Persistence of Low-Income Weatherization Savings

Andy Lick; Maddie Koolbeck; Scott Pigg, Slipstream, Madison, WI

ABSTRACT

Weatherizing low-income homes was among the first energy efficiency efforts in the U.S. Many studies have documented the first-year impacts of these programs, but far fewer have examined the savings persistence even though the measures that dominate the savings are assumed to provide benefits for many years. Examining the persistence of savings can help establish the extent to which long-lived measures can be counted on to deliver energy savings and inform policy.

Persistence at the program-level, is the examination of whether the savings achieved in the year following weatherization are maintained over time. This paper summarizes an analysis of the persistence of savings for homes treated by the Wisconsin low-income weatherization program for up to 13 years following weatherization. The analysis utilized utility consumption histories for six utilities and program tracking data covering 14 years.

This analysis consisted of both a simple comparison of year-over-year trends and regression modeling that compared results for four model-fitting procedures, including ordinary least-squares, robust regression, quantile regression, and mixed-effects models. The analysis shows that natural gas savings are stable while electricity savings persistence is more ambiguous given the data at hand.

Although the study is limited in geographic scope, the program on which it was based is implemented similarly in many areas. Thus, the methodological approach could be repeated elsewhere, and the general results offer insight for other programs.

Introduction

Weatherizing low-income homes was among the first energy efficiency efforts in the U.S. and has more than a 40-year track record. While many studies have documented the first-year impacts of these programs, far fewer have examined the persistence of their savings over time even though the insulation, air-sealing and mechanical-system measures that dominate the savings are assumed to provide benefits for many years. Examining the persistence of energy savings can help establish the general extent to which long-lived shell and mechanical-system measures can be counted on to deliver savings and inform policy adaptations for the short and long term.

Persistence at the program-level, is the examination of whether the savings achieved in the year following weatherization are maintained over time. We measured savings as the change in weather-normalized annual consumption of energy (NAC). To date, persistence of energy savings from weatherization has rarely been evaluated due to cost and data barriers (Vine 1992; Hoffman et al. 2015). The majority of the studies performed on this subject were completed in the 1990s, which a number of programs utilized to make the assertion that changes in savings over time are rare (Hoffman et al. 2015). A number of these studies are based on short timeframes, or a limited dataset (Violette 2013). For example, the most applicable prior study was conducted for Wisconsin in the early 1990s and was able to examine the persistence of energy savings for 8 years post-weatherization (Narum, Pigg, and Schlegel 1992). This limits the ability to make conclusions about long-lived measures.

In the last several years, a renewed scrutiny around the persistence of energy savings has emerged. With the previous studies dating several decades back, and often lacking long data histories, there is a need for new studies around persistence of energy savings (Hoffman et al. 2015). This paper seeks to address this gap by exploring the persistence of savings using data from Wisconsin's low-income weatherization program. It addresses two key research questions:

1. Can program-level gas and electricity savings be detected years after weatherization was completed?
2. How persistent are the energy savings by fuel type among different housing types? (e.g. manufactured homes, single family, and 2-4 unit multifamily)

This analysis utilizes data from Wisconsin's low-income weatherization program, Home Energy Plus (HE+), which blends funding from the federal Weatherization Assistance Program with state-level public-benefits funds derived from utility charges to weatherize 5,000 to 6,000 homes annually. The program has been evaluated annually since 2009 and has gathered extensive data in a consistent manner since the early 2000's. The program provides weatherization services through 19 agencies and 20 service areas across the state and is available to households meeting program eligibility requirements. This is defined as a household income of 60 percent or less than the state's median income for a similar-size household. The program targets homes with a high energy burden as well as those with elderly, very young, or disabled occupants. The main objectives of the HE+ program are to (1) reduce home energy bills, (2) save energy, and (3) make homes warmer in the winter and cooler in the summer.

Weatherization measures under this program fall under four general categories: space conditioning, water heating, shell measures, and other. Examples of these each category include a heating system replacement with a higher-efficiency model of gas furnace (space conditioning), gas power vent from conventional gas (water heating), air sealing and additions of attic insulation (shell), and lightbulb replacements with light emitting diode (LED) bulbs (other).

This study takes advantage of the fact that the Wisconsin program maintains a detailed database of participants and regularly assesses actual energy savings based on obtaining and analyzing utility consumption histories. With some additional work to update consumption histories for prior participants, a relatively large sample size was available for the study to examine overall persistence of savings for participants dating back to 2006.

The remainder of this paper details the methodology for this study, the results for gas and electricity savings, and general implications of the results.

Methodology

We define persistence at the program-level, as the examination of whether the savings achieved in the year following weatherization are maintained over time. The study used annual estimates of weather-normalized consumption derived from monthly billing histories for each home and involved two types of analysis: (1) a fairly simple review of changes in annual consumption over time by year of program participation (NAC method), which used pre-weatherization consumption as a baseline against which to measure savings; and (2) more complex regression modeling (regression method) that attempted to account for potential underlying non-program trends in consumption over time. Below, we describe the data-processing and analysis steps involved for both.

Data processing – NAC Analysis

The analysis utilized 13 years of weatherization records from the program databases as well as monthly utility billing data from the six largest utilities in the state of Wisconsin.

The service territories of the six utilities covers nearly the entire state of Wisconsin and the billing data spans from 2006 to 2018. We sought to include all available homes in the analysis regardless of occupancy changes over the period of analysis, as excluding these has long been known to introduce attrition bias into savings estimates (Blasnik 1989). However, not all utilities were willing to provide data for customers other than the original weatherization applicant. This resulted in a dataset that varied in billing-history completeness for individual buildings. To maintain a balanced dataset over the entire analysis period, for the main analysis we removed buildings that did not have a full 13-year billing history. Although this requirement decreases the available sample size – and likely introduces some overall attrition bias to the analysis – it avoids what we considered to be a larger problem of trying to discern persistence trends in imbalanced data where the composition of the study group varies from year to year. We also removed a small number of buildings where the weatherization period itself spanned multiple years, as this also created problems with the analysis, which relied on annual NAC estimates.

After all the data preparation and screening steps, this dataset includes about 8,000 weatherized homes heated with natural gas and about 8,700 weatherized homes heated with electricity, with single-family homes making up over 80 percent of all units. The homes in the dataset were all weatherized between 2007 and 2017. The final dataset represents around 25 percent of all the housing units in the original dataset, with the main attrition factor being the requirement that each unit have a full 13-year history of billing data. This dataset provides a sufficient sample to analyze the impacts of persistence based on NAC trends over time, however it is somewhat skewed towards the service territories of utilities that were willing to provide consumption data regardless of occupancy changes that occurred after weatherization.

We then combined the billing data with information from the tracking databases to determine when weatherization occurred for each home and to categorize homes by housing type. We analyzed the consumption data by calendar year, grouping homes according to the year in which they were weatherized and then defining years relative to weatherization accordingly. Table 1 shows the basic data structure laid out by weatherization cohort and calendar year. Note that the nature of the data inevitably means that while persistence in the first few years following weatherization can be gauged across participants treated in many program years, longer-term persistence is strongly determined by early program participants. The analysis is thus potentially sensitive to changes in program policies over time that might affect the types of homes treated or the measures installed. However, with one exception—a major program policy change related to weatherization of manufactured homes in 2015—our analysis of program tracking data showed the program to be fairly stable over the period of interest in terms of the types and incidence of installed measures.

Table 1. Basic data structure of analytic dataset

Wx Cohort	Calendar Year												
	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018
2007	Pre	**Wx**	Post 1	Post 2	Post 3	Post 4	Post 5	Post 6	Post 7	Post 8	Post 9	Post 10	Post 11
2008	Pre	Pre	**Wx**	Post 1	Post 2	Post 3	Post 4	Post 5	Post 6	Post 7	Post 8	Post 9	Post 10
2009	Pre	Pre	Pre	**Wx**	Post 1	Post 2	Post 3	Post 4	Post 5	Post 6	Post 7	Post 8	Post 9
2010	Pre	Pre	Pre	Pre	**Wx**	Post 1	Post 2	Post 3	Post 4	Post 5	Post 6	Post 7	Post 8
2011	Pre	Pre	Pre	Pre	Pre	**Wx**	Post 1	Post 2	Post 3	Post 4	Post 5	Post 6	Post 7
2012	Pre	Pre	Pre	Pre	Pre	Pre	**Wx**	Post 1	Post 2	Post 3	Post 4	Post 5	Post 6
2013	Pre	Pre	Pre	Pre	Pre	Pre	Pre	**Wx**	Post 1	Post 2	Post 3	Post 4	Post 5
2014	Pre	Pre	Pre	Pre	Pre	Pre	Pre	Pre	**Wx**	Post 1	Post 2	Post 3	Post 4
2015	Pre	Pre	Pre	Pre	Pre	Pre	Pre	Pre	Pre	**Wx**	Post 1	Post 2	Post 3
2016	Pre	Pre	Pre	Pre	Pre	Pre	Pre	Pre	Pre	Pre	**Wx**	Post 1	Post 2
2017	Pre	Pre	Pre	Pre	Pre	Pre	Pre	Pre	Pre	Pre	Pre	**Wx**	Post 1

Wx = Weatherization; Pre = Pre weatherization; Post *x* = *x* years after weatherization

We then weather normalized the data at the unit-level to account for the influence of year-to-year weather variation on household energy use. The models disaggregate the energy use each calendar year into space heating, cooling (on the electric side) and non-space-conditioning components based on the relationship between consumption and heating- and cooling-degree days, and then adjust heating and cooling use to long-term average weather, using a procedure that individually determines the best balance-point temperatures to create heating and cooling degree-days. Fitting the models to individual households versus the entire group of treated homes captures the unique energy-temperature relationship of each home and allows for a more accurate adjustment of observed whole-home-energy use to long-term average weather conditions. The result is a weather-normalized annual whole-home consumption value for each participant.

Data processing – Regression Analysis

To complement the data processing for the NAC analysis, we included some steps to address the particularities of the regression analysis. For this, we created two variants of the dataset described above and looked at results for these with and without screening for outliers. The first variant included all available data, including homes that did not have a full 13-year billing history. A second variant included only weatherization jobs with NAC values for the full 14-year (2006-2019) instead of 13-year (2006-2018) time span, without removing 2019 as was done above. Changes in the makeup of the participant population over time could also affect this analysis, but we lacked complete data on characteristics such as home size and age, so could not control for these.

Analysis methods

NAC method

We analyzed the data by examining the time trend of mean consumption for each weatherization-year cohort, both in terms of calendar year and by years before and after the year of weatherization. We also calculated the percentage change for each post-weatherization year compared to the average pre-weatherization-annual usage. Both methods provided a general view of how consumption patterns changed over the period from 2006 to 2018 for each weatherization cohort. This approach tells a basic story of the effect of weatherization on the population of weatherized homes in the program over the study period.

Regression method

To test the validity of a NAC-based approach, we used regression-based models to tease out persistence effects across cohorts while controlling for potential non-weatherization influences beyond weather variation. We also implemented upfront outlier screening and employed robust regression procedures to examine how NAC outliers might influence the results. Finally, for the regression analyses we relaxed the screening requirement for a complete billing history and examined how the results change if all available data are included. In all, the regression methods include four different models with and without limits to outliers and with all available data versus restriction to only those homes with full billing histories.

These persistence models are all based on the year-to-year change in weather-normalized annual consumption (ΔNAC), both before and after weatherization. Analyzing consumption *changes* instead of consumption *levels* in each year eliminates some complex issues associated with different homes or entire weatherization-year cohorts having different levels of consumption and focuses the analysis more directly on the goal of the analysis, which is to estimate how savings change over time.

The main predictor in the model is the number of years after weatherization (PostYr), which we represent as a series of indicator variables that are set to 0 or 1, with $PostYr_1$ taking a value of 1 for ΔNAC values associated with the first-year impact from weatherization, $PostYr_2$ representing the change between the first and second years after weatherization, etc.[1] These represent the difference in energy consumption from weatherization in reference to a base category of the year-to-year changes in energy consumption during pre-weatherization for later participants, represented as $PostYr_0$. The remaining categories ($PostYr_1$, $PostYr_2$, etc.) represent the change from this reference.

When the model is fitted, the coefficients associated with each PostYr term (β_{yr1}, β_{yr2}, etc.) thus capture the average change in NAC between two post-weatherization years, with β_{yr1} representing the average first-year impact of weatherization on usage, β_{yr2} representing the average change in savings from post-weatherization Year 1 to Year 2, etc. Shown in Equation 1, the cumulative persistence of savings for any given number of years after weatherization can then be calculated as the sum of the PostYr coefficients. For example, the overall persistence of savings in post-weatherization Year 5 is calculated as:

[1] Note that for the first-year savings from weatherization to be properly represented, the ΔNAC associated with $PostYr_1$ is calculated somewhat differently as the difference in NAC between the year immediately following weatherization and the year immediately preceding weatherization, thus skipping the one to two years associated with weatherization itself.

$$\text{Year 5 cumulative impact} = \beta_{yr1} + \beta_{yr2} + \beta_{yr3} + \beta_{yr4} + \beta_{yr5} \qquad \text{Equation 1}$$

In addition to the main PostYr coefficients of interest, several additional variables are included to help control for non-weatherization influences on ΔNAC as well as changes in the group composition over time. These include indicator variables for:

- Calendar year
- Weatherization-year cohort
- Agency

There are several ways to fit this type of model and we found that the results can be sensitive to the fitting procedure, so we provide comparative results for four model-fitting procedures:

- Ordinary least-squares (OLS)
- Robust regression
- Quantile regression
- Mixed-effects model

Ordinary least-squares (OLS) is what is typically employed in statistical regression modeling when the dependent variable is continuous. It fits the model by minimizing the sum of the squared differences between the observed data and the model. Modeling results from OLS show the most basic explanation of trend in changes in NAC across years and a theoretical baseline against which to hold more sophisticated models.

Robust regression is a variant of OLS that identifies and down-weights outliers in the data. We used this as a check on both our basic NAC model and OLS by showing the effect of homes where extreme high and low energy consumption was observed. We used a robust regression routine implemented by Stata (Version 15.0).

Quantile regression is another way of dealing with outliers. Instead of modeling average effects, quantile regression models medians, which are more resistant to extreme datapoints. We employed this method as a check on more basic models to discern the impact of extreme data on the results.

Mixed-effects modeling considers some predictors as “fixed effects” and others as “random effects,” the former being predictors of interest to the analysis and the latter being considered factors that, while potentially systematically influencing the data, are random influences that are not of interest to the analysis. Here we treat the PostYr_x terms as the fixed effects, and consider Calendar Year, Weatherization Year and Grantee as random effects in the context of measuring overall average persistence of savings. For the mixed-effects model, we also included weatherization job as a random effect to help control for time trends in consumption associated with individual homes. We treated Calendar Year, Weatherization-Year Cohort and Grantee as overall random effects (Level 1), and Weatherization Job as a Level 2 random effect within Grantee. Overall, mixed-effects allowed us to account for the possibility that correlation within levels like Calendar Year, Weatherization Year and Grantee could systematically influence the results.

In addition to the two dataset choices—all available data and only weatherization jobs with NAC values for the full 14-year time span—we implemented the analysis with and without upfront screening for outliers. The outlier screening involved dropping individual data points where the year-to-year change in consumption exceeded 40 percent for ΔNAC years that did not involve weatherization itself or 75 percent for the ΔNAC associated with weatherization (since weatherization itself can have a significant impact on consumption). These screens eliminated about 4 percent of the gas data and 7 percent of the electric data. Note that we did not run the robust or quantile models on the outlier-screened data, since those fitting procedures are already intended to be resistant to outliers.

Overall, the four fitting procedures, two dataset selections and two options for outlier screening yield 12 sets of results each for natural gas and electricity, as shown in the Table 2 below.

Table 2. Regression models and data

	All Available Data		Full-Span Data Only	
Fitting Procedure	Untrimmed	Trimmed	Untrimmed	Trimmed
OLS	Model 1	Model 2	Model 3	Model 4
Mixed Effects	Model 5	Model 6	Model 7	Model 8
Robust	Model 9		Model 10	
Quantile	Model 11		Model 12	

Results

NAC Method

Estimates of weather-normalized whole-home natural gas consumption throughout the study period suggests that weatherization savings persist through time. Figure 1 illustrates NAC as a percent change relative to years since weatherization separated by housing type. Weatherization has an obvious and significant impact on annual gas consumption, with a large percent change occurring in the first year after weatherization.

Across housing types, the story is slightly more nuanced. Single-family, site-built homes—which make up the bulk of the cases—show consistent and stable savings over time. Small multifamily buildings show a similar pattern, though with more variation in average savings across weatherization-year cohorts, perhaps due to the smaller number of cases in each cohort.

On the other hand, manufactured homes show much more variation across weatherization-year cohorts. Manufactured homes weatherized in 2015 or later actually show a small but distinct *increase* in gas consumption following weatherization. This is a result of the change in the program approach for gas-heated, manufactured homes during that period. The program switched from a computer-audit-driven approach that targeted major measures to a prescribed-measures-list approach that only called for minor gas measures—but that also allowed for fuel switching of electric water heaters to natural gas. For the most part, homes weatherized prior to this major program-policy change show persistence of savings. In contrast, the 2007 weatherization cohort shows a notable erosion of savings, which is likely attributable to the particularly small sample size for 2007 among manufactured homes.

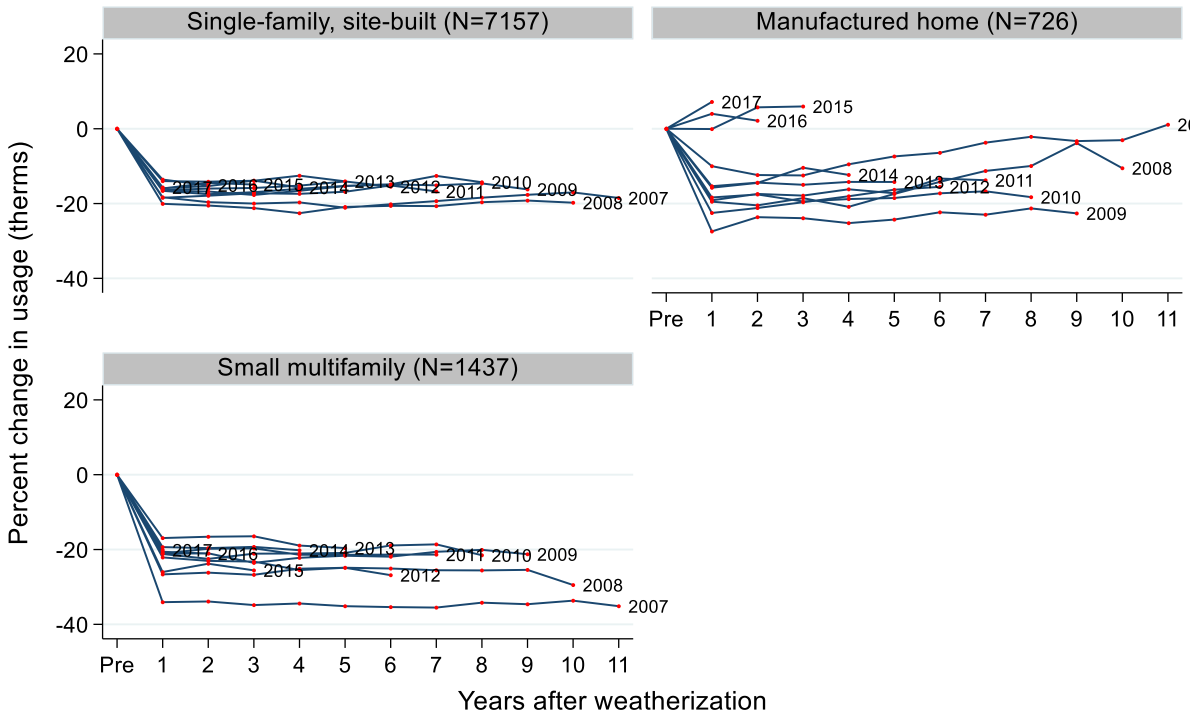

Figure 1. Percent change in mean NAC relative to pre-weatherization consumption by housing type and number of years after weatherization. Each line represents the average for homes weatherized in a given year.

The analysis of whole-home electric-savings trends over time shows results that were less clear than gas-savings trends (Figure 2). In general, compared to the gas results, there is more year-to-year variation in savings within weatherization-year cohorts as well as more variation in savings across cohorts. The general sawtooth pattern evident in the data is likely the result of calendar-year weather variation that is not adequately captured in the weather-normalization process.

Looking beyond the year-to-year variability, the results for single-family, site-built, and small multifamily homes do suggest some erosion in electricity savings over time, on the order of 2 to 5 percentage points per year. Manufactured homes show more year-to-year differences (likely due to smaller study-group sizes) and the 2007 to 2009 weatherization-year cohorts seem to lose most of their savings by the end of the period.

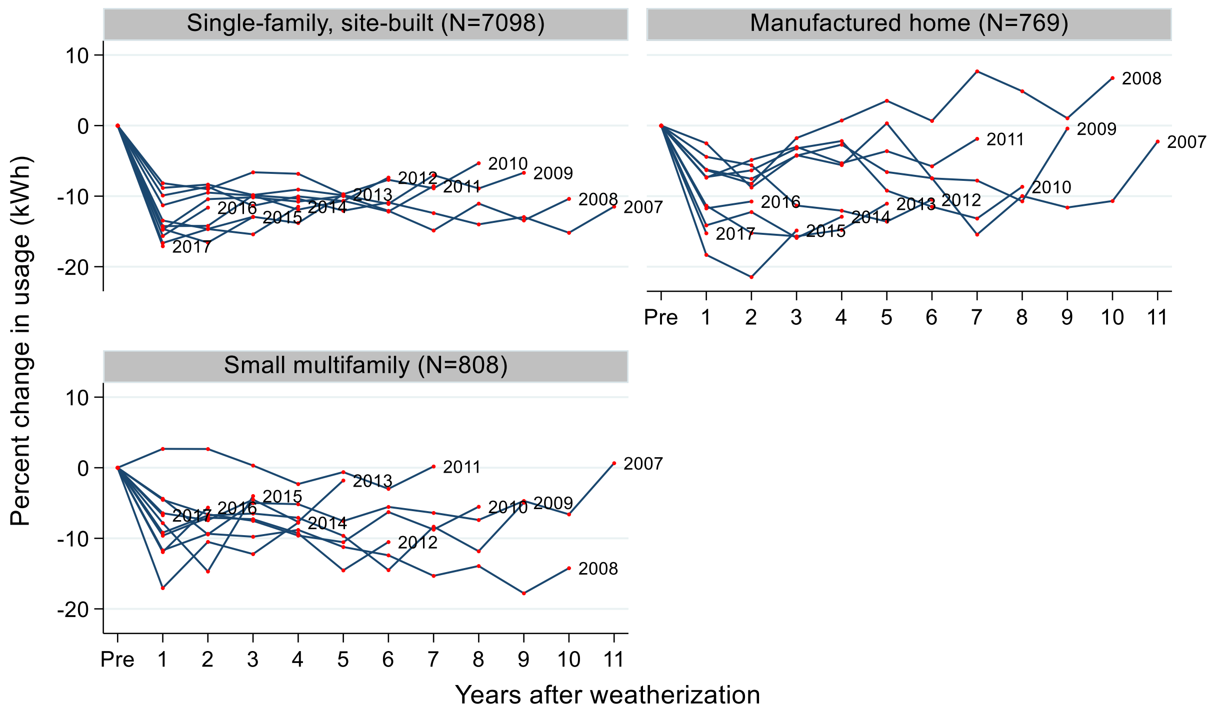

Figure 2. Percent change in electricity usage compared to pre-weatherization by housing type and number of years after weatherization.

Regression Methods

As described in the methods section, we ran 12 separate models that account for four fitting procedures, two dataset selections and two options for outlier screening (Table 1). Overall, the gas model runs generally suggest no erosion of savings over the analysis period, except for manufactured homes where a number of the models suggest savings erosion within the first 13 years following weatherization. The results for electricity, on the other hand, are more ambiguous, with single-family homes in particular having some models indicative of erosion of savings and others indicative of no erosion.

Additionally, comparing across the various model runs, outliers can strongly affect the results. The general effect of outliers is to implausibly suggest that savings increase strongly over time. For this reason, we favor Model 8, as the mixed-effects approach has conceptual appeal, the full-span dataset avoids issues with unbalanced data and the outlier trimming helps make the results more robust.

Figure 4 plots the gas results from Model 8, and strongly suggests persistence of gas savings for single-family and small multifamily homes and some evidence of erosion for manufactured homes.

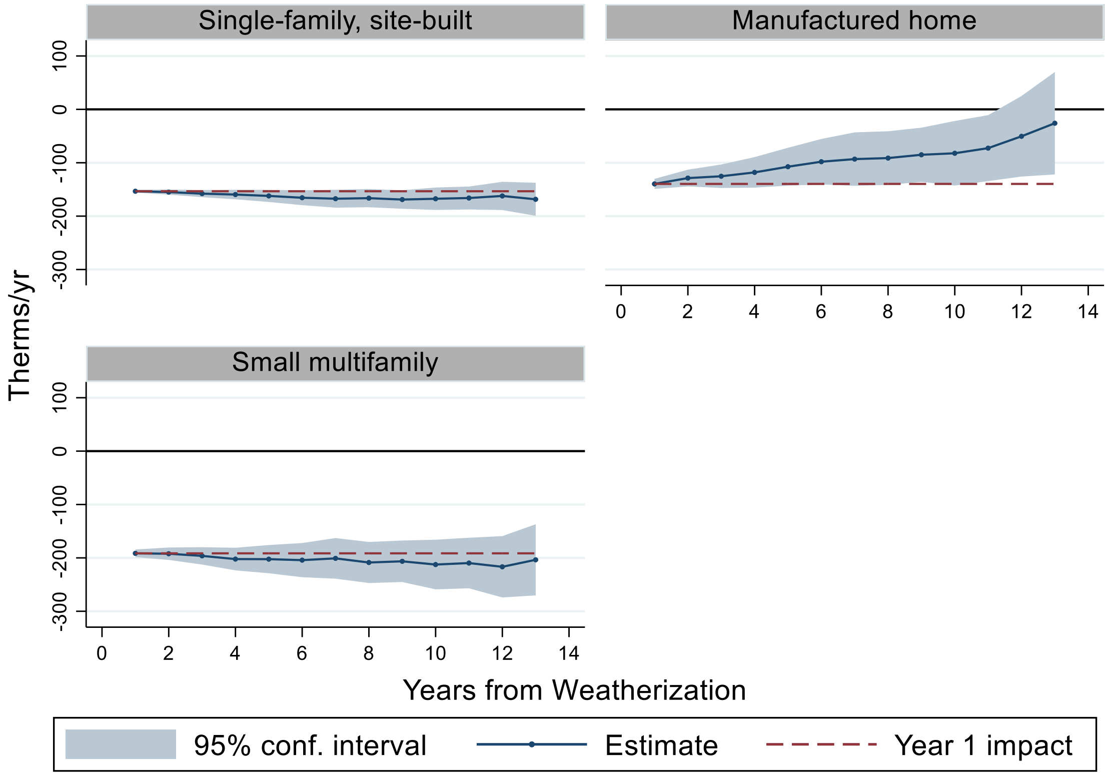

Figure 3. Estimated gas persistence for Model 8, by housing type, based on mixed-effects regression modeling.

For electricity, Model 8 suggests erosion of savings for single-family homes, though this is subject to the ambiguity from the varied results across models. Persistence of electricity savings for small multifamily and manufactured homes is also ambiguous under Model 8 due to the wide confidence bands for these housing types.

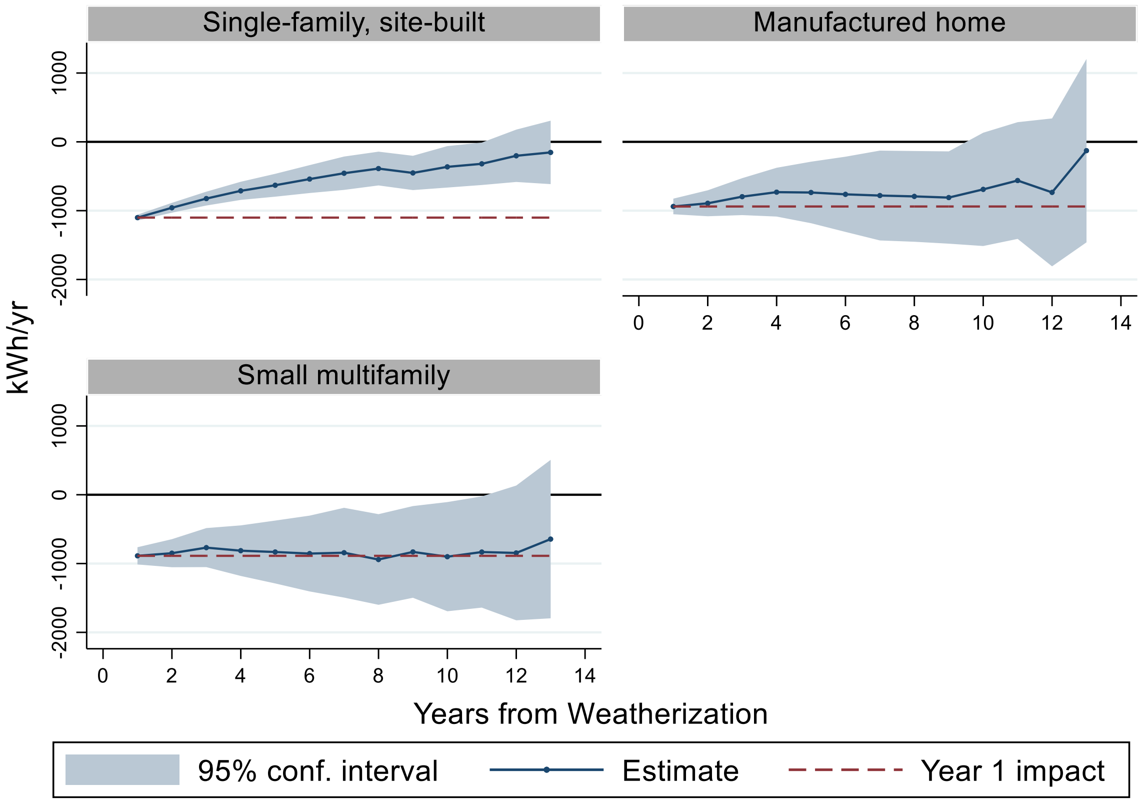

Figure 4. Estimated electricity persistence for Model 8, by housing type, based on mixed-effects regression modeling.

The results from the regression-based modeling supports our findings from the basic NAC method, bolstering the conclusion that gas savings are persistent across time while electric savings are less clear. Finally, we examined variants of Model 8 (mixed effects, full-span data, trimmed for outliers) that included binary predictors for selected individual measures installed under the program. The measures selected included those known to be key contributors to impacts from the program as well as measures such as water heater temperature reduction that have a short measure life and/or are easily disabled. Measure-level estimates of the persistence of impacts over time are obtained by including terms that interact the measure indicators with the PostYr indicator variables. For the most part, the results are either statistically ambiguous due to wide confidence intervals or show stable savings over time.

Conclusion

The results from this analysis suggest that natural gas savings for the program persist for at least twelve years after weatherization. Basic NAC trends show natural-gas-savings persistence for single-family homes to stabilize between 15 to 20 percent compared with pre-weatherization. However, the persistence narrative for electric savings is less clear, starting between 8 to 15 percent savings but show signs of erosion across the study period. Reasons for this remain unclear but potential paths for future research include investigating occupant behavior, maintenance, material degradation, and the success of targeting electricity measures in home weatherization compared to natural gas. An additional avenue for future research will be to control for participant population characteristics over time. Single-family homes show the most stability across weatherization cohorts for both natural gas and electricity. While the study provides no direct evidence of savings beyond the first dozen years, the fact that little erosion is seen after even 12 years for gas suggests that substantial savings are likely to be seen beyond the first dozen years. This is an important finding for low-income-energy-efficiency programs as it supports life-cycle cost-effectiveness calculations that presume longevity of installed gas measures, particularly mechanical systems, and shell measures.

Still, due to the scarcity of peer studies for comparison in methods and results, more research is needed to better determine the durability of benefits delivered by low-income weatherization programs.

References

Blasnik, M. 1989. "Attrition Bias in Fuel Savings Evaluations of Low Income Energy Conservation Programs." In Energy Program Evaluation: Conservation and Resource Management – The Fourth International Conference in Chicago. Chicago, Illinois.

Hoffman, Ian, Steven Schiller, Annika Todd, Megan Billingsley, Charles Goodman and Lisa Schwartz 2015. "Energy Savings Lifetimes and Persistence: Practices, Issues, and Data." Berkeley, California: Lawrence Berkeley National Lab. https://eta-publications.lbl.gov/sites/default/files/savings-lifetime-persistence-brief.pdf

Narum, David, Scott Pigg and Jeff Schlegel 1992. "Looking Past the First Year: Do the Savings Last? A Study of the Persistence of Energy Savings in Low-Income Wisconsin Residences." Madison, WI: Wisconsin Energy Conservation Corporation. https://www.osti.gov/servlets/purl/7084606

Vine, Edward. 1992. "Persistence of Energy Savings: What Do We Know and how Can It Be Ensured?" Berkeley, CA: Lawrence Berkeley National Laboratory. https://www.osti.gov/servlets/purl/10180664

Violette, Dan 2013. "Chapter 13: Assessing Persistence and Other Evaluation issues Cross-Cutting Protocols." Golden, CO: National Renewable Energy Laboratory. https://www.energy.gov/sites/prod/files/2013/11/f5/53827-13.pdf