Research proposal: *"Winter is coming: do energy efficiency improvements reduce households energy burden?"*

Guillaume Wald, CERNA, Mines Paris - PSL
September 2023

1 Research question

We want to measure the effect of home energy efficiency improvements on household energy burden, *all other things equal*. Indeed, a growing number of policies and programs aim to increase investment in energy efficiency, because conventional wisdom suggests that people fail to take up these investments even though they have positive private returns and generate environmental benefits. We want to test whether the conventional wisdom that energy efficiency investments decrease energy burden is correct. We won't dig into the causal mechanisms behind the (under) performance of such investments: both the *rebound effect* [1] and moral-hazard issues in work provision [2] would be included in our estimate.

2 Data

The *American Housing Survey* is a longitudinal biannual survey conducted since 1976. It gathers self-reported information from observations drawn from a pool of representative housing units. The housing unit is the baseline unit of observation: it should not be confused with the household. Hence, a given housing unit can be inhabited by different households across time. We will gather years 2015, 2017, 2019 and 2021 so as to track specific housing units. Your first task is to download the data and consolidate a panel with all housing units present continuously across the 4 years.

The AHS contains a lot of variables (see the Codebook), and we will focus on the following ones:

1. **Yearly energy burden**:
   We will combine *ELECAMT*, *GASAMT* and *OILAMT* (electricity, gas and oil burden in USD) to define each household's energy burden.

2. **Energy-efficiency investments**:
   The AHS provides information on energy-efficiency investments such as replacement or repair of roof, windows, as well as installation of insulation, central air system, built-in heating system and water heater. Feel free to expand this list with all relevant retrofits. For the years 2017-2021, one can select home improvements of interest by looking at the variable *JOBTYPE* in the research area. You may also take a look at related variables in the subtopic *Home improvements*, such as *JOBCOMPYR* (Year home improvement job was completed), *JOBCOST* (Cost of home improvement job) or *JOBFUNDS* (Main source of funding for home improvement job). Notice that the variable *RAS* up to 2013 also recorded the "type of alteration or repair" (which could be used for linkages with previous years).

3. **Socio-economic characteristics** of household:
   We look for all socio-economic variables that are possibly varying within a given housing unit across years. We will use the householder age, sex and race (1 among 21 possible choices), as well as ownership status. The AHS also includes household level information, such as its year of entry in the unit, the size of the household and the annual income. Again, feel free to expand this list with all relevant varying characteristics.

4. **Housing unit's characteristics and environment**:
As for household characteristics, we want to control for varying determinants of energy use at the housing unit level. We include the square footage, the year of construction, the number of housing units in the building, as well as the main heating equipment of the unit (1 among 14 possible choices). Again, feel free to expand this list with relevant characteristics. Do not forget to keep the SMSA information (Metropolitan area) to match each housing unit with the additional variables below.

Other data sources to be matched on the SMSA code are as follow:

- **Contemporaneous and past Heating Degree Days (HDD)**
Because we look at variations in energy use, we need to control for the level of HDD within each Metropolitan Area in the U.S. We will also use past HDD as an instrument in our IV strategy later on. HDD can be found on the NASA website, (see the *Gridded Monthly Temperature Anomaly Data* section). We will use the Surface air temperature (no ocean data) compressed NetCDF files.

- **Localized energy prices**
We will need to control for local prices in gas, electricity and heating oil. You may find times series data at the State level for electricity, gas and heating oil (download the *Excel file*) on the EIA website.

3 Econometric approach and challenges

3.1 Baseline binary treatment

We estimate the effect of energy efficiency investments on the housing unit energy burden. To do so, we use a Two-Ways Fixed-Effet (TWFE) model. Our analysis therefore resembles a lot that of Blaise and Glachant (2019) [3], which also uses a TWFE model. Here, we exploit the *within* housing unit variation across years (μ_i is the HU fixed effect and η_t the year FE).

We will start with a very simple framework in which retrofits will just be considered as a dummy variable taking value one if investment occured at time $t-1$:

$$R_{i,t} = \begin{cases} 1 & \text{if a retrofit has been implemented in } i \\ & \text{between the household's entry into the unit and } t \\ 0 & \text{otherwise.} \end{cases}$$

Thus, the TWFE model writes as follows:

$$y_{it} = \beta_1 R_{it} + \beta_2 X_{it} + \beta_3 W_{it} + \mu_i + \eta_t + u_{it}$$

We will carefully include the control variables related to household (X_{it}) and housing unit (W_{it}) characteristics, so as to rule out confounding variations within HU across years, which could bias our estimate of β.

3.1.1 First-Difference (FD) vs Within (FE) estimation

As $Cov(\mu_i, R_{i,t}) \neq 0$, we are in the case of **correlated individual effects** (see these M1 APE slides by *Luc Behaghel* (PSE) on the different estimators and their design requirements) To get rid of μ_i, we can use two strategies:

1. First approach: first-difference model:

   $$y_{i,t} - y_{i,t-1} = \beta_1(R_{i,t} - R_{i,t-1}) + \beta_2(X_{i,t} - X_{i,t-1}) + \beta_3(W_{i,t} - W_{i,t-1}) + \eta_t - \eta_{t-1} + u_{i,t} - u_{i,t-1}$$

   often denoted as:

   $$\Delta y_{i,t} = \beta_1 \Delta R_{i,t} + \beta_2 \Delta X_{i,t} + \beta_3 \Delta W_{i,t} + \Delta \eta_t + \Delta u_{i,t}$$

2. Second approach: within model:

   $$y_{i,t} - \overline{y_{i,t}} = \beta_1(R_{i,t} - \overline{R_{i,t}}) + \beta_2(X_{i,t} - \overline{X_{i,t}}) + \beta_3(W_{i,t} - \overline{W_{i,t}}) + u_{i,t} - \overline{u_{i,t}}$$

The main pitfall of this strategy is that the *within* and *FD* models do not enable us to estimate the coefficients on fixed-effects (μ_i).

Because these are two methods aiming at getting rid of the same empirical issue, one may ask which one should be preferred. In the case where $R_{i,t}$, our regressor of interest, is exogenous (that is, $Cov(R_{i,t}, u_{i,t}) = 0$, the two estimators $\hat{\beta_1^{FE}}$ and $\hat{\beta_1^{FD}}$ should be equal (modulo the sampling variations imputable to the FD transformation which looses one year of observation). Hence, a statistically significant difference between the two estimators indicates a possible endogeneity of $R_{i,t}$. More precisely, Woolridge (in Econometric Analysis of Cross Section and Panel Data, Chapter 10.7.1), shows that the FE estimator is more efficient when the errors are serially uncorrelated (namely, $Cov(u_{i,t}, u_{i,t-1}) = 0$), while the FD estimator is more efficient when the error follows a random walk (i.e $u_{i,t} = \phi u_{i,t-1} + \varepsilon$), with $|\phi| < 1$ and ε a random coefficient.

At this early stage of the research project, it is a good idea to think about which estimation method is the most appropriate. Then, you can compare both estimators using the Hausman test as described by Woolridge and conclude on the exogneity of $R_{i,t}$.

3.1.2 Heterogeneous effects

The econometric literature has recently shown that so-called Two-Ways Fixed-Effect regressions may produce *misleading estimates* if the policy's effect is **heterogeneous** *between groups* or *over time*. To understand the bias implied by these recent findings, see section 1, 2.1 and 2.2.1 in *de Chaisemartin and D'Haultfœuille* (The Econometrics Journal, Sept. 2023) [4]. According to section *2.1*, the risk of *sign reversal* of the coefficient is low in the case of β_1 as the treatment embedded in $R_{i,t}$ is staggered and binary, with a lot of untreated units even in the last period. However, section *2.2.1* indicates that the treatment effect (TE) may be biased (no sign reversal, but still negative weights) if we face heterogeneity in the treatment across time (i.e exactly what has been documented by *Peñasco and Anadón* (2023) [5]. Such a dynamic pattern cannot be tested using an Event-Study à la *Borusyak et al.* (2021) [6], as we only have 4 periods of data while we would at least need 5 (two before, two after and one at the time of the treatment).

To account for this risk of bias implied by dynamic TE, one can try to implement the *twowayfeweights* R command (see *Zhang and de Chaisemartin* (2021)) to compute the weights $W_{g,t}$ in section 2.2 of *de Chaisemartin and D'Haultfœuille* (2023) as a diagnosis for potential bias. This can be done for both estimators of β_1 relying on either the *within* or the *FD* transformation.
Building on the result of this test, one could push forward and implement $E(DID_+)$ for both estimation methods as in *de Chaisemartin and D'Haultfœuille* (2023) section 3.1, and compare the results to those obtained with the standard estimators of β_1^{FE} and β_1^{FD}.

3.2 Continuous treatment

An alternative specification would rely on a continous regressor accounting for the cumulative cost of each operation, with increments from $t-1$ increasing the value of the treatment at t. More precisely, this regressor is built as the cumulative sum across years within housing units, of the variable *JOBCOST* accounting for the USD cost of the retrofits. Hence, instead of $R_{i,t} = \{0, 1\}$, one could define $R_{i,t}$ as follows:

$$\mathrm{R}_{i,t} = \sum_{k=1}^{t-1} JOBCOST_{\,i,k}$$

This new definition does not have specific implications for the *FD* vs *FE* debate. You can follow the steps in 3.1.1 to define your favorite specification. Nevertheless, we will stick to the *FE* estimation for the remaining of the project, as it is the most common way employed in the literature.

We then look at the consequences of such a continuous treatment on the heterogeneous effects of the treatment. We follow subsection *2.2.2 More forbidden comparisons when the design is not staggered or treatment is not binary* in *de Chaisemartin and D'Haultfœuille* (2023). In the previous, binary treatment case, heterogeneity of the treatment effect between groups would not bias the estimation provided that the treatment effect was constant across time for each given group. With a continuous treatment, time-constancy of the treatment effect for each group is not enough anymore: one would need an equal treatment effect between groups (which may still receive different treatment intensities) so as to ensure a convex combination of treatment effects.

To ensure a consistent estimation, one can rely on the extension of the DID_M estimator to continuous treatments detailed in *de Chaisemartin, C., X. D'Haultfoeuille, F. Pasquier, and G. Vazquez-Bare* (2022): *Difference-in-differences estimators of the effect of a continuous treatment*. This new estimator relies on the distinction between *movers* (units which treatment changes from t to $t+1$) and *stayers* (units which treatment does not change from t to $t+1$). Then, it compares the outcome evolution of *movers* and *stayers* with the same period-one treatment. The difference in treatment effect is then compared to the treatment evolution, eventually reweighting stayers by propensity score weighting or by adjusting movers' outcome change using a regression of the outcome change (the treatment effect) on the period one treatment among the *stayers*.

3.3 Instrumental Variations

Based on the comparison between $\beta_1^{\hat{FE}}$ and $\beta_1^{\hat{FD}}$, we can evaluate the endogeneity of our regressor $R_{i,t}$, either in its bonary or continous form. Indeed, despite a very detailed AHS database and a cautious analysis of heterogeneous treatment effects, we cannot control for all the within housing unit, across years variations that are *not imputable to variations* in R_{it}. Our TWFE $\hat{\beta}$ will remain a biased estimate. More precisely, we face an **endogeneity issue** because energy use is directly affected by shocks that may also affect energy efficiency investments: For instance, if a housing unit turns *Green*er, it could choose to consume less, and at the same time invest in energy efficiency. Hence, the marginal effect of the retrofit would be contaminated by the contemporaneous change in energy use. Conversely, if a household head anticipates a future unemployment period (which is unobservable in, say, income data), she would probably decrease energy use. However, we do not know the effect of such anticipation on energy efficiency investment. Notice that one could multipliate such examples of within-HU variations affecting both our outcome and our treatment.

To tackle this endogeneity issue, we will rely on an IV strategy [7] based on local HDD vintages, which are correlated with investments in energy efficiency but not with current determinants of the energy burden. We will mostly focus on the estimation of a clean TWFE model and on its instrumentation with HDD vintages. Because we are somewhat constrained on the time dimension of our panel dataset, we will stick to the *Within* estimation for both steps of our $2SLS$ strategy. We begin with the binary treatment, and then turn to the continuous one. Finally, we will deal with heterogenous treatment effects, implementaing the appropriate estimator for each specification as defined above.

References

[1] Meredith Fowlie, Michael Greenstone, and Catherine Wolfram. "Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Program". In: *The Quarterly Journal of Economics* 133.3 (2018), pp. 1597–1644. URL: `https://academic.oup.com/qje/article-abstract/133/3/1597/4828342`.

[2] Louis-Gaëtan Giraudet, Sébastien Houde, and Joseph Maher. "Moral Hazard and the Energy Efficiency Gap: Theory and Evidence". In: *Journal of the Association of Environmental and Resource Economists* 5.4 (2018), pp. 755–790. URL: `https://EconPapers.repec.org/RePEc:ucp:jaerec:doi:10.1086/698446`.

[3] Gaël Blaise and Matthieu Glachant. "Quel est l'impact des travaux de rénovation énergétique des logements sur la consommation d'énergie". In: *La Revue de l'énergie* 646 (2019), pp. 46–60.

[4] Clément de Chaisemartin and Xavier D'Haultfœuille. "Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: a survey". In: *The Econometrics Journal* 26.3 (June 2022), pp. C1–C30. ISSN: 1368-4221. DOI: `10.1093/ectj/utac017`. eprint: `https://academic.oup.com/ectj/article-pdf/26/3/C1/51707976/utac017.pdf`. URL: `https://doi.org/10.1093/ectj/utac017`.

[5] Cristina Peñasco and Laura Díaz Anadón. "Assessing the effectiveness of energy efficiency measures in the residential sector gas consumption through dynamic treatment effects: Evidence from England and Wales". In: *Energy Economics* 117 (2023), p. 106435.

[6] Kirill Borusyak, Xavier Jaravel, and Jann Spiess. "Revisiting event study designs: Robust and efficient estimation". In: *arXiv preprint arXiv:2108.12419* (2021).

[7] Joshua D Angrist, Guido W Imbens, and Donald B Rubin. "Identification of causal effects using instrumental variables". In: *Journal of the American statistical Association* 91.434 (1996), pp. 444–455.