cooling degree-day model implemented in Python by Open EEMeter (Linux Foundation Energy EM2 and Joint Development Foundation 2018). The EEMeter model optimizes for a cooling and a heating outdoor air temperature balance point, and estimates the effect of heating and cooling on energy use for each degree above or below the balance points.
- **Time-of-week and temperature (TOWT).** TOWT models use indicators for each hour of the week and the estimated occupancy status, and temperature bins to model the continuous temperature relationship. This case study has used the three-month weighted TOWT implemented in Python by Open EEmeter (Linux Foundation Energy EM2 and Joint Development Foundation 2018).
- **Gradient boosting machine (GBM).** This GBM model, implemented in Python, follows the GBM methods used in the Berkeley Lab's RMV 2.0 tool (LBNL / ETA 2020; Touzani, Granderson, and Fernandes 2018). This model also uses time of week, temperature, and federal holidays as an indicator. To avoid over-fitting, cross-validation is used during model training and parameter tuning processes.

To understand model goodness of fit and the uncertainty of savings estimates for a given project, the study team used the following metrics and thresholds (Granderson, Crowe, Fernandes, and Touzani 2019):

- **Coefficient of variation of the root mean squared error (CV(RMSE)).** This metric describes the variation unexplained by the model and indicates the predictive accuracy. The lower the CV(RMSE), the lower the savings uncertainty, and the closer the predictions are to the actual values. Threshold: CV(RMSE) < 0.25.
- **Normalized mean bias error (NMBE).** The NMBE is the total difference between the predicted and observed energy. Threshold: $\pm$0.5 percent.
- **Factional savings uncertainty (FSU).** This is an ASHRAE Guideline 14-defined FSU metric with adjustment for autocorrelation to estimate uncertainty in energy savings. Threshold: At 90 percent confidence, uncertainty is within 50 percent.

Even if models pass the model validity checks, other factors can influence the validity of the results of an Option C analysis—aside from baseline uncertainty. If, during the installation period, a facility makes other changes not related to the energy conservation measure (such as adding a new production line), this could greatly bias the energy savings, even if the models are a good fit to the data. To detect these types of errors, the team might have to inspect the data manually. Automating non-routine event detection and adjustments into analysis pipelines is an important area of active research; however, it was not an objective for this paper. Results with modeled energy savings that are outside the savings possible for that project should be removed. The team's professional judgment is that projects in which the magnitude of modeled savings was more than 3 times the magnitude of the claimed savings are outside what is possible for that project. Thus, in such instances, the team has removed them from the comparison between claimed and modeled savings.[3]

[3] This approach is based on non-routine event detection method No. 7 from the IPMVP *Application Guide on Non-Routine Event and Adjustments* (EVO 2020). The guide recommends checking that reporting period savings are close to project savings claims, and suggests this approach be supplemented by other non-routine event checks, because of possible variability in the claimed savings.

Results

Uncertainty and Bias by Model Type

The following primary findings show that M&V 2.0 methods can model building use and compute savings with acceptable levels of statistical certainty for many of Efficiency Vermont's custom projects, making this method a candidate for use on a wide scale. Table 2 shows the percent and number of projects that were within the thresholds for each metric. The hourly methods of TOWT and GBM showed lower uncertainty and bias compared to the change-point model using daily frequency energy data. The GBM model improved slightly from the TOWT model in all metrics.

Table 2. Percent and number of projects with goodness-of-fit metrics within industry standard thresholds. Number of projects within threshold are recorded in parentheses.

Model type	FSU *	Baseline NMBE **	Reporting NMBE **	Baseline CV(RMSE) ***	Reporting CVRMSE ***	All metrics
Change-point	41% (58)	100% (142)	100% (142)	48% (68)	56% (79)	23% (33)
TOWT	77% (110)	87% (124)	80% (114)	46% (66)	51% (73)	37% (52)
GBM	80% (113)	100% (142)	100% (142)	49% (70)	56% (80)	39% (56)

* FSU < +/- 50% at 90% confidence. ** NMBE < +/- 0.5%. *** CVRMSE < 0.25.

Uncertainty by Savings Magnitude

Another research objective was to understand the magnitude of identifiable energy savings, given the uncertainty of these methods. The inability to detect very small amounts of savings is often seen as a limitation of Option C. Figure 2 shows the proportion of projects within the FSU thresholds for energy savings under 10 percent.

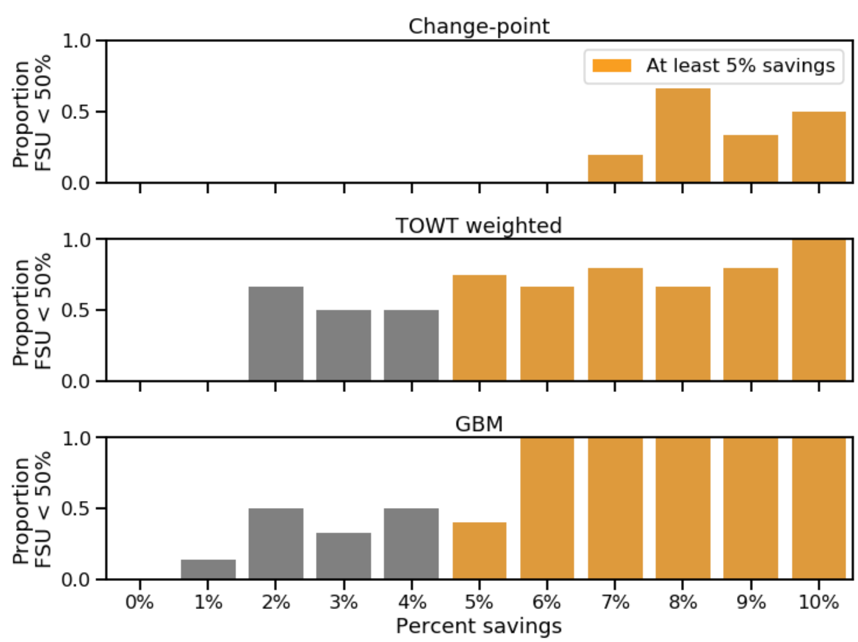

Figure 2. Proportion of projects with FSU thresholds grouped by percent savings, relative to baseline energy use.

For the hourly methods, we estimated savings for most projects with savings at or above 5 percent of baseline period energy use. This result aligned with IPVMP's M&V 2.0 guidance, and with results from other studies (Webster 2020; Kelly and Sinnamon 2020). The change-point model does not meet requirements as frequently for savings lower than 10 percent of baseline period energy use. This result emphasizes the benefit of having high-frequency energy data, and how advanced M&V methods are not applicable for all projects, especially those with lower magnitude of savings relative to baseline.

Goodness of Fit by Building End Use

The study team used Standard Industrial Classification (SIC) codes to classify buildings, thus offering insights on the primary type of business activity at each site. Figure 3 shows the proportion of projects that meet the uncertainty requirements for normalized savings for each end use.[4] Retail trade businesses showed the highest proportion of projects that met the criteria. Manufacturing businesses showed the least promising results, primarily because facilities can be difficult to model without having production data. That is, most of the load is process driven and the facility might not operate on a consistent schedule. The algorithmic complexity and flexibility of the GBM model showed stronger performance compared to the TOWT and change-point models in the difficult-to-model businesses: manufacturing, public administration, and finance and real estate facilities.

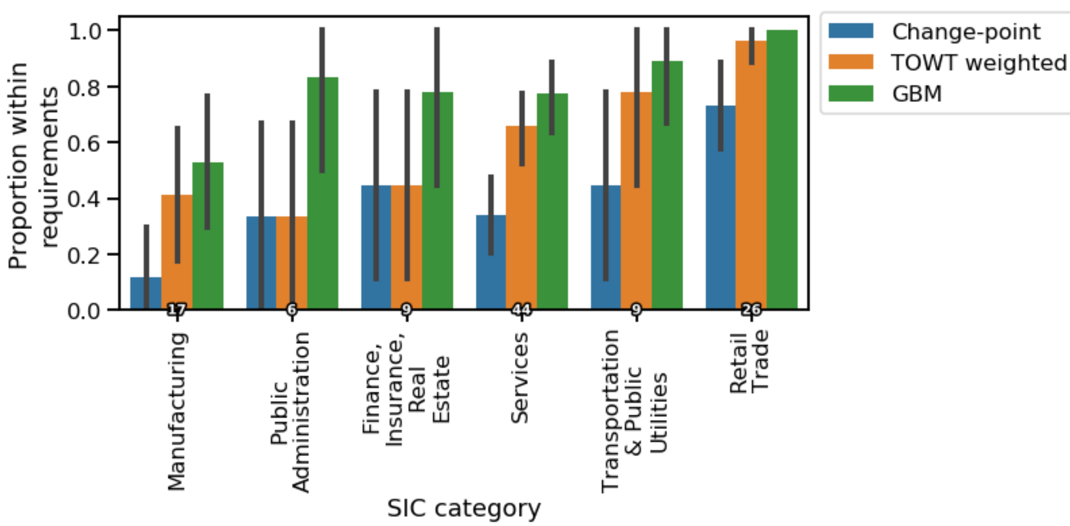

Figure 3. Proportion of projects within acceptable savings uncertainty, and bias criteria grouped by model type and business end use. Sample sizes are shown for each category of SIC codes.

To further understand how model uncertainty varies across different building end uses, Figure 4 shows the proportion of projects within goodness-of-fit criteria, grouped by model type and statistically significant weather parameters from the change-point model (cooling and heating degree-days). Of this sample, 10 projects had no statistically significant relationships with weather. The remaining projects had at least one statistically significant regression parameter for heating or cooling degree-day. The proportion of projects meeting requirements was lowest for sites with no significant weather relationship or only a significant heating degree-day relationship. The projects with significant cooling and heating degree-day relationships showed higher proportions meeting the requirements. This result additionally supports that process-driven loads, instead of weather- and occupancy-driven loads, have higher model uncertainty.

[4] Each category had to have a sample size of 5+ projects, to increase statistical validity of inferences from the results.

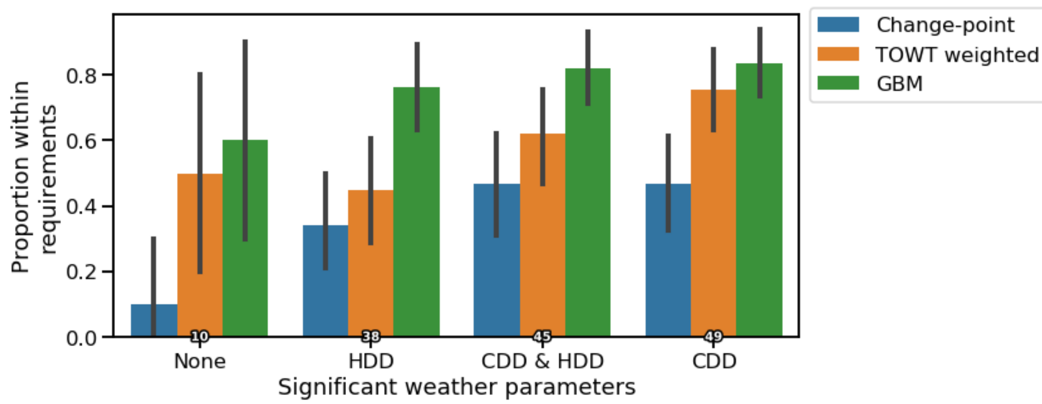

Figure 4. Proportion of projects within savings uncertainty and bias criteria, grouped by model type and statistically significant heating and cooling degree-day regression parameters, with sample sizes for each set of parameters.

Claimed Savings Comparison

To identify where M&V 2.0 is appropriate, the team filtered a subset of projects by data requirements, model validity checks, and feasibility of the results. Of these projects, the team compared model-derived results to Efficiency Vermont *ex ante* savings claims, calculated from methods such as direct metering, regression analysis, engineering calculations, and stipulated load shapes. The team grouped results by magnitude of savings. Stratum 1 shows projects with low savings, and Stratum 4 shows projects with high savings. The team computed realization rates for projects meeting the filtering requirements. The realization rate is the ratio between the whole-building, model-derived savings and the Efficiency Vermont *ex ante* savings claims. A rate of 1 means the modeled or evaluated savings equal the claimed savings; a rate less than 1 means the modeled or evaluated savings are less than claimed savings. Table 3 and Table 4 show the rates for the GBM and TOWT model strata. The GBM rate is 0.75 for winter demand savings and 0.65 for summer. The TOWT's is 0.71 for winter and 0.62 for summer. These rates show that, on average, modeled savings are less than claimed savings for this sample.

Table 3. Realization rates of claimed savings to modeled savings computed from the GBM model

		Winter demand savings (kW)				Summer demand savings (kW)			
Stra-tum	N	Realiza-tion rate	Margin of error	Average claimed	Average modeled	Realiza-tion rate	Margin of error	Average claimed	Average modeled
1	17	0.57	0.44	2.58	1.47	0.65	0.47	3.52	2.28
2	20	0.74	0.35	9.35	6.95	0.44	1.31	12.82	5.63
3	3	0.89	0.20	50.55	44.97	0.91	0.15	55.60	50.48
4	2	0.69	0.58	108.13	74.26	0.79	0.54	112.55	89.08
All	42	0.75	0.21	14.25	10.66	0.65	0.37	16.86	11.45

Table 4. Realization rates comparing claimed savings modeled savings computed from TOWT model

Stra-tum	N	Winter demand savings (kW)				Summer demand savings (kW)			
		Realiza-tion rate	Margin of error	Average claimed	Average modeled	Realiza-tion rate	Margin of error	Average claimed	Average modeled
1	15	0.55	0.99	2.59	1.43	0.57	0.71	3.50	1.98
2	19	0.71	0.56	9.40	6.69	0.36	2.09	13.24	4.82
3	3	0.88	0.26	50.55	44.69	0.91	0.19	55.60	50.36
4	2	0.62	1.22	108.13	66.62	0.86	0.45	112.55	96.26
All	39	0.71	0.35	15.01	10.66	0.62	0.46	17.85	11.92

The margin of error is higher in the TOWT model, indicating more variation in claimed savings. Even after filtering projects for goodness of fit and reliability of estimates, the margin of error is higher for individual strata. Figure 5 and Figure 6 compare claimed and modeled savings of the filtered subset. Thus, even when filtering projects for the applicability of an Option C analysis, at the individual project level there can be substantial variation compared to project claims.

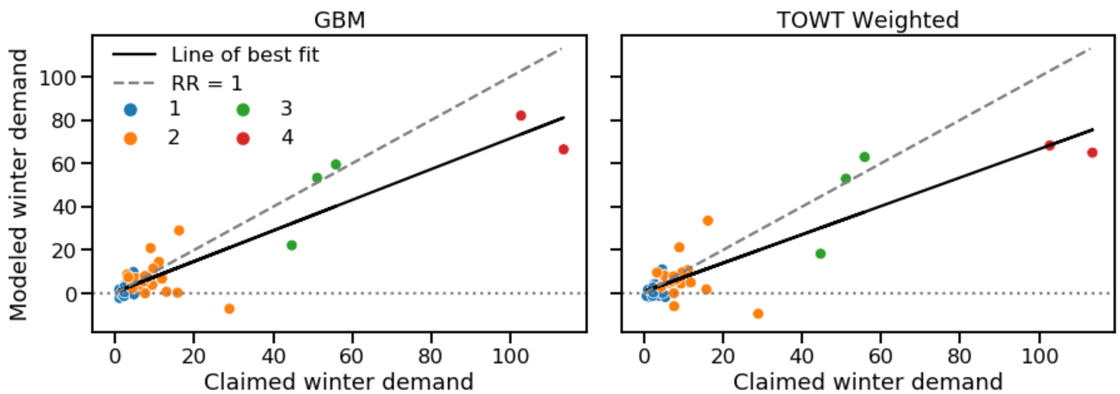

Figure 5. Claimed winter demand savings, compared to modeled savings.

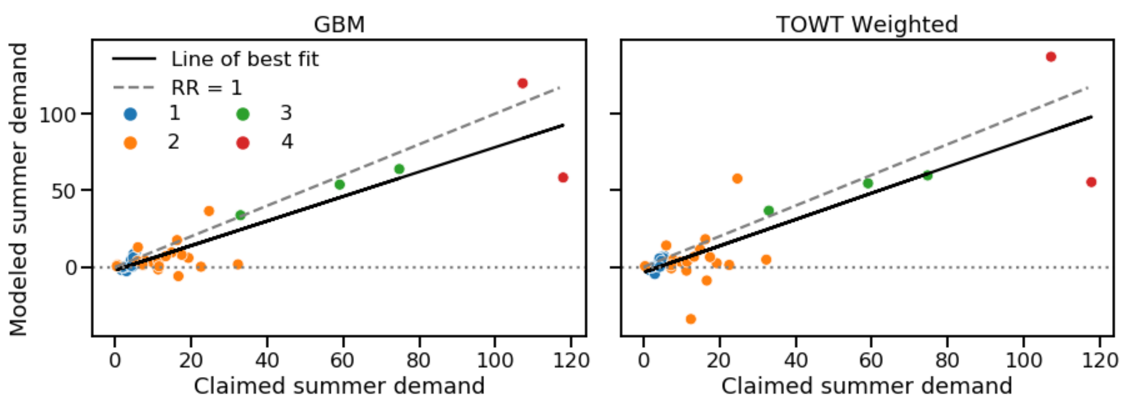

Figure 6. Claimed summer demand savings, compared to modeled savings.

The realization rates for all strata in Table 3 and Table 4 are similar to realization rates reported by third-party evaluators for this sector of the portfolio. For example, the custom C&I realization rate for program year 2018 was 0.64 for winter demand savings, and 0.85 for summer. Direct comparison of these results is not appropriate because the samples are not comparable; this is further discussed, below.

The study's primary objective was to determine how M&V 2.0 could be applied more broadly to evaluation of efficiency portfolio performance. The team originally thought that the automation behind M&V 2.0 would allow evaluation of larger sample of projects, at a lower cost than the current sample process that uses on-site metering for evaluation. The team reasoned that the central tendency of a larger sample would be a stable estimate for population realization rates. For an evaluation to meet M&V requirements, a representative population sample must be evaluated and meet precision requirements. The team found that where the data requirements, modeling requirements and results feasibility requirements were *not* met, the data were far too noisy and biased to meet M&V evaluation standards. Using M&V 2.0 for projects where requirements are not met can introduce substantial bias into the evaluation, even when applied at scale. The team filtered the projects on data availability and model validity checks from the population of custom C&I retrofit projects. The factors driving filtering are not random, and thus the resulting project subset was no longer representative of the entire custom C&I retrofit population. Verification results from this filtered sample would thus not be appropriate to apply across all custom C&I retrofit projects. To meet M&V requirements, Option C would need to be supplemented by other M&V methods. This does not mean that Option C has no applicability to FCM M&V and other evaluation use cases, but it does mean that Option C feasibility should be determined at the appropriate step in the evaluation scoping process. Following the creation of a representative project sample, those projects could be pre-screened to test appropriateness for Option C. For inappropriate projects, alternative M&V methods could be used. This flexibility allows for the benefits of Option C, while ensuring the evaluation results can be applied to the project population under evaluation.

M&V 2.0 Pre-screening

Using M&V 2.0 to pre-screen projects involves gathering relevant data, verifying data requirements, and performing whole-facility savings analysis with model diagnostics. Data requirements and diagnostic metrics from this analysis can be an effective first pass of an evaluation sample to identify appropriate projects and sites for using M&V 2.0. Figure 7 shows how pre-screening outputs can inform project evaluation plans, by determining whether on-site metering is required.

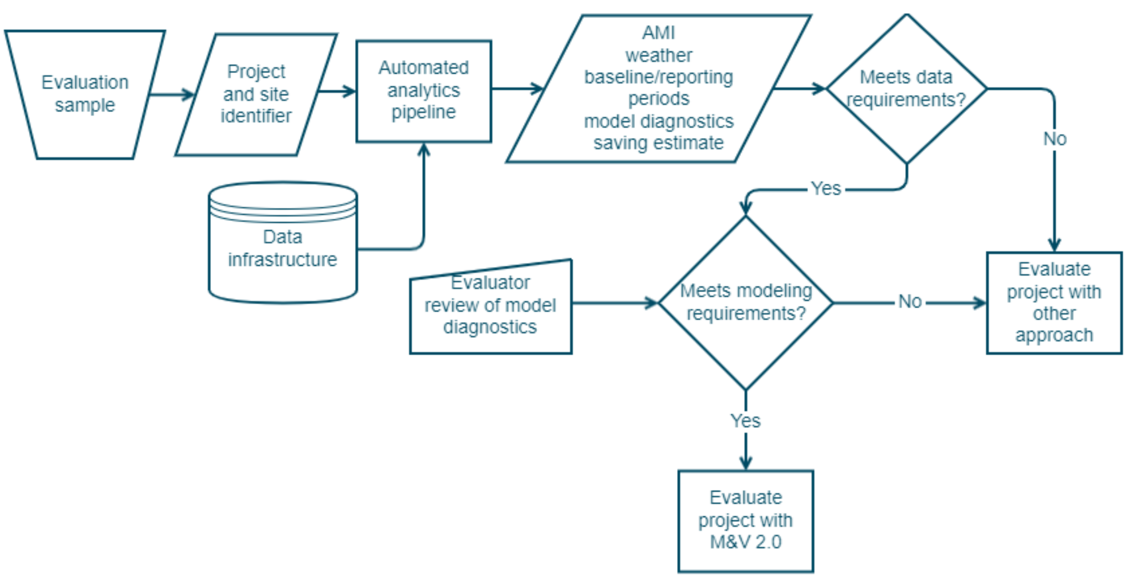

Figure 7. Project evaluation process using M&V 2.0 pre-screening.

Pre-screening can reduce the amount of support needed to conduct a third-party evaluation. It enables:

- Automated collection of project details, AMI data, and weather data for whole-facility analysis.
- Cost-effective prioritization of direct metering resources for projects not suitable for whole-facility analysis.
- Cost-effective preliminary insights on project savings potential. Results can be easily updated throughout reporting period.
- Ensuring AMI data are used wherever whole-facility analysis is valid. This minimizes the amount of necessary, costly on-site metering.
- Automatically generated diagnostics, such as residual and cumulative savings plots, providing insight into non-routine events that could require modifications for valid Option C analysis.

Based on these findings, VEIC and Efficiency Vermont plan to work with regulators and third-party evaluators to use M&V 2.0 to provide efficiencies in the evaluation processes. Pre-screening with M&V 2.0 can indicate potential M&V 2.0 program participants and provide near real-time insights into efficiency program or load management effects. VEIC plans to continue to work with program implementers to integrate near real-time M&V 2.0 to allow program staff to obtain preliminary insights into savings potential, verify savings persistence, and monitor building operations.

Conclusion

The study shows strong potential for using advanced, whole-building M&V methods to estimate savings for many different types of projects within Efficiency Vermont's C&I retrofit project portfolio. Approximately 40 percent of the C&I retrofit projects with sufficient AMI data met all uncertainty requirements; more than 60 percent of projects met the fractional savings uncertainty and bias requirements, using hourly models. Therefore, this analysis method can provide cost-effective and reliable insights into many of Efficiency Vermont's savings measures.

Efficiency Vermont's investments in data infrastructure have allowed high levels of automation to be integrated into M&V 2.0 pre-screening proposed in this research. The data infrastructure also allows easy expansion of M&V 2.0 across many other use cases beyond the FCM evaluation. Making M&V 2.0 results more readily available and interpretable for energy consultants, efficiency program staff, and utilities can inform ways to obtain the most impact from energy efficiency or load management.

Although M&V 2.0 can cost-effectively scale to larger numbers of projects, it should not be relied on as the sole method for evaluation. On-site metering is still required to supplement M&V 2.0 in meeting M&V requirements. M&V 2.0 can be particularly effective for pre-screening projects for appropriateness for Option C methods and can support substantial process improvements. This study has highlighted the strengths and limitations of applying M&V 2.0 to real-world evaluation use cases. Although the limitations do not diminish the promise and efficacy of M&V 2.0, they inform implementation, design, and scoping. M&V 2.0 is a cost-effective and information-rich first look at data, and can inform evaluation strategies, program design, and insights into measure performance in ways no other method has yet done.

References

Crowe, Eliot, Jessica Granderson, and Samuel Fernandes. 2019. "From Theory to Practice: Lessons Learned from an Advanced M&V Commercial Pilot." Presented at the 2019 International Energy Program Evaluation Conference (IEPEC). https://www.iepec.org/wp-content/uploads/2020/03/Crowe-Eliot-From-Theory-to-Practice-Lessons-Learned-from-an-Advanced-MV-Commercial-Pilot.pdf.

EVO (Efficiency Valuation Organization). 2020. Application Guide on Non-Routine Events and Adjustments: International Performance Measurement and Verification Protocol®. Washington, DC: EVO. https://evo-world.org/images/denisdocuments/NRE_TOC.pdf.

EVO (Efficiency Valuation Organization). 2016. Core Concepts: International Performance Measurement and Verification Protocol®. Washington, DC: EVO. https://evo-world.org/en/products-services-mainmenu-en/protocols/ipmvp.

Franconi, Ellen, Matt Gee, Miriam Goldberg, Jessica Granderson, Tim Guiterman, Michael Li, and Brian Arthur Smith. 2017. "The Status and Promise of Advanced M&V: An Overview of 'M&V 2.0' Methods, Tools, and Applications." Berkeley, CA: Lawrence Berkeley National Laboratory. https://eta-publications.lbl.gov/sites/default/files/lbnl-1007125.pdf.

Granderson, Jessica, Eliot Crowe, Samuel Fernandes, and Samir Touzani. 2019. "Suggested Guidance Criteria for Documentation and Acceptance of Whole-Building Savings Results." Berkeley, CA: Lawrence Berkeley National Laboratory. https://buildings.lbl.gov/sites/default/files/AcceptanceCriteriaAndDocumentation%202019%2006%2030.pdf.

Kelly, Anna, and Craig Sinnamon. 2020. "Detecting Savings Under 10% Using IPMVP Option C." M&V Focus 7 (November). Washington, DC: EVO. https://evo-world.org/en/news-media/m-v-focus/883-october-2020-m-v-focus-issue-7/1192-detecting-savings-under-10-using-ipmvp-option-c.

LBNL-ETA (Lawrence Berkeley National Laboratory Energy Technologies Area). 2020. "RMV2.0." https://github.com/LBNL-ETA/RMV2.0.

Linux Foundation Energy EM2. 2019. "EEmeter: Tools for Calculating Metered Energy Savings." San Francisco: Linux Foundation. https://eemeter.openee.io/.

Linux Foundation Energy EM2 and Joint Development Foundation. 2018. "CalTRACK Methods." San Francisco: Linux Foundation. https://docs.caltrack.org/en/latest/methods.html.

Touzani, Samir, Jessica Granderson, and Samuel Fernandes. 2018. "Gradient boosting machine for modeling the energy consumption of commercial buildings." Energy and Buildings 158 (6). Berkeley, CA: Lawrence Berkeley National Laboratory: 1533-43. https://www.sciencedirect.com/science/article/abs/pii/S0378778817320844?via%3Dihub.

Webster, Lia. 2020. "IPMVP's Snapshot on Advanced Measurement and Verification." Washington, DC: EVO. https://evo-world.org/images/corporate_documents/NRE-NRA_White_Paper_Final_2701.pdf.