Yale SCHOOL OF FORESTRY & ENVIRONMENTAL STUDIES

Yale Carbon Charge Pilot: A Statistical Analysis

Fall 2016

Luke Elder, MEM 2018
Brenda Meany, MEM 2018

Table of Contents

Key Conclusions

- **Pilot Participation**
  Buildings that participated in the Carbon Charge pilot had a statistically significant reduction in their emissions from their baseline relative to the buildings that did not participate in the pilot.

- **Treatment Type**
  All treatment groups that participated in the Carbon Charge (Treatment 1- Treatment 4) experienced a reduction in emissions relative to the baseline, although we only found this reduction to be statistically significant in Treatment 4.

- **Building Size**
  It may be more difficult for larger buildings to reduce their emissions and, conversely, it may be easier for energy intensive buildings to reduce their emissions.

- **Month Effect**
  The magnitude of change in emissions for all buildings was largest in December. However, the winter during the pilot (academic year 2015-2016) was warmer on average than the five winters during the baseline period.

1. Introduction

Climate change is one of the greatest challenges of our time, as its threats are imminent and its impacts are significant. The most recent report from the Intergovernmental Panel on Climate Change (IPCC) brings to light the dangerous effects that climate change is already having on water resources, built environments, agriculture, human health, and natural ecosystems. If countries, states, regions, companies and institutions do not act to keep warming below two degrees Celsius, these dangers will only be exacerbated. In order to keep climate change in check, the IPCC recommends that the world reach zero net emissions before the end of this century (Cubasch et al. 2013). One strategy to reaching zero net emissions, in addition to carbon offsetting, is to internalize the cost of carbon through various market-based solutions.

Fortunately, there has been increased momentum in recent years to create such a solution to climate change in the public and private sectors in the form of carbon pricing. In December 2015, Yale University conducted a carbon charge pilot on twenty building to test the efficacy and feasibility of carbon pricing on Yale's campus. The university's pilot found local illustrations of many issues observed in global efforts to price carbon. The study provided a framework for others to develop carbon pricing models within an educational institution and more broadly across other sectors aimed at achieving a low-carbon economy.

We will be analyzing the effectiveness of the carbon charge pilot conduct from December 2015 to May 2016. The twenty buildings that participated in the carbon charge were divided into 4 pilot groups, with five buildings per pilot group. 258 other Yale campus buildings served as the control group. The control group did not participate in any pricing scheme and did not receive additional information or engagement. The four pricing schemes were:

1. Information: These five units received a building energy report along with information about the \$40 per mtCO_2e, without financial consequences.

2. Target: These five units had a reduction goal of 1% below their baseline. They were charged or received a rebate of \$40 per mtCO_2e for emissions above or below their baseline, respectively.

3. Redistribution: These five units were subject to a revenue-neutral scheme in which each unit was compared to the group's overall percent change in emissions from the group's baseline. These units incurred a charge or a rebate based on emissions relative to the baseline value.

4. Investment: These five units received a subsidy equal to 20% of their baseline carbon charge for investment on energy efficiency. This scheme was an attempt at simulating the second year of the program during which those units that performed well would spend a portion of their carbon charge revenue on building efficiency, education and outreach, and energy conservation initiatives.

For six months, each of the twenty pilot units received a monthly utility report which included energy use and cost information. Yale applied a $40 per metric ton of carbon dioxide equivalent ($mtCO_2e$) carbon price based on the Federal Government's estimate of the Social Cost of Carbon (SCC) (Laemel and Milikowsky 2016). The pilot included all direct greenhouse gas (GHG) emissions (Scope 1) and indirect GHG emissions from the generation of purchased energy (Scope 2). The buildings' representatives received a bill for carbon on their monthly energy report, but they did not incur the charges until the end of the pilot. The buildings were not allowed to purchase carbon offsets to reduce emissions at a carbon price lower than the $40 cost set by the Carbon Charge Task Force.

2. Design and Primary Questions

Primary Questions:
The focus of the 2015-2016 pilot was to gather qualitative data to inform the design of Yale's ultimate carbon pricing scheme and to use the university as a "living laboratory" to figure out the most effective scheme design. The pilot structure and analysis did not prioritize statistical significance. However, in order to generate insights about the administrative feasibility, effectiveness, and impact of Yale's carbon charge, carbon charge design and analysis is critical.

After the pilot's completion, the Carbon Charge Program staff analyzed the data collected and published a white paper detailing the preliminary results. The introduction to the "Data Collection and Analysis" section of this report presented a call to action: "Yale looks forward to updating this document with a more rigorous statistical analysis of the pilot results after consulting with student and faculty experts" (Laemel and Milikowsky 2016, p. 30). This project is an effort to answer this call.

This project seeks to answer three primary questions related to the effect of the Carbon Charge Program at Yale University:

1. What was the relative impact of each variable on the change in emissions between baseline emissions and pilot emissions?

2. With what degree of confidence did instituting a carbon charge, no matter the design, influence the emissions of participating Yale buildings?

3. Did one carbon pricing scheme have a greater impact on change in emissions than another? In other words, which was the most successful carbon pricing scheme at reducing emissions?

Our analysis on the three primary questions is structured as follows:

1. What was the relative impact of each variable on the change in emissions between baseline emissions and pilot emissions?

a. Test: Multiple Regression
b. Schemes: All treatment groups (1-5)
c. Variables:
    i. Log of difference in monthly energy use (continuous);
    ii. Log of baseline energy use (continuous);
    iii. Log of gross square footage (continuous);
    iv. Calendar month (categorical); and
    v. Carbon charge treatment type (categorical)
d. Design of Experiment: We chose to run a Multiple Regression to determine the relative impact of each factor on the per month difference in emissions between the baseline and the pilot. We used December as our reference month and Treatment 5 (control) as our reference treatment.

2. With what degree of confidence did instituting a carbon charge, no matter the design, influence the emissions of participating Yale buildings?
    a. Test: Two-Sample T-Test
    b. Schemes used: Treatment 1-4 vs. Treatment 5
    c. Variables:
        i. Log of difference in monthly energy use (continuous); and
        ii. Participated in the charge (yes/no) (binary-categorical)
    d. Design of Experiment: We chose to run a Two Sample T-Test to determine the degree of confidence instituting a carbon charge influenced the change in emissions of participating Yale buildings. This test will show us, within a specific confidence interval, the mean change in emissions for all the buildings who participated in the pilot and those that did not, as well as the estimate for the difference in these means.

In order to validate the results from a Two-Sample T-Test, we will run a multiple regression that will include compounding effects of additional variables:

a. Test: Multiple Regression
b. Schemes used: Treatment 1-4 vs. Treatment 5
c. Variables:
    i. Log of difference in monthly energy use (continuous);
    ii. Log of baseline energy use (continuous);
    iii. Log of gross square footage (continuous);
    iv. Calendar month (categorical); and
    v. Participated in the charge (yes/no)
d. Design of Experiment: In addition to the Two-Sample T-Test, we chose to run a multiple regression to answer question 2 so that our analysis would account for the compounding effect of all other variables. Our conclusion will determine if simply participating in the carbon charge pilot had an effect on emissions. We used December (12) as our reference month.

3. Did one carbon pricing scheme have a greater impact on percent change of emissions than another? In other words, which was the most successful carbon pricing scheme at reducing emissions?
    a. Test: Multiple Regression
    b. Schemes used: Treatment 1 vs. Treatment 2 vs. Treatment 3 vs. Treatment 4
    c. Variables:
        i. Log of difference in monthly energy use (continuous);
        ii. Log of baseline energy use (continuous);
        iii. Log of gross square footage (continuous);
        iv. Calendar month (categorical); and
        v. Carbon charge treatment type (categorical)
    d. Design of Experiment: We chose to run a Multiple Regression to determine if one carbon pricing scheme had a greater impact on the absolute change of emissions than another scheme. Our conclusion will determine if there was one carbon pricing scheme that was more effective than the others (if at all). We used December (12) as our reference month. We decided to run a multiple regression rather than an ANOVA in order to properly account for the effect of all variables.

3. Data

This project analyzed data that had previously been collected by Carbon Charge Program staff. The dataset contained information on:

- Building name;
- Facility ID;
- Calendar month;
- Total baseline emissions;
- Average baseline emissions;
- Emissions in FY16;
- Percent change in emissions between the baseline and pilot;
- Absolute change in emissions between the baseline and pilot;
- Treatment group;
- Gross square footage;
- First year energy data was collected for the unit; and
- Whether or not the unit participated in the pilot.

This dataset went through a complex process of compilation on behalf of the Carbon Charge Program staff. At Yale, end-users have no control over their energy source which means that the Carbon Charge had to help to rectify the balance among buildings with different energy sources. Yale used a campus- wide electric conversion factor that was developed to calculate buildings' carbon emissions and account for the differences between buildings served by Yale natural gas power plants and local utilities. These emission factors were then applied to monthly energy use data from the three previous fiscal years 2013-2015 as a baseline to determine change in emissions.

Twenty buildings were chosen for the pilot from approximately 270 Yale buildings based on their carbon footprint, building and budget types, and data quality. These twenty buildings were stratified into four treatment types and the remaining 250 buildings were placed into a control group (Treatment 5). An analysis was conducted of quarterly metering data from last three years to ensure that the data of the chosen buildings were of high quality. Throughout the six-month pilot, energy data was collected by meters in each of the buildings and usage was reported on monthly energy reports. Energy use on the monthly bill was converted to tons of carbon equivalents, or tC02e, using the aforementioned campus-wide electric conversion factor.

We worked with Casey Pickett, the Director of the Yale Carbon Charge, to get access to all the data needed for this project. For each building, we only used energy data corresponding with the six months of the pilot (December to May). Additionally, our project employed an updated baseline of a five-year average for each month during the pilot from 2011 to 2015. This revised baseline will be used by the Yale Carbon Charge for all analysis in the future. After cleaning the original data, our final data set included a total of 1,490 data points.

<u>Deleted Data</u>:

- Energy data for months 6-11 (calendar year). These data were outside the pilot timeframe (months 12-5) and thus not relevant to this analysis.
- Summary energy data (months 7-6). We removed these summary data and analyzed monthly data instead.
- Total Baseline 2011-2015 with values of zero. There were 56 data points that had zero for the total baseline 2011-2015, which were removed in order to run the regression.
- Emissions in FY16 with values of zero. 9-10 data points had zero for emissions in FY16, which could be due to energy meter failure. Because this data is missing and therefore incorrect, it was removed before our analysis.
- Gross square footage with values of zero. Energy data associated with buildings that are now owned by Yale were assigned a gross square footage of zero. We deleted these buildings from our dataset as they are outside the scope of the Carbon Charge project.

<u>Transformations</u>:

- Average Baseline Emissions. We used a natural log transformation on average baseline emissions data.
- Gross Square Footage. We used a natural log transformation on GSF data.
- Difference of Emissions.
    - The data for percent change of emissions between the pilot and the baseline is exponentially distributed and there is evidence of heteroscedasticity in this distribution: As shown in the area graph below labeled "PercentChange," there are many small percent changes and a few very large percent changes in emissions. In order to account for this and for the fact that we see large percent changes in buildings that are relatively small, we used the absolute difference in emissions to normalize our data and best capture emissions reductions as seen in the area graph below labeled "logchange." To further explain, as gross square

footage increases, baseline emissions increase as well and thus any absolute change in emissions will be small as a percent change.

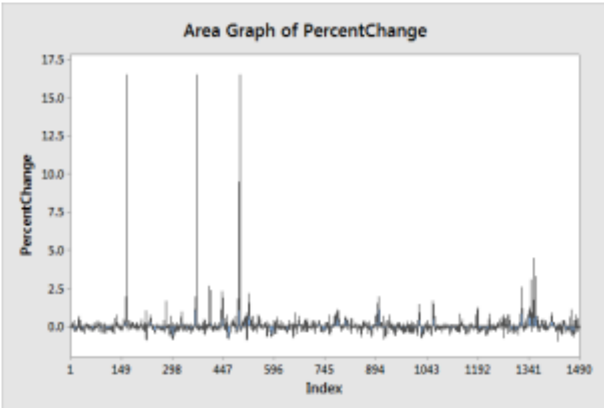

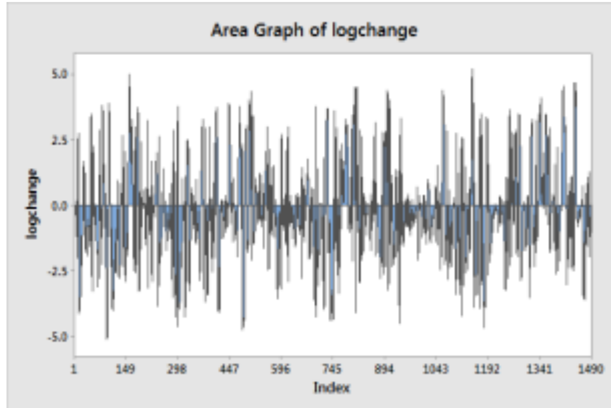

- We used a complex transformation on our response variable described by the following equation: *sign(Emissions in FY16-Average Baseline)*ln(abs(Emissions in FY16-Average Baseline)+1)*. This transformation took the natural log of the absolute value of the difference in pilot emissions and baseline emissions for each month. A one was added to the absolute value of this difference before taking the natural log because the natural log of zero is undefined. The addition of one still preserved the data where the difference in emissions was zero because the natural log of one is zero. The first half of the equation was included to preserve the sign of the difference, with negative values showing a decrease in emissions during the pilot.

Other Changes:

- There was no emissions data for months 7 and 8 for one unit, Temple St 40. This could be due to a broken energy meter or collection error. This building was not included in any of the four treatment groups in the Carbon Charge pilot, but it was included in the control group. To address this missing data, we used the average energy use from that unit for month 7 and month 8.

4. Descriptive Plots and Observations

a. Descriptive Statistics: logchange, logGSF, logbaseline, Treatment, Month (calendar)

Variable	N	N*	Mean	SE Mean	StDev	Minimum	Q1	Median	Q3
logchange	1490	0	-0.2390	0.0518	1.9988	-5.1530	-1.5535	-0.3725	0.8253
logGSF	1490	0	10.151	0.0319	1.231	6.454	9.223	10.041	11.088
logbaseline	1490	0	3.0593	0.0364	1.4034	0.7504	1.8982	2.8683	4.1512
Treatment	1490	0	4.7987	0.0195	0.7509	1.0000	5.0000	5.0000	5.0000
Month (calendar)	1490	0	4.5074	0.0932	3.5962	1.0000	2.0000	4.0000	5.0000

```
Variable           Maximum
logchange           5.2070
logGSF             13.192
logbaseline         6.3885
Treatment           5.0000
Month (calendar)   12.0000
```

b. Histograms:

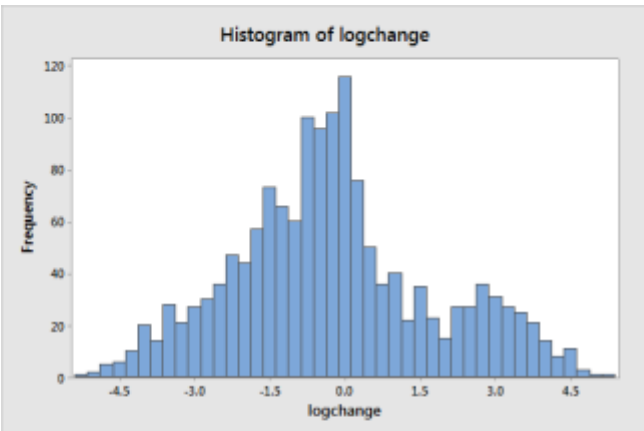

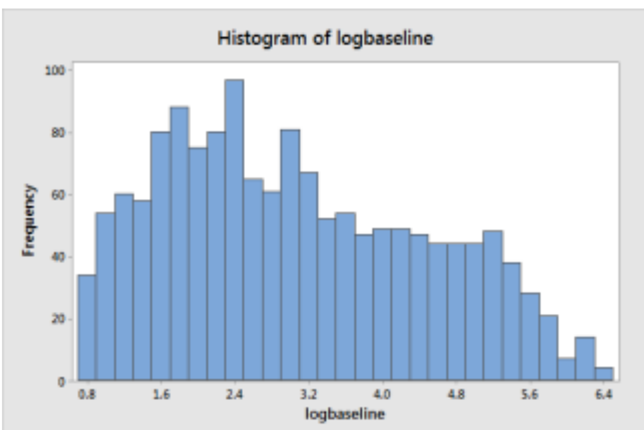

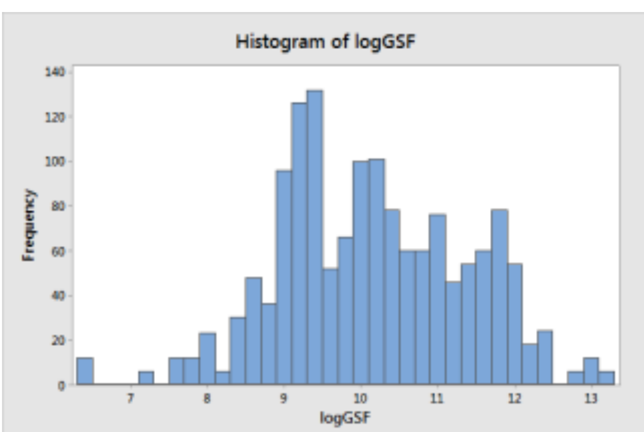

c. Normal Quantile Plot (probability plot of logchange):

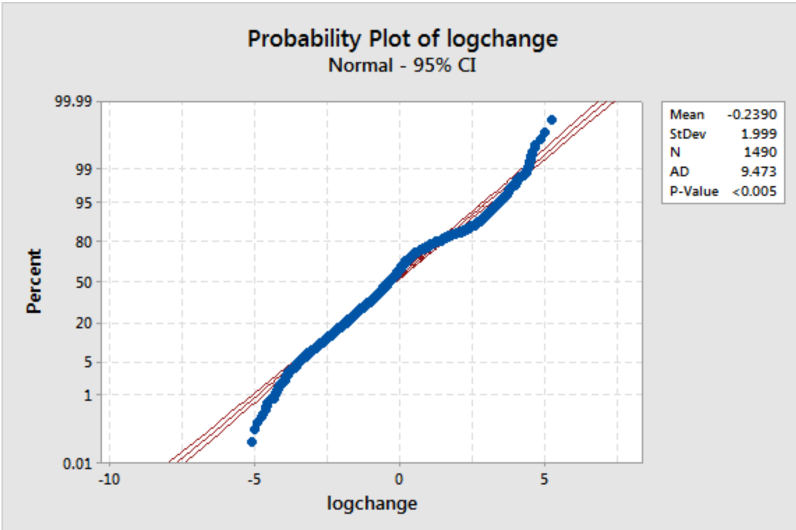

Based on a 95 percent confidence interval, the normal distribution appears to fit the sample data fairly well. The data is skewed from the probability line at the end in both directions, which is due to extreme data points that are evidence of the exponential relationship between baseline emissions and gross square footage.

d. Matrix Plot:

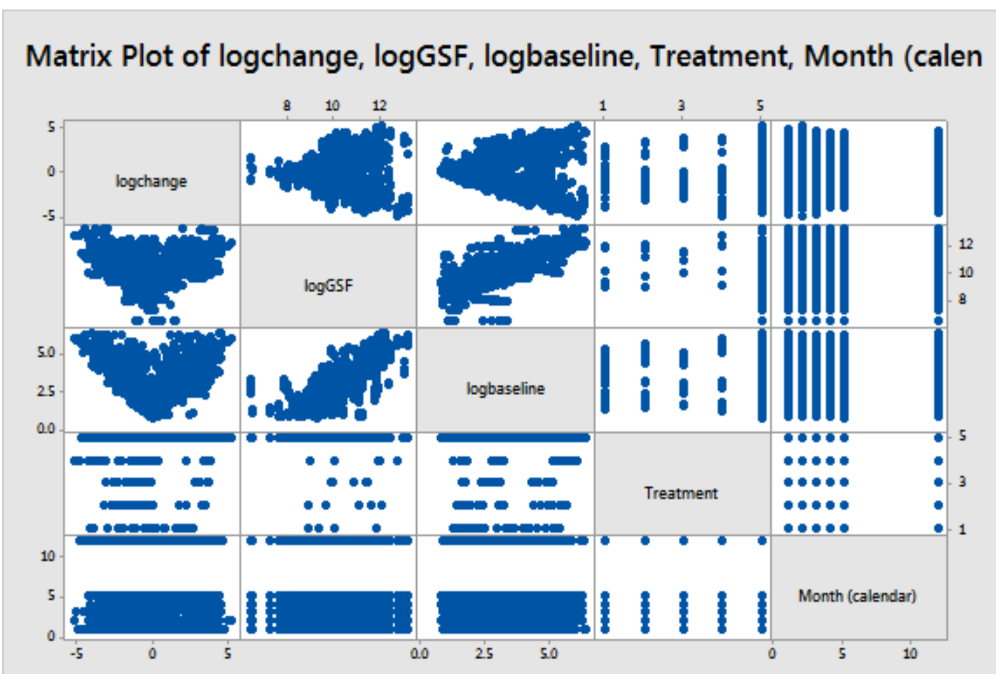

In this matrix plot, there is evidence of multicollinearity between logbaseline and logGSF (which we examined individually in the scatterplot below). There is a smaller spread of log change, logGSF, and logbaseline among treatment groups 1-4 as opposed to treatment 5. This is due to the difference in sample size. There appears to be a relatively even spread of different building sizes (logGSF) and energy intensity (logbaseline) across treatment groups 1-4.

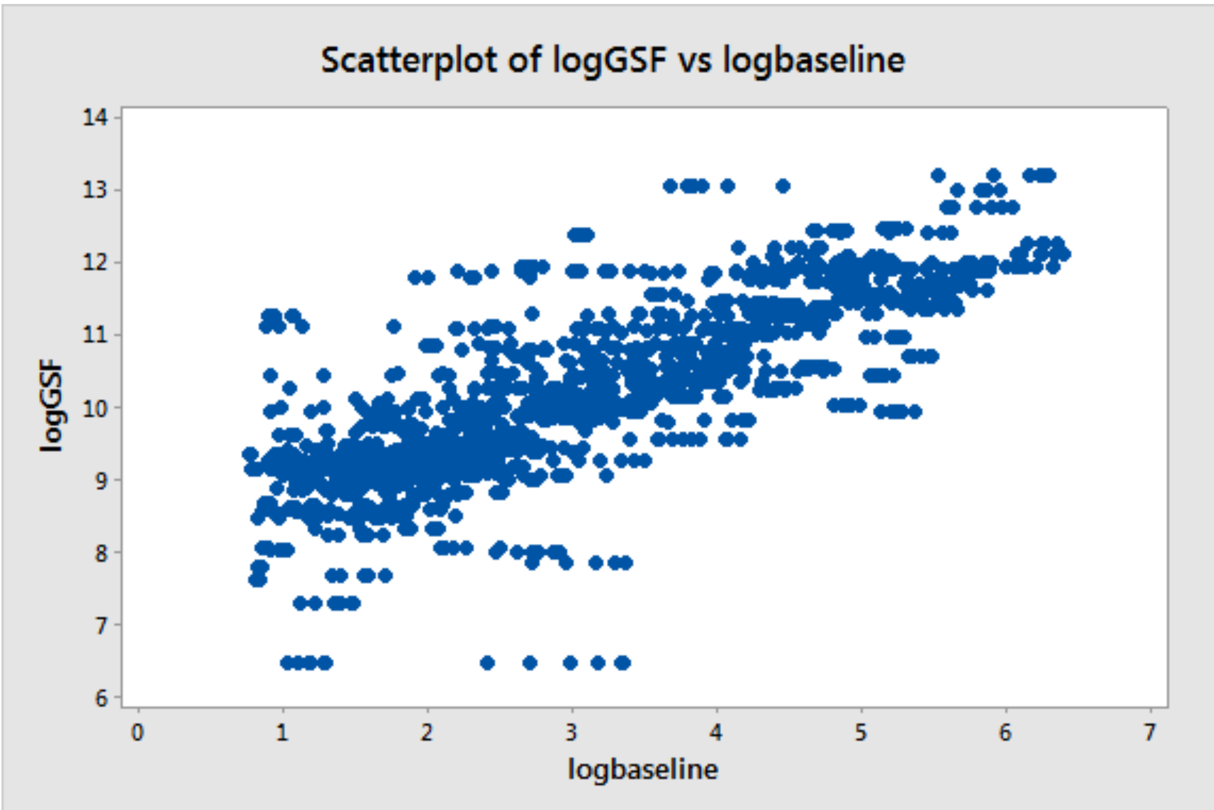

The scatterplot of log GSF (gross square footage) and log baseline shows a clear correlation between the two variables. If baseline emissions are high, then the building unit was already using more energy before the pilot and thus they have further to go in terms of absolute change in emissions. This plot shows that baseline energy use is actually a proxy for building size (gross square footage) as these two variables are multicollinear.

After the various transformations and data cleaning described previously, we concluded that the assumptions of this model had been sufficiently met, and that there were no other transformations needed.

5. Analysis

<u>Question 1</u>: Which variables had the greatest impact on the per month change in emissions between the baseline and the pilot?

<u>Test Used</u>: *Multiple Regression*

<u>Results</u>:

```
Regression Analysis: logchange versus logGSF, logbaseline, Treatment, Month (calendar)

Method

Categorical predictor coding  (1, 0)


Analysis of Variance

Source               DF    Adj SS   Adj MS  F-Value  P-Value
Regression           11    333.59   30.326     7.98    0.000
  logGSF              1     14.29   14.288     3.76    0.053
  logbaseline         1     20.33   20.333     5.35    0.021
  Treatment           4     58.50   14.625     3.85    0.004
  Month (calendar)    5    223.96   44.792    11.79    0.000
Error              1478   5615.26    3.799
Total              1489   5948.85

Model Summary

      S    R-sq  R-sq(adj)  R-sq(pred)
1.94916   5.61%      4.91%       3.96%


Coefficients

Term                 Coef  SE Coef  T-Value  P-Value   VIF
Constant           -1.648    0.568    -2.90    0.004
logGSF             0.1326   0.0684     1.94    0.053  2.78
logbaseline       -0.1393   0.0602    -2.31    0.021  2.80
Treatment
  1                -0.424    0.360    -1.18    0.240  1.00
  2                -0.512    0.361    -1.42    0.156  1.01
  3                -0.284    0.361    -0.79    0.431  1.01
  4                -1.263    0.363    -3.48    0.001  1.02
Month (calendar)
  1                 0.263    0.175     1.50    0.134  1.66
  2                 0.472    0.175     2.70    0.007  1.66
  3                 0.407    0.175     2.33    0.020  1.67
  4                 1.142    0.175     6.52    0.000  1.68
  5                 0.953    0.176     5.41    0.000  1.69

Regression Equation

           Month
Treatment  (calendar)
```

1	1	logchange = -1.809 + 0.1326 logGSF - 0.1393 logbaseline
1	12	logchange = -2.071 + 0.1326 logGSF - 0.1393 logbaseline
1	2	logchange = -1.599 + 0.1326 logGSF - 0.1393 logbaseline
1	3	logchange = -1.665 + 0.1326 logGSF - 0.1393 logbaseline
1	4	logchange = -0.929 + 0.1326 logGSF - 0.1393 logbaseline
1	5	logchange = -1.119 + 0.1326 logGSF - 0.1393 logbaseline
2	1	logchange = -1.897 + 0.1326 logGSF - 0.1393 logbaseline
2	12	logchange = -2.160 + 0.1326 logGSF - 0.1393 logbaseline
2	2	logchange = -1.688 + 0.1326 logGSF - 0.1393 logbaseline
2	3	logchange = -1.753 + 0.1326 logGSF - 0.1393 logbaseline
2	4	logchange = -1.018 + 0.1326 logGSF - 0.1393 logbaseline
2	5	logchange = -1.207 + 0.1326 logGSF - 0.1393 logbaseline
3	1	logchange = -1.669 + 0.1326 logGSF - 0.1393 logbaseline
3	12	logchange = -1.932 + 0.1326 logGSF - 0.1393 logbaseline
3	2	logchange = -1.460 + 0.1326 logGSF - 0.1393 logbaseline
3	3	logchange = -1.525 + 0.1326 logGSF - 0.1393 logbaseline
3	4	logchange = -0.790 + 0.1326 logGSF - 0.1393 logbaseline
3	5	logchange = -0.979 + 0.1326 logGSF - 0.1393 logbaseline
4	1	logchange = -2.648 + 0.1326 logGSF - 0.1393 logbaseline
4	12	logchange = -2.911 + 0.1326 logGSF - 0.1393 logbaseline
4	2	logchange = -2.439 + 0.1326 logGSF - 0.1393 logbaseline
4	3	logchange = -2.504 + 0.1326 logGSF - 0.1393 logbaseline
4	4	logchange = -1.769 + 0.1326 logGSF - 0.1393 logbaseline
4	5	logchange = -1.958 + 0.1326 logGSF - 0.1393 logbaseline
5	1	logchange = -1.385 + 0.1326 logGSF - 0.1393 logbaseline
5	12	logchange = -1.648 + 0.1326 logGSF - 0.1393 logbaseline
5	2	logchange = -1.176 + 0.1326 logGSF - 0.1393 logbaseline

5	3	logchange = -1.241 + 0.1326 logGSF - 0.1393 logbaseline
5	4	logchange = -0.506 + 0.1326 logGSF - 0.1393 logbaseline
5	5	logchange = -0.695 + 0.1326 logGSF - 0.1393 logbaseline

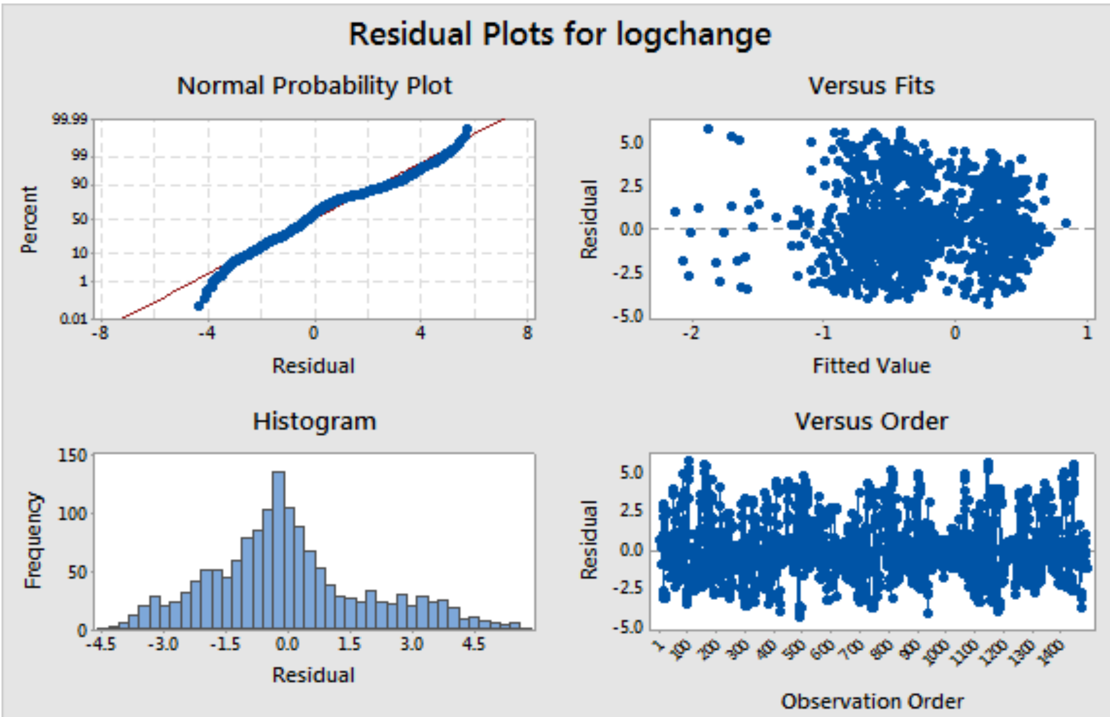

Our normal probability plot shows our data is relatively normally distributed, although the points are skewing off the probability line at the end in both directions, which is evidence of the exponential distribution within the data. Our histogram appears to be normally distributed, although it is slightly skewed to the right. Our versus fit plot shows evidence of heteroskedasticity, which is what we expected to see based on our matrix plot. Our "versus order" plot appears to be normally distributed. We concluded that the assumptions of this model had been sufficiently met according to the residual plots.

Question 2: With what degree of confidence did instituting a carbon charge, no matter the design, influence carbon emissions of affected Yale buildings?

Test #1 used: *Two-Sample T-Test.* The units that were in the control group were coded with a value of zero and the 20 units that did participate were coded with a value of one.

Results:

```
Two-Sample T-Test and CI: logchange, Participated in Charge

Two-sample T for logchange

Participated
in Charge        N   Mean  StDev  SE Mean
0             1370  -0.19   1.98    0.053
1              120  -0.80   2.13     0.19

Difference = μ (0) - μ (1)
Estimate for difference:  0.611
95% CI for difference:  (0.212, 1.010)
T-Test of difference = 0 (vs ≠): T-Value = 3.03  P-Value = 0.003  DF = 137
```

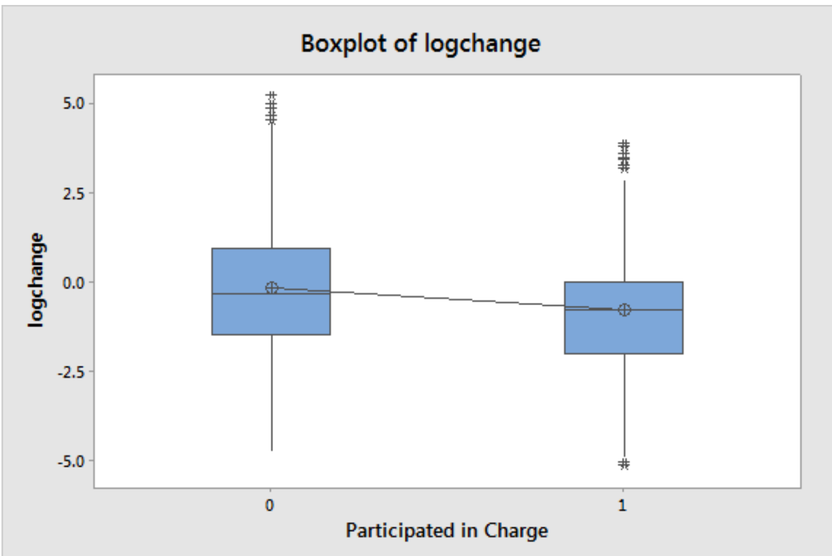

Test #2 used: *Multiple Regression.* Units that participated in the pilot program and Month 12 (December) were used as the reference groups for the test.

Results:

```
Regression Analysis: logchange versus logGSF, logbaseline, Participated in , Month
(calendar)

Method

Categorical predictor coding  (1, 0)


Analysis of Variance

Source                      DF   Adj SS   Adj MS  F-Value  P-Value
Regression                   8   316.43   39.553    10.40    0.000
  logGSF                     1    15.13   15.132     3.98    0.046
```

logbaseline	1	23.18	23.178	6.09	0.014
Participated in Charge	1	41.34	41.339	10.87	0.001
Month (calendar)	5	222.62	44.525	11.71	0.000
Error	1481	5632.42	3.803		
Total	1489	5948.85			

Model Summary

S	R-sq	R-sq(adj)	R-sq(pred)
1.95016	5.32%	4.81%	4.10%

Coefficients

Term	Coef	SE Coef	T-Value	P-Value	VIF
Constant	-1.658	0.568	-2.92	0.004	
logGSF	0.1363	0.0683	1.99	0.046	2.77
logbaseline	-0.1484	0.0601	-2.47	0.014	2.79
Participated in Charge					
1	-0.618	0.188	-3.30	0.001	1.02
Month (calendar)					
1	0.264	0.175	1.50	0.133	1.66
2	0.473	0.175	2.70	0.007	1.66
3	0.406	0.175	2.33	0.020	1.67
4	1.140	0.175	6.50	0.000	1.68
5	0.949	0.176	5.38	0.000	1.69

Regression Equation

Participated in Charge	Month (calendar)	
0	1	logchange = -1.394 + 0.1363 logGSF - 0.1484 logbaseline
0	2	logchange = -1.185 + 0.1363 logGSF - 0.1484 logbaseline
0	3	logchange = -1.252 + 0.1363 logGSF - 0.1484 logbaseline
0	4	logchange = -0.518 + 0.1363 logGSF - 0.1484 logbaseline
0	5	logchange = -0.708 + 0.1363 logGSF - 0.1484 logbaseline
0	12	logchange = -1.658 + 0.1363 logGSF - 0.1484 logbaseline
1	1	logchange = -2.013 + 0.1363 logGSF - 0.1484 logbaseline
1	2	logchange = -1.804 + 0.1363 logGSF - 0.1484 logbaseline
1	3	logchange = -1.870 + 0.1363 logGSF - 0.1484 logbaseline
1	4	logchange = -1.136 + 0.1363 logGSF - 0.1484 logbaseline
1	5	logchange = -1.327 + 0.1363 logGSF - 0.1484 logbaseline
1	12	logchange = -2.276 + 0.1363 logGSF - 0.1484 logbaseline

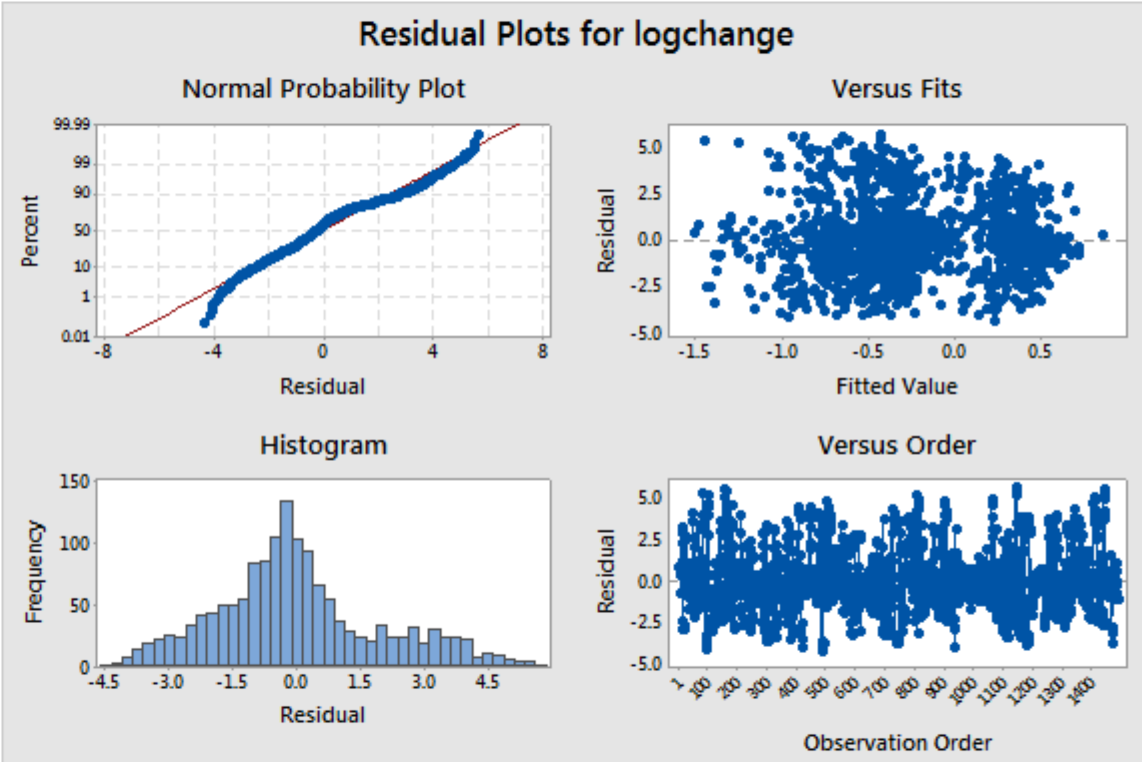

<u>Question 3</u>: Did one treatment type have a greater impact on carbon emissions than another? In other words, which was the most successful carbon pricing scheme at reducing Yale building emissions?

<u>Test used</u>: *Multiple Regression.* Treatment 5 (control group) and Month 12 (December) were used as the reference groups for the test.

<u>Results</u>:

```
Coefficients

Term                  Coef  SE Coef  T-Value  P-Value   VIF
Constant            -1.648    0.568    -2.90    0.004
logGSF              0.1326   0.0684     1.94    0.053  2.78
logbaseline        -0.1393   0.0602    -2.31    0.021  2.80
Treatment
  1                 -0.424    0.360    -1.18    0.240  1.00
  2                 -0.512    0.361    -1.42    0.156  1.01
  3                 -0.284    0.361    -0.79    0.431  1.01
  4                 -1.263    0.363    -3.48    0.001  1.02
Month (calendar)
  1                  0.263    0.175     1.50    0.134  1.66
  2                  0.472    0.175     2.70    0.007  1.66
  3                  0.407    0.175     2.33    0.020  1.67
  4                  1.142    0.175     6.52    0.000  1.68
  5                  0.953    0.176     5.41    0.000  1.69
```

6. Conclusions and Discussion

<u>Question 1</u>: What was the relative impact of each variable on the change in emissions between baseline emissions and pilot emissions?

There are a few general conclusions we can draw from the results of our multiple regression model:

1. It may be more difficult for larger buildings to reduce their emissions and, conversely, it may be easier for energy intensive buildings to reduce their emissions. The coefficient for logGSF is 0.1326, which indicates a positive relationship between gross square footage and emissions during the pilot: for all units, if the gross square footage was larger, emissions during the pilot were larger. In other words, bigger buildings had a harder time reducing their emissions than smaller buildings. This result indicates that buildings that are poor energy stewards had an easier time becoming better energy stewards and reducing their emissions.

   Additionally, the coefficient for logbaseline is -0.1393, which indicates an inverse relationship between baseline emissions and emission during the pilot. In other words, units that are more energy intensive were able to reduce their emissions by more than units that had a smaller baseline energy use.

   These results combined indicate that larger buildings on Yale's campus that had relatively low baseline emissions had the hardest time reducing their emissions during the pilot.

2. All treatment groups that participated in the Carbon Charge (Treatment 1- Treatment 4) experienced a reduction in emissions relative to the baseline. According to our Two-Sample T-Test, the coefficients for each treatment group are all negative (-0.424 for Treatment 1, -0.512 for Treatment 2, -0.284 for Treatment 3, -1.263 for Treatment 4). However, the coefficient for the constant is also negative (-1.648) (the constant refers to Treatment 5 (control group) and month 12 (December), meaning that the emissions for both the treatments that participated in the program and the control group went down.

   According to our multiple regression, the coefficient for buildings that participating in the charge was -0.618, meaning that change in emissions for all buildings that participated in the pilot decreased relative to the buildings that did not participate in the pilot ($p = 0.001$).

3. The winter during academic year 2015-2016 probably was warmer on average than the previous five winters. As stated above, the coefficient for the constant in our model is -1.648, meaning that the emissions in the control group went down during the pilot relative to their baseline emissions. The descriptive statistics for our data show that the median value of logchange (the log of the difference between emissions during the pilot

and baseline emissions) is negative (-0.3725). This means that more than half the changes were negative, indicating that the winter was probably warmer overall and thus all units used less energy for heating.

4. The magnitude of change in emissions was largest in December. This result could lead us to believe that it is easier to reduce your emissions during cold months, perhaps by heating slightly less. While this makes logical sense assuming that heating is one of the largest components of energy use, in this report we are not focusing on the effect of one individual month during the six-month pilot. However, this result is informative for individual units that are designing an energy reduction strategy to focus their efforts on the coldest months.

<u>Question 2</u>: With what degree of confidence did instituting a carbon charge, no matter the design, influence carbon emissions of affected Yale buildings?

Based on the results of our Two-Sample T Test, we can conclude that Carbon Charge participation did have a statistically significant influence on the emissions of Yale buildings ($p<0.05$). Our test showed a statistically significant difference of 0.611 between the log of the absolute change in emissions based on whether the building participated in the carbon charge or not. The mean log change in emissions for the control group was -0.19, while the mean log change in emissions for the 20 units that participated in the program was -0.8. The 95% confidence interval for this mean difference is (0.212, 1.01). It is somewhat surprising that we found a statistically significant difference between the units that participated in the pilot and those that did not, if only given the difference in the sample size of each group (120 data points for the units that participated versus 1,370 data points for the units that did not). These results are positive for the Carbon Charge Program at Yale, justifying their efforts at instituting a price on carbon in order to decrease campus emissions. We suggest that Yale adopt a campus wide Carbon Charge in order to achieve the university goal of carbon neutrality by 2050.

<u>Question 3</u>: Did one treatment type have a greater impact on carbon emissions than another? In other words, which was the most successful carbon pricing scheme at reducing Yale building emissions?

We ran regression to determine if one treatment type had a greater impact on carbon emissions than another group and based on these results we can conclude that Treatment 4 (Investment) was the most statistically significant in reducing emissions relative to the control group, Treatment 5. While all treatments groups decreased emissions relative to the control group, the units that were placed in Treatment 4 had the greatest change in carbon emissions. The coefficient for Treatment 4 in our regression is -1.263 which was more than twice as much as the emissions of the other treatment groups (-0.424 for Treatment 1, -0.512 for Treatment 2, and -0.284 for Treatment 3). The P-value for Treatment 4 is .001, relative to the control group, which is also the only statistically significant P-value for all treatment groups (0.240 for Treatment 1, 0.156 for Treatment 2, 0.431 for Treatment 3).

Our results conclude that investment in spending on self-guided energy actions did lead to the greatest emission reduction during the pilot. However, this result is based on a limited sample size, and therefore more research is needed before definitively deciding on one carbon charge scheme over another using this statistical model.

7. Points for Further Analysis

We have identified two variables that would be beneficial to Yale for future Carbon Charge analysis:

1. Weather. It would be beneficial to know the effect weather variability has on energy use data, if any. Knowing this information could help better relate energy use data and Carbon Charge Program design.
2. Building envelope. It would be beneficial to know how insulated each building is in order to analyze the effect of weather on energy use. This is done by conducting a blower test to assess how "tight" the building is: the tighter the building, the more energy efficient it is. This would also help address the difference in insulation in buildings or how many years they have gone without being renovated, because their climate envelope score could be directly compared between buildings.

8. References

Cubasch, U., D. Wuebbles, D. Chen, M.C. Facchini, D. Frame, N. Mahowald, and J.-G. Winther, 2013: Introduction. In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Laemel, R. and J. Milikowsky. "Yale University's Carbon Charge: Preliminary Results from Learning by Doing." Yale Carbon Charge Program, October 10, 2016, URL: http://carbon.yale.edu/sites/default/files/files/Carbon_Charge_Pilot_Report_20161010.pdf