

Massachusetts Green Retrofit Initiative Savings Analysis Report

John Rutherford, Daniel Teague, and Barun Singh
WegoWise

July 1, 2015

Summary

The Massachusetts Green Retrofit Initiative (MAGRI) provides technical assistance for multifamily building owners to reduce their energy and water usage. MAGRI is funded by The Barr Foundation and the Department of Housing and Urban Development, and operated jointly by Local Initiatives Support Corporation (LISC) and New Ecology, Inc (NEI).

This document is the final analysis of pre- and post-retrofit utility usage, performed by WegoWise, comprising data as recent as May 2015. It outlines the latest average and cumulative savings information, an analysis of potential correlations, a few notable case studies, and a description of the WegoWise process for measurement and verification.

The results of MAGRI are significant. Over half of the apartments affected by upgrades saw greater than 20% savings. Retrofits affecting gas heating systems saw an average of 23% yearly savings, while electric retrofits, comprised primarily of lighting upgrades, saw 29% average yearly savings. Savings have even been extracted from already-efficient buildings, beyond the level predicted by another similar analysis of multifamily retrofits. Additionally, savings correlate moderately with project cost: the more capital-intensive a project, the more it tends to reduce usage and CO₂. In all, 3700 tons of CO₂ have been saved to date thanks to the MAGRI efforts.

Contents

1 Upgrade Analysis Overview	3
2 Gas Upgrades	5
2.1 Savings Trends with Usage	5
2.2 Trends with Building and Upgrade Characteristics	7
3 Electric Upgrades	8
3.1 Savings Trends with Usage	8
3.2 Trends with Building and Upgrade Characteristics	8
4 Carbon Impact	9
5 Savings Trends with Project Cost	11
6 Savings Summary	13
7 Recommendations	14
8 Appendix: End-Use-Specific Case Studies	15
8.1 End-Use-Specific Retrofits	16
8.2 Changing Building Conditions	16
9 Appendix: Measurement & Verification Methods	18
9.1 M&V Procedure	18
9.2 Data Quality	19
9.3 Calendarization	19

1 Upgrade Analysis Overview

WegoWise analyzed utility data from 133 retrofits covering 129 unique buildings and a variety of cases: new gas heating systems, common area lighting upgrades, oil-to-gas conversions, and more. Some retrofits affect multiple buildings, depending on the utility metering setup, while others can be merged together, such as the installation of a new boiler and the soon-after commissioning.

For all retrofits, NEI provided WegoWise with the upgrade scope of work, construction dates, and building information. Since all of these buildings are tracked in the WegoWise application already, additional characteristics could be added as needed, such as age and size.

Data integrity checks were performed as the first task for the analysis. The WegoWise team ensured that the monthly utility data was up-to-date and accurate, adding new data when automatic data-fetching had failed. Due to privacy concerns, tenant data for most buildings is not available, so upgrades affecting tenant living spaces could not be evaluated. (Most measures were performed on master-metered buildings, though, so such tenant data was not required.)

Water system retrofits are omitted from this analysis, due to incomplete data.

Table 1 shows the distribution of upgraded properties across building type, age, and retrofit type. The statistics show that lighting retrofits are just as popular as more substantial heating system upgrades, across all types of buildings. Older pre-War-era buildings represent only one fifth of the total retrofits, but this is over twice the proportion of such buildings in Massachusetts (10% of multifamily buildings in the state, according to the WegoWise database). Modern buildings are also overrepresented: 29% in the state, 39% in this sample of retrofits.

Of these upgrades, roughly half occurred more than a year ago. These retrofits with more than a year of savings information are the most valuable for informing future projects, as the buildings have undergone a full four seasons of Massachusetts weather. Most of the analysis presented here focuses on such retrofits, although cumulative savings are reported for all of the available retrofits with sufficiently clean data and good fits. (Section 9 describes the quality checks in greater detail.)

	Total	High-rise	Mid-rise	Low-rise
Buildings	129	5	33	91
Apartments	2879	667	846	1366
Bedrooms	4571	763	1260	2548
Apartment Distribution	100%	23%	30%	47%
Age:	Percent of apartments:			
Pre-1946	21%	0%	18%	3%
1946-1979	40%	13%	5%	22%
Post-1979	39%	10%	7%	22%

Table 1: Breakdown by building type. High-rise is defined as 10 or more stories, mid-rise as 5 to 10 stories, and low-rise as 4 or fewer stories. Multiply retrofitted buildings (e.g., new boiler and new lighting) are counted only once.

The savings analysis, which is detailed further in an appendix (Section 9), proceeded as follows: a model of weather-dependent pre-retrofit usage was created for the monthly meter data for the meters affected, then savings were computed as the post-retrofit difference between the model and the actual usage. Retrofits with utility data that did not pass quality requirements were removed from the analysis. Approximately one tenth of the retrofits had to be removed, and are not included in the savings reported below.

Figure 1 shows a successful gas boiler upgrade, and is indicative of the quality of the modeling, even with coarse monthly data.

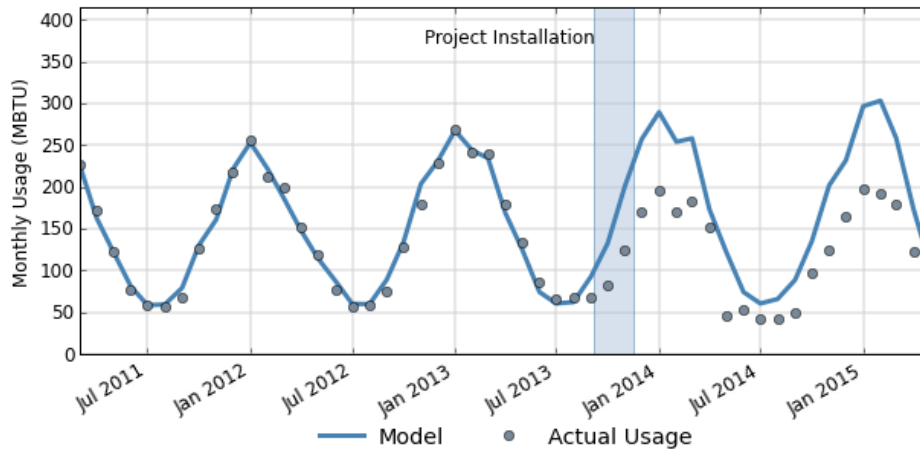


Figure 1: The linear degree-day model, detailed in Section 9, fits the monthly data well. The deviation from the model in the post-retrofit period shows the savings.

2 Gas Upgrades

On average, the gas heating retrofits save 23% on energy, with a standard deviation of 12%. For both gas and electric retrofits, NEI is investigating cases with low or negative savings. Continual, on-site monitoring is invaluable for completing individual stories of successes and failures, in a way that summary graphs and statistics cannot.

2.1 Savings Trends with Usage

Figures 2 and 3 show two views of how savings vary with pre-retrofit energy use intensity (EUI). (Only upgrades with 12 months of post-retrofit data are included.) In Figure 2, we can see a slight linear dependence of savings EUI on pre-retrofit EUI, a trend which has been reported in similar analyses, such as the 2012 Deutsche Bank multifamily analysis¹. The gray line on Figure 2 shows the fit from the Deutsche Bank study. It ‘predicts’ zero savings below roughly 60 kBTU/sqft/year. However, the MAGRI data show savings extending well below that mark, demonstrating that even more efficient buildings have savings potential. The linear fit shown in Figure 2 predicts zero savings below 17 kBTU/sqft/year, and a weaker dependence on pre-retrofit EUI (i.e., a shallower slope).

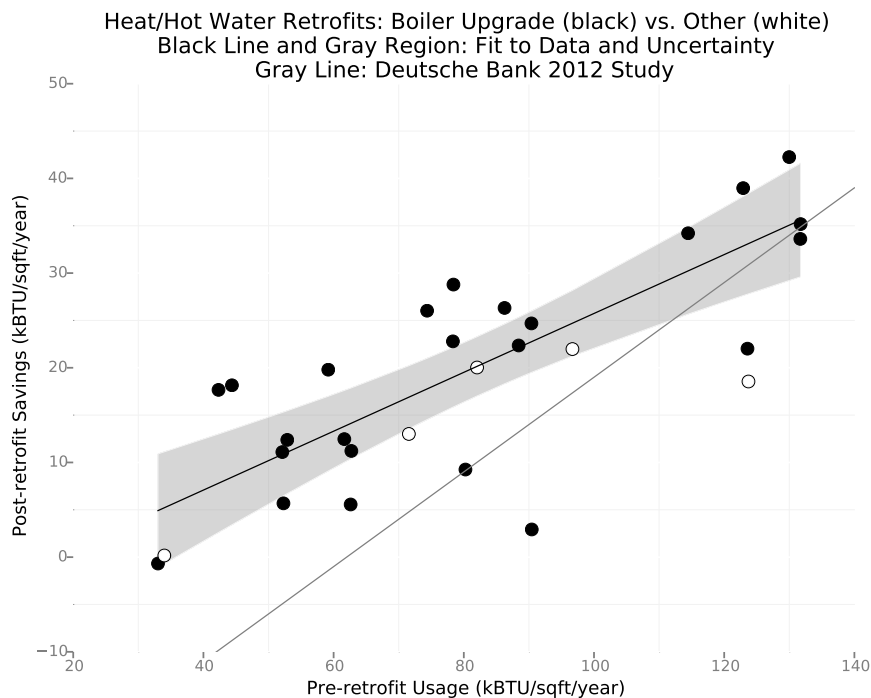


Figure 2: Savings from gas retrofits (in terms of EUI) show a weak linear dependence on pre-retrofit EUI: savings = $0.31 \times \text{usage} - 5.3$ kBTU/sqft/year. The gray region shows the 95% confidence interval around the linear fit (black line). The Deutsche Bank study’s fit to similar data (heating and hot water retrofits of gas, oil, and steam systems) is shown as the gray line: savings = $0.51 \times \text{usage} - 31$ kBTU/sqft/year.

¹‘Recognizing the Benefits of Energy Efficiency in Multifamily Underwriting’
https://www.db.com/usa/img/DBLC_Recognizing_the_Benefits_of_Energy_Efficiency_01_12.pdf

Figure 3 shows a different view of the same result, with mean percent savings from the latest year on the y -axis. In this view, the dependence of savings on pre-retrofit EUI is very weak. The low predicted x -intercept makes the slope—the proportion—more informative; if it were zero, we could apply a constant percent savings to buildings across the whole range of EUI. Percent savings—as opposed to absolute savings—can be a more natural statistic to consider, since building upgrades typically deal with multiplicative changes, such as the increased efficiency of a boiler or the reduction in percent usage of LED bulbs.

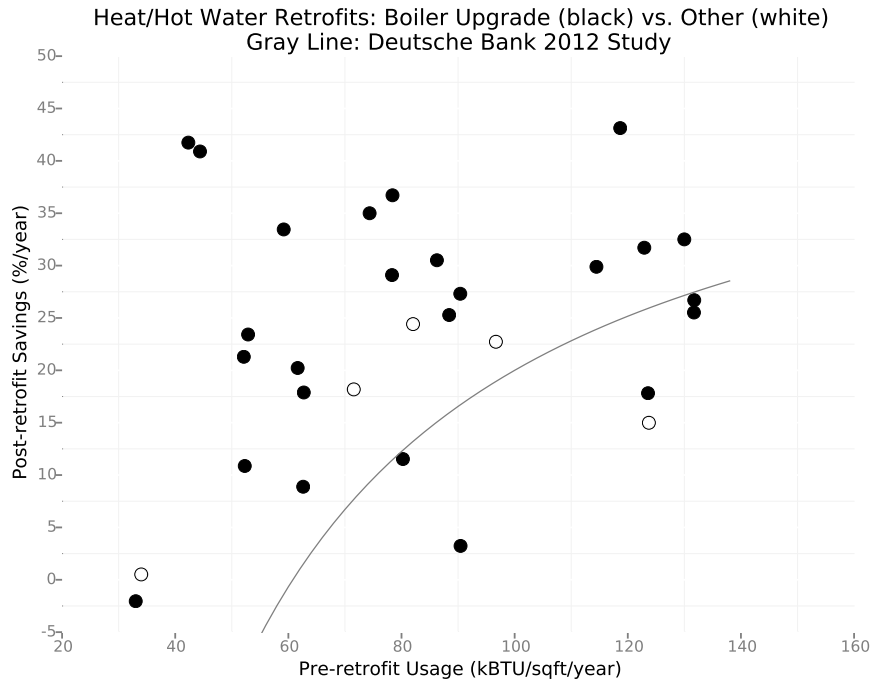


Figure 3: The gas retrofits show consistent savings ($23 \pm 12\%$ per year), with only a weak dependence on pre-retrofit EUI. When plotted as percent savings, the deviation from the Deutsche Bank results becomes pronounced.

2.2 Trends with Building and Upgrade Characteristics

A variety of building and retrofit characteristics were examined to see what possible factors could influence (or at least correlate with) the success of an upgrade. Few patterns emerged, though. No significant trend in savings can be seen with building age (neither with year built or grouped into vintages). Building type—low-, mid-, and high-rise—also didn't correlate with savings, nor did building size or number of bedrooms.

Neither retrofit duration, which may indicate the depth of work, nor the retrofit date, which could be a proxy for cumulative experience, showed any trend with the percent savings. Only a few moderate rehabilitation projects ('mod rehabs') were performed², and those six showed similar median monthly percent savings to the rest of the retrofits.

Only one distinct trend emerges: fuel conversion retrofits display significantly less percent savings: only a median of 12% savings, compared to 24% for all other heating system upgrades. However, this difference could be due entirely to the circumstances of the individual projects: one project—comprising most of the fuel-conversion retrofits—switched the (tenant-paid) electric hot water heaters to the new (owner-paid) gas system, resulting in higher baseload usage; another project could not upgrade the indirect hot water tanks alongside the heating system upgrades due to lack of funding. Yet, these projects yield twice the CO₂ savings, since they replace oil heating systems with natural gas (see Section 4).

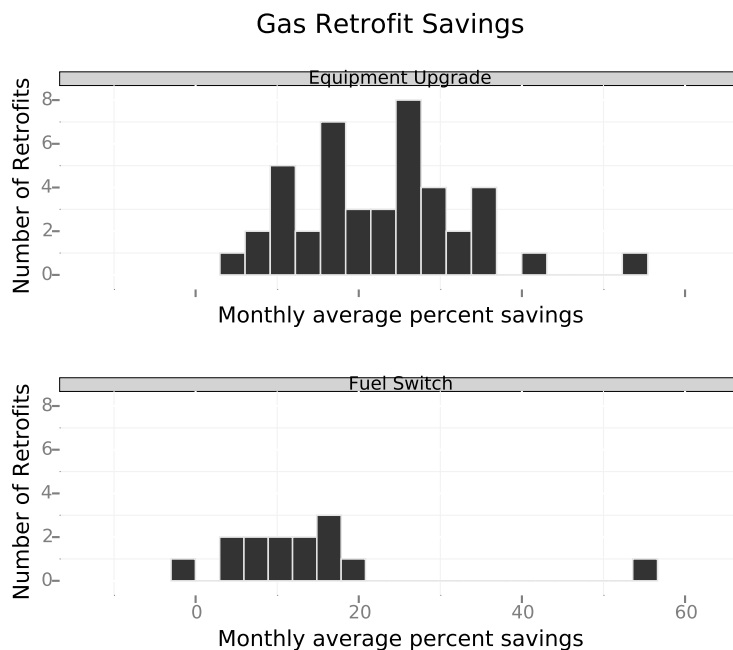


Figure 4: Oil-to-gas retrofits tend to save less, percent-wise, but Figure 8 shows that they save more CO₂. The results here are averaged per-month since not many fuel switching projects have matured to one year.

²Mod rehabs are characterized by less invasive upgrades than complete gutting and rebuilding; for instance changing the HVAC or lighting system alone, without changing the skin of the building.

3 Electric Upgrades

The electric retrofits, which are mostly lighting upgrades, perform better than gas heating retrofits—29% yearly savings on average—but with much greater variability, nearly 22%. (Successful retrofits—those with positive savings—yield savings of $33 \pm 15\%$.) However, since Massachusetts buildings overwhelmingly use gas heat, the total electricity usage is lower, and thus so is the impact of these successful retrofits.

Six electric upgrades affected the heating systems: one building received new pumps, trimming the winter electric usage, and the others had heat pumps installed. These retrofits saw a median monthly savings of 18%, closer in line with the gas heating upgrades (23%) than the electric lighting upgrades (37%).

3.1 Savings Trends with Usage

The electric percent savings display no trend with pre-retrofit usage intensity (Figure 5). The Deutsche Bank study found average electric savings of 7%, but did not report any trend with EUI. Additionally, no difference is observed when the same buildings’ gas systems have been upgraded. (Such a trend could reveal an owner- or contractor-dependent savings effect.)

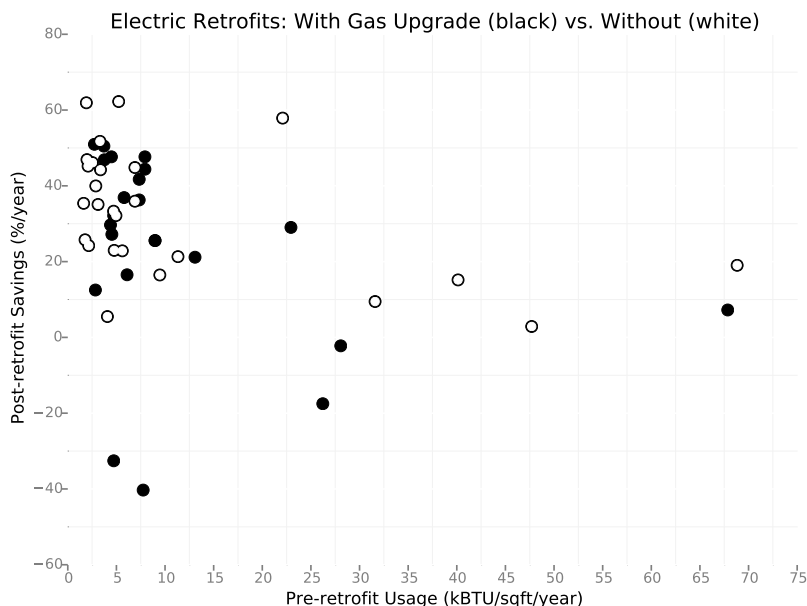


Figure 5: Electric upgrades show similar average success to gas upgrades, but with greater variability ($29 \pm 22\%$ per year). No significant trend with usage intensity is discernible.

3.2 Trends with Building and Upgrade Characteristics

As with the heating system upgrades, no significant trend was found between the percent savings for electric retrofits and building age, building type, building size, retrofit duration, retrofit timing, or mod rehab status.

4 Carbon Impact

In terms of environmental impact, the savings demonstrated above translate to over one ton of CO₂ per apartment per year³. Figure 6 shows the distribution of CO₂ savings for gas heating and hot water upgrades. The CO₂ savings from electric upgrades in Figure 7 average to about one tenth that of the gas retrofits.

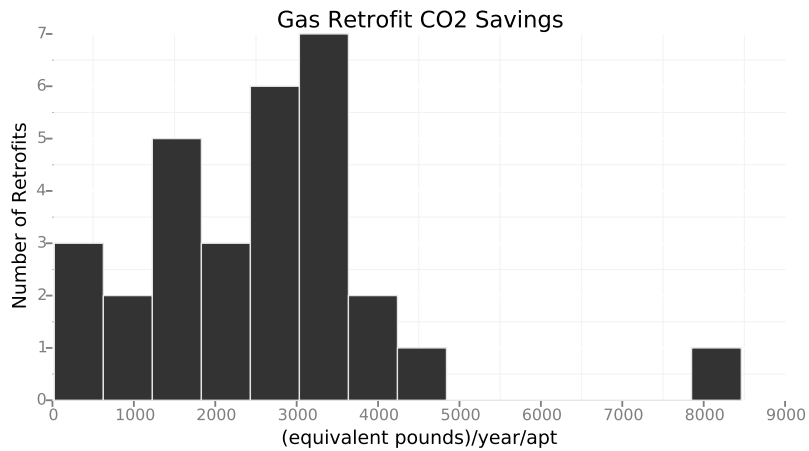


Figure 6: The CO₂ savings for gas retrofits average to over one ton per bedroom per year. The outlier is a very successful (43% yearly savings) upgrade on a heat/hot water system that served only 10 apartments, resulting in higher per-apartment savings.

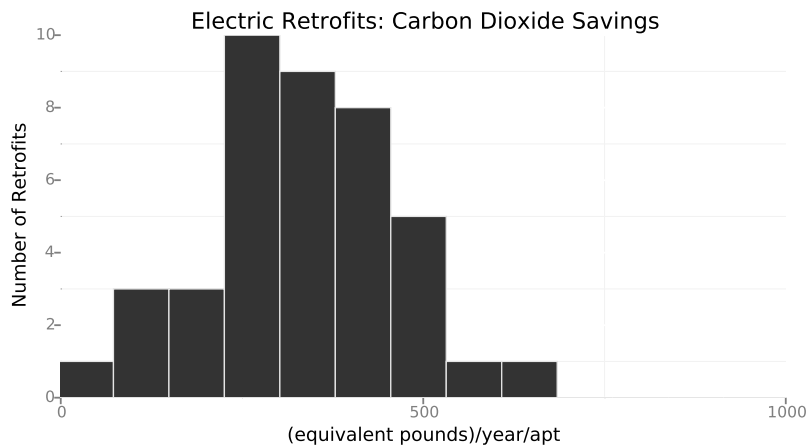


Figure 7: Due to the lower raw usage, electric retrofits save only about one-tenth the CO₂ as gas retrofits. Four unsuccessful upgrades are not shown.

³The CO₂ conversions were made with values provided by the Energy Information Administration: http://www.eia.gov/environment/emissions/co2_vol_mass.cfm, <http://www.theclimaterestory.org/downloads/2012/05/eGRID2012.pdf>

Figure 8 shows the marked impact of switching a heating system from oil to gas. The typical CO₂ savings are approximately double those of the rest of the gas heating upgrades: a median of 350 pounds per apartment per month, compared to 180. These savings could be even larger, were it not for the limitations of the individual fuel-conversion projects (noted in Section 2.2).

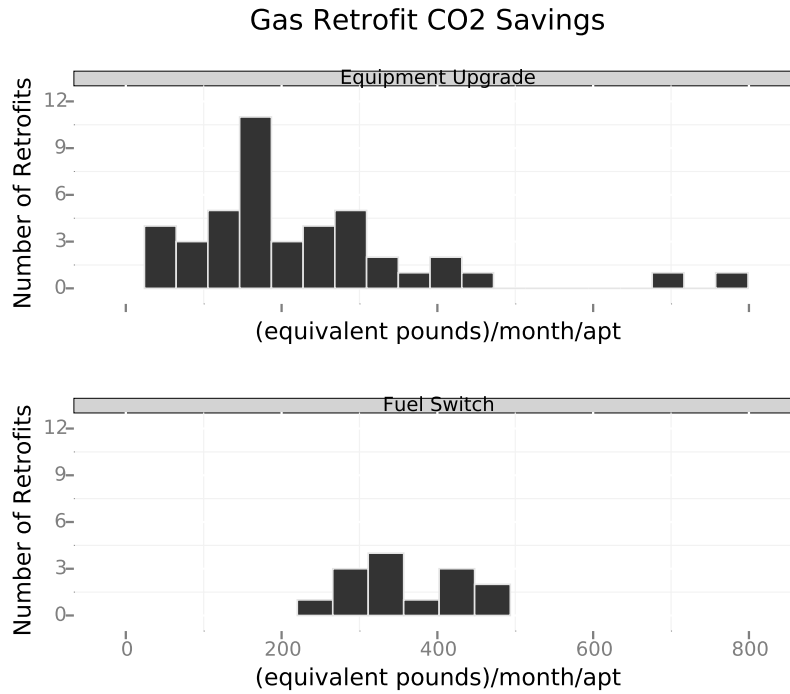


Figure 8: While oil-to-gas retrofits save less raw energy (about half) compared to the other gas system upgrades in this analysis, they save twice the CO₂. The results here are averaged per-month since not many fuel switching projects have matured to one year.

5 Savings Trends with Project Cost

Funding information, provided by NEI, allows us to test the adage ‘the more you spend, the more you save.’ The retrofits analyzed in this program seem to prove the adage true, to some degree.

Below are two views of the same conclusion, showing the trends of increasing percent savings and CO₂ savings with increased project funding. These results group individual retrofits into projects, as funding is allocated at that level. All retrofits are included in the project-wide savings, even if work completed within the last 12 months.




Figure 9: Project-wide percent savings correlate with the investment. The correlation coefficient is 0.45, halfway between no correlation (0) and perfect correlation (1). A linear fit to the data yields the relation: percent savings = $17 + 1.8 \times \text{cost}$, where the cost is measured in \$1000's/bed. (The grey region encloses the 95% confidence interval around the linear fit.)




Figure 10: Project-wide CO₂ savings also correlate with the investment. The correlation coefficient is 0.45, stronger than the relationship between energy savings and funding (0.40; not shown). A linear fit to the data yields the relation: $\text{lbs/bed/month} = 47 + 10 \times \text{cost}$, where the cost is measured in \$1000's/bed.

6 Savings Summary

Tables 2 and 3 summarize the above analysis and show the progress towards the savings target. The project lifetime savings are aggregated for all upgrades of a given fuel, and are summed together even when one year of post-retrofit data is not available. The average savings, however, are reported only for the latest year, and omit any projects which occurred less than a year ago.

The electric savings surpassed the 30%-better-than-30% program goal, while the gas savings have nearly achieved it as well (40% exceed 25% savings). The goal could be surpassed as retrofits for more buildings mature past 12 months. Overall, this project has saved the equivalent amount of CO₂ of burning 1800 tons of coal, passenger cars driving 8 million miles over a year, or the carbon sequestration power of 2800 acres of forest per year⁴.

	Savings To Date		Average Savings, Latest Year	
	Usage	CO ₂	Usage	CO ₂
Gas	43,000 MMBTU	6,100,000 lbs	23%	2600 lbs/apt
Electric	6500 MMBTU	1,400,000 lbs	29%	370 lbs/apt

Table 2: Overall savings results.

Annual Energy Savings	Gas (Percent of 1056 apts.)	Electric (Percent of 1379 apts.)
> 10%	72%	65%
> 20%	54%	50%
> 30%	18%	34%

Table 3: Progress towards savings targets, for projects with at least one year of post-retrofit data.

⁴<http://www.epa.gov/cleanenergy/energy-resources/calculator.html>

7 Recommendations

Over the course of the analysis, several possibilities arose for improving the nuance and breadth of the results, for any future program.

1. Collect additional categorical variables, such as contractor name or whether staff training for new equipment was performed. Controlling for additional variables relating to the implementation and ongoing operations could improve our ability to understand underlying correlations between retrofit results and building characteristics.
2. Improve access to and standardization of utility company energy and water data.
 - (a) For this analysis, WegoWise relied on utility company websites to collect energy and water data. In some instances, utility companies created barriers to prevent building owners—and consequently, WegoWise—from easily accessing historical usage data. For example, some utilities require owners to sign up for e-bills in order to view usage history. This practice disrupts building owners’ normal accounting and bill payment process, and it forces them to choose between getting a paper bill and having electronic access to utility history.
 - (b) Access to aggregated and anonymous tenant energy data should be readily available. A portion of the buildings in this analysis split their electric use between owner- and tenant-paid electric meters. Data from tenant-paid meters is much more difficult to access, limiting the ability to analyze measures which affect tenants. Utilities should offer an easy way for building owners to request tenant usage information, while respecting tenants’ privacy.
3. Track costs at the retrofit level. This would allow a finer-grained look at savings trends with funding, which Section 5 could only examine at the project level.

8 Appendix: End-Use-Specific Case Studies

The WegoWise method to calculate savings only reports *total* savings after the retrofit (see Section 9.1). For some cases, though, only one particular end-use component is expected to change; e.g., a hot water upgrade may only affect the gas baseload, and not the gas heating component. In order to investigate such cases, this section departs from the usual methodology, employing the procedure detailed in the next two paragraphs.

To determine end-use-specific savings, two regression models must be built: one for before the retrofit, and one for after. Then, the model parameters (typically: baseload, heating setpoint, and heating slope) can be compared. The dependence of the savings on these mathematical models has thus *increased*: instead of fitting the pre-retrofit data with three parameters and comparing the post-retrofit data to the model-predicted data, we are fitting all the data with six parameters and comparing these fitted parameters directly. Moreover, these parameters are highly correlated with each other: for instance, a model with a higher baseload and lower setpoint can fit the data just as well as a model with a lower baseload and higher setpoint. Such tradeoffs could bias the savings calculation: continuing the example, did the building reduce its baseload consumption, or did the retrofit increase its setpoint? (For these reasons, WegoWise typically avoids end-use-specific M&V.)

An additional component must be added to the analysis to account for the parameter correlations. We employ a method called ‘bootstrapping,’ which re-samples the monthly data and creates a regression model many times. This statistical method is widely used to determine the uncertainty of fitted parameters⁵. The reported uncertainties enclose the middle 68% of the re-sampled results, corresponding to ‘plus-or-minus one sigma’ for a Normal distribution. Figure 11 shows an example bootstrap analysis used to calculate end-use savings. Some of these uncertainties can be quite large, revealing the pitfalls of simply comparing two best-fit regression models.

The results below compare the parameters from pre- and post-retrofit regression models, as outlined above. For ease of reference, the normal M&V method will be referred to as the ‘one-model method’, and the end-use-specific method with the bootstrapped uncertainties will be called the ‘two-model method’.

⁵‘Bootstrap Methods: Another Look at the Jackknife’, B. Efron, 1979.
<http://projecteuclid.org/euclid.aos/1176344552>

8.1 End-Use-Specific Retrofits

- Margolis: The lighting retrofit in this electrically heated building should only affect the baseload. Indeed, while the one-model method calculates 6% savings total, the two-model method isolates an $18 \pm 4\%$ savings in the baseload.
- Harvard Place: This property received new boilers and a new hot water system, for a very successful 32% savings, according to the one-model method. The two-model method shows that most of the savings likely comes from the hot water: the baseload decreases by $50 \pm 7\%$, but the heating component only decreases by $9_{-39}^{+28}\%$.
- Pondview Apartments: In addition to heating system upgrades, a solar thermal hot water system was installed. Thanks to the solar collectors, the gas baseload decreased by $51_{-14}^{+19}\%$, and, due to the newer heating system, the heating slope decreased by $34 \pm 3\%$. However, the heating setpoint after the retrofit is higher by 6 ± 3 degrees F, washing out the heating energy savings, which change by $+13_{-30}^{+46}\%$ after the retrofit. The one-model method found 26% total savings.

8.2 Changing Building Conditions

- 76 Byers at New Court Terrace: The ventilation was increased after a lighting retrofit, resulting in a 19% growth in usage, according to the one-model method. The two-model method identifies a $41 \pm 12\%$ increase in the baseload usage, which can be attributed to the ventilation.
- 62 Elm Hill at Blue Mountain: After lighting upgrades were performed, a new pump was installed (outside of NEI's purview), adding to the electric usage during winter months. The one-model method found an 18% increase in total usage. The two-model method saw an increased, but highly uncertain change in the baseload, by $+93_{-51}^{+78}\%$. The heating slope change, which indicates how much more the heating system must work to keep the building warm with increased influx of cold outside air, rose $9 \pm 20\%$.

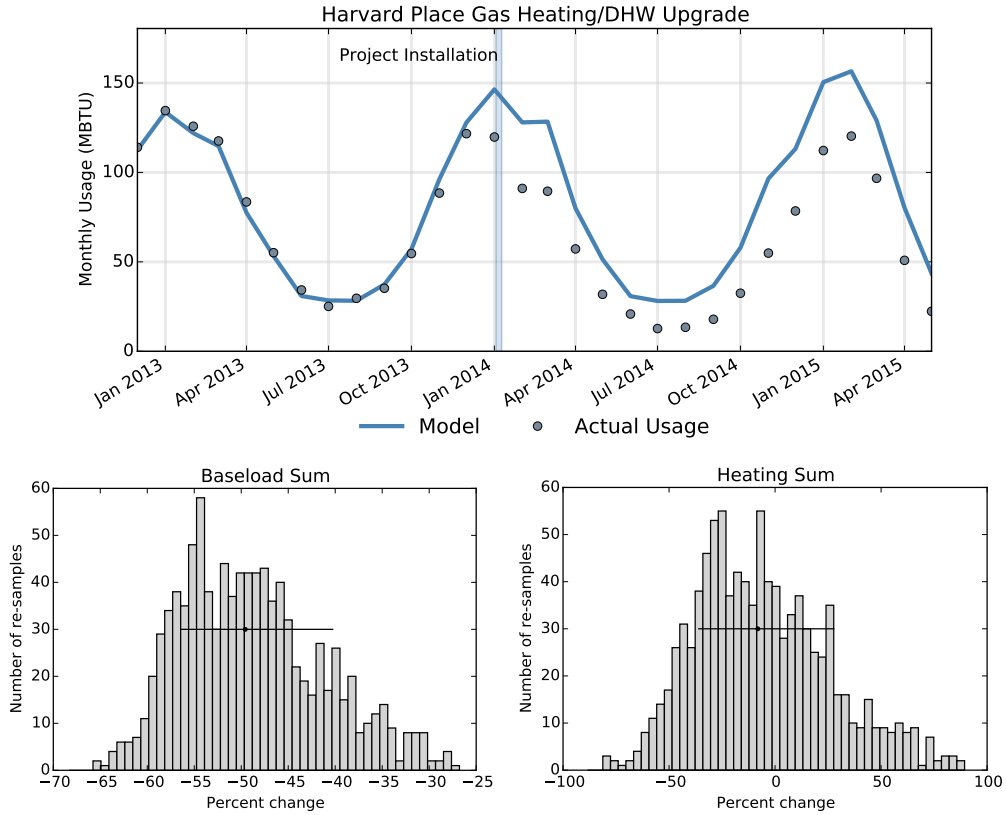


Figure 11: This upgrade to the entire heat and hot water system resulted mostly in baseload savings alone. The top graph shows the change in monthly usage, with the pre-retrofit model underlaid for reference. The left panel shows a histogram of the calculated change in baseload, after fitting models 1000 times to bootstrap samples. The histogram approximates the sampling distribution of the statistic (in this case, the difference between the modeled baseloads post- and pre- retrofit). The horizontal whiskers enclose the middle 68% of the distribution. The right panel shows the same histogram, but for the heating component, which can be calculated from the fitted heating setpoints and heating slopes.

9 Appendix: Measurement & Verification Methods

9.1 M&V Procedure

The WegoWise procedure for Measurement and Verification (M&V) consists of two steps.

1. A variable degree-day regression model is fit to the pre-retrofit data.
2. The post-retrofit savings are reported as the projected usage (calculated using the regression fit parameters and the actual post-retrofit weather) minus the building's actual usage.

These steps are briefly detailed below. The method is agnostic to the type of input data, whether it be gas usage from several buildings, one electric meter normalized by square feet, or outdoor water usage.

The degree-day model is an industry-wide-standard PRISM model⁶ with a variable heating (and/or cooling) balance point. That is, a range of balance points are explored, and the one yielding the best fit to the resulting degree-day-versus-usage linear model is chosen. Daily temperatures from the building's ZIP code are used to calculate degree days. The model fits weather-independent data (baseload electric usage, water usage) with the baseload parameter only.

When performing M&V analysis, a range of pre-retrofit baselines of at least 12 contiguous months is explored. The pre-retrofit range with the best fit (to the data within that range) determines the PRISM model for the building. Post-retrofit savings are determined by running the weather through the model to produce expected usage, then simply subtracting the actual usage.

Note, we do not perform end-use-specific M&V by creating a pre- and post-retrofit PRISM model. That is, the WegoWise method does not report savings for heating only, or for baseload only. Instead, we report total savings for several reasons:

1. Not many buildings have a full year of post-retrofit data to build a second regression model.
2. The model parameter uncertainties are highly correlated and may produce spurious results without careful error analyses which are beyond the scope of this project. Moreover, estimates of the usage are more precise than estimates of the underlying parameters (see 'Calendarization' below).
3. Reporting end-use-specific savings tends to inflate the success of a project.

The regression model and M&V automation software have been used to deliver hundreds of customized analyses to WegoWise customers.

⁶http://www.marean.mycpanel.princeton.edu/~marean/images/prism_intro.pdf

9.2 Data Quality

Several data quality requirements are enforced. First, before the modeling can even begin, all affected utility meters must be accounted for, and their data must look ‘reasonable’ to a trained energy analyst. (Unfortunately, this step is ill-defined, but it is necessary: analysts can more readily detect when a gas meter is mis-attributed as an electric meter, or when a reading is incorrect, for instance.) Second, all weather-dependent models with $R^2 < 0.7$ are removed from the final results. This level represents our qualification for a ‘bad fit’ by eye. Third, any weather-independent model (which is not well characterized by an R^2 value) must be evaluated by eye. Any weather-independent data with outliers or drifting trends in the pre-retrofit year is disqualified, since that period is not representative of the usage, and cannot be used to calculate savings.

9.3 Calendarization

The raw meter data is ‘calendarized’ to align with the first day of each month, by prorating utility bills which straddle calendar months. This practice makes data storage and viewing much easier across the WegoWise application. However, this convenience could possibly undermine the accuracy of a PRISM analysis, since the temperatures will not align perfectly with the usage. The calculated degree days for swing months, which have days on either side of the balance point, could deviate from the true value, especially when the billing date is halfway through a month.

To allay the fears of calendarization, we ran many model-fitting simulations for different billing cycles, using real weather data and realistic building parameters. The recovered model parameters differed from the ‘true’ input parameters by less than 10%, even as the phase between the billing cycle and first of the month increased to 15 days. More importantly, the residuals between the model and the monthly values stayed extremely tight, below 0.1%, while the *absolute value* of the residuals remained below 3%⁷. The predicted model usage is thus more reliable than the underlying parameters. The algorithm’s ability to faithfully model the usage, even when mis-aligned to the billing cycle, convinced us that calendarization does not weaken this analysis.

⁷The absolute residuals are around 1% for billing cycles which align with the calendar. Taken together, these two statistics for the residuals show that, as the phase between the billing cycle and the calendar grows, the model still threads through the middle of the actual usage on average (very low total residuals), but that it ‘tilts’ slightly around that average (low but growing absolute residuals).