



Analysis Plan

Project Name: Decreasing energy costs in Federally assisted housing

Project Code: 1742

Date Finalized: February 12, 2018

This document serves as a basis for distinguishing between planned (confirmatory) analysis and any unplanned (exploratory) analysis that might be conducted on project data. This is crucial to ensuring that results of statistical tests will be properly interpreted and reported. In order that the Analysis Plan fulfill this purpose, it is essential that it be finalized and date-stamped before we begin looking at the data — ideally, before we take possession of the data. Once this plan is finalized, a date is entered above, and the document is posted publicly on our team website.

Outcome variables to be analyzed:

NYCHA measures daily kwh of electricity usage at the apartment unit level. Our primary outcome of interest is average daily kwh of electricity usage during a two-week follow up period post intervention. The intervention began with flyers mailed from Washington, DC to the treatment units in New York, NY on September 12, 2017. The follow up period will run for 14 days from two days after the intervention letters were sent (September 14, 2017). To capture the same days of the week, the baseline period will end 6 days before the intervention letters were sent and cover the previous 14 days. Thus, the baseline period will include the time period from August 24, 2017 to September 6, 2017, and the follow up period will include September 14, 2017 to September 27, 2017.

Secondary outcomes of interest are average daily kwh of electricity usage on hot or cold days - measured using average outside temperature. We also will plot treatment effects for each intervention treatment (compared to the control group) for each day during the two-week follow up period and examine effects for a longer time period (i.e., one month).

Indices:

We calculate average daily electricity as the mean of daily total electricity use over the two-week follow up period. Daily average usage is calculated as total energy usage (sum of average daily electricity use) during the follow-up period divided by the number of days of energy use recorded.

Total electricity use will be calculated as the sum of daily total electricity use over the same period.

Statistical models:

We will run a multiple regression model with our primary outcome regressed on an indicator for received any letter (multiple-tips letter (n=929) and single-tip letter (n=935)), a control variable

for average daily energy use during the baseline period, and block fixed effects to account for how random assignment occurred. Block fixed effects are based on one of the seven developments and apartment size. Robust standard errors will be clustered by randomization block.

Our primary specification is:

$$Y_{ib} = \beta_0 + \beta_1(MT_{ib} | ST_{ib}) + \beta_2 Z_{ib} + \beta_3 M_b + u_{ib}$$

Where:

- Y_{ib} : Outcome Y for housing unit i, in block b;
- MT_{ib} : Indicator for multiple-tip letter group;
- ST_{ib} : Indicator for single-tip letter group;
- Z_{ib} : Unit-level controls for baseline energy usage;
- M_b : Block fixed effects;
- u_{ib} : Individual unit error term.

The regressor of interest is β_1 , which estimates the causal impact of being in the multiple-tip or single-tip letter conditions compared to the control condition, controlling for unit-level baseline energy usage and randomization block.

Clarification June 6th, 2018

We reported that we would use robust standard errors clustered by randomization block. We realized that we should not cluster at the block level, since treatment was at the apartment unit level. Thus, we instead constructed unclustered heteroskedasticity-robust standard errors in all specifications.

Although our regression specification did not change, we updated the notation to better align with common notation practices. Our primary specification is:

$$Y_{ib} = \beta_0 + \beta_1 MT_{ib} + \beta_2 ST_{ib} + \pi Z'_{ib} + \alpha_b + \varepsilon_{ib}$$

Where:

- Y_{ib} : Outcome Y for housing unit i, in block b;
- MT_{ib} : Indicator for multiple-tip letter group;
- ST_{ib} : Indicator for single-tip letter group;
- Z'_{ib} : A vector of unit-level controls for baseline energy usage¹;
- α_b : Block fixed effects; and
- ε_{ib} : Individual unit error term

¹ Unit-level controls for baseline energy use include mean daily energy use for the baseline period and indicators for imputed values for each day during the baseline period.

Transformations:

Heating-degree and cooling-degree days are calculated using daily average temperature from the closest weather tower to each apartment site. These data come from the [National Oceanic and Atmospheric Administration](#) and weather towers in Central Park.

- Cooling-degrees are calculated as the greater of 0 or mean daily temperature minus 65 degrees.
- Heating-degrees are calculated as the greater of 0 or 65 degrees minus mean daily temperature.

In exploratory analysis, we will use cooling-degree days and heating-degree days (if applicable, given the time period) as additional controls and/or for subgroup analysis. Cooling- (heating-) degree days are calculated as the sum of cooling- (heating-) degrees on each day over the study time period.²

Follow-up analyses:

We will run the follow-up analysis using the same multiple regression framework described above. Our follow-up analysis will compare impacts of the single-tip letter to the multiple-tip letter. These will be conducted regardless of the estimated treatment effect found in the primary analysis to examine whether pooling the treatment arms masks any differential effect on energy use of the two versions of the letter. If we find differences between the intervention groups, we will also compare single-tip letter to control and multiple-tip letter to control.

The pooled test in our primary analysis will likely provide slightly more power than these secondary comparisons; however, we would not expect to find effects for the pooled treatment without also finding treatment effects for the individual intervention arms.

Note that our sample size and related power analysis make finding statistically significant effects between treatment arms unlikely. We hope that this analysis will help us learn about the direction and magnitude of the effects of the two treatment arms (i.e., single-tip and multiple-tip conditions).

Inference criteria:

We will use two-tailed tests in our analysis. We are running one primary regression and test that examines the impact of receiving any letter. Follow up regressions will then test whether the

² The definitions for cooling- and heating-degrees and degree-days are commonly used as indicators of energy requirements for air conditioning or space heating, respectively. See: Monthly Energy Review, January 2018. U.S. Energy Information Administration, Office of Energy Statistics, Washington, DC. Available at <https://www.eia.gov/totalenergy/data/monthly/>, accessed Feb. 5, 2018.

single-tip letter or multiple-tip letter impacts are different from one another. We will use a Tukey HSD that includes a family-wise error rate adjustment to account for multiple hypothesis tests.

Clarification June 6th, 2018

We originally planned to adjust for multiple comparisons only for the pairwise comparisons of the effects of each treatment arm. Since our analysis plan describes conducting four primary statistical tests regardless of the findings of being sent any energy-saving flyer compared to the control condition, we determined that multiple comparison adjustments for all four statistical tests was more appropriate.

Moreover, we chose to adjust for multiple comparisons using the holm-bonferroni family-wise error rate rather than the tukey family-wise error rate, since that adjustment is more appropriate when the multiple comparisons come from different regressions.

Data exclusion:

We will exclude outliers that appear to be the result of malfunctioning sensors. Previous studies defined an outlier as a daily reading above 1500 kwh (Allcott & Rogers, 2014).³

- Prior to running outcome analysis, we will examine the distribution of energy usage for apartments of the same size in the same housing complex to determine the outlier threshold. We will define outliers at the daily energy reading level.
- Households with negative average daily kwh electricity used (i.e., negative electricity meter reads) also will be excluded.

Clarification June 6th, 2018

We reported that we would exclude outliers that appear to be the result of malfunctioning WEMs. From our examination of the data distribution, we classified the outlier threshold as 200 kwh of energy use or above on a given day. We treated these observations as missing data.

From our NYCHA data contact, we also learned that we should consider recordings of zero daily energy to be due to malfunctioning WEMs. Thus, we treated energy use reports of zero as missing data.

³ Allcott, H. and T. Rogers. 2014. The short-run and long-run effects of behavioral interventions: experimental evidence from energy conservation. *American Economic Review*, 104(10): 3003-3037.

Treatment of missing data:

We will impute baseline values for apartment units with missing baseline data or units that meet the data exclusion criteria described above. If the missing data are baseline observations, then we will impute the value and include an imputed data indicator. The value will be imputed as the mean daily use for apartments of the same size in the same building on the same day. However, we do not expect to have missing baseline data based on study sample criteria.

Since we are reporting average daily electricity usage, we will be able to include apartment units with partially missing outcome data in our analysis. We will include units in our analysis if they have valid energy usage observations for 75% or more days (i.e., 11 or more days in our primary two-week specification). If some units in the analysis are observed with missing energy use on up to 3 days, the analysis will use observation weights based on the number of days observed for each unit (ranging between 11 and 14).

We will run sensitivity analysis that bound our estimates. In this analysis, we will impute missing values to the minimum and maximum observed values for apartments of the same size, building, and day.

Clarification June 6th, 2018

We learned that in some cases that we could not impute baseline missing data to the day by unit size by building level, because all data at the day by unit size by building level were missing. In these cases, we imputed values to the day by unit size by development level.

Limitations:

We expect that we are underpowered to detect statistically significant effects of the intervention letters, particularly differential effects by letter type.

Exploratory analysis (optional):

Exploratory analysis will examine the effects of the intervention for each post-intervention day. This analysis will help examine intervention fade out, differential effects by day of the week, and differential effects on hot or cold days. In addition, our follow up analysis will aggregate usage by different categories of post-intervention days (e.g., first week only, second week only, weekdays only, weekends only, hot days only). This exploratory analysis will provide descriptive insight as to whether any detected effects fade out over time or are driven by other day characteristics.

Subgroup analysis will examine whether effects are concentrate in apartments of different sizes or have relatively high or low baseline energy usage. Low-baseline energy users may be more

receptive to strategies to reduce usage, while the letter may be more informative for high-baseline energy users. We may also expect smaller/larger effects based on apartment size (e.g., fewer people to coordinate energy consumption in smaller units, greater potential for change in larger units).

We will also extend the analysis to include a longer follow up period to examine: potential for intervention fade out, potential for recursive habit-forming effects, potential for effects during different types of weather, and account for late mail delivery or opening patterns.

We will cluster by unit for analysis that includes multiple observations per unit (i.e. days).

We also have considered controlling for daily temperature to potentially improve precision and will run models with daily energy use as the outcomes that control for heating- and cooling-degree days for exploratory analysis.

There are a number of potential exploratory outcomes that we could consider depending on available data. These exploratory outcomes could include:

- Standard deviation in daily energy usage, which would provide evidence of how the intervention changed variability in energy usage throughout the day.
- Energy use during daytime hours that are most closely tied to when behavioral changes would occur based on the intervention tips.
- Differential effects based on apartment location (i.e., north/south facing).
- Differential effects based on household composition.