The Efficacy of Energy Efficiency: Measuring the Returns to Home Insulation

Linde Kattenberg* Piet Eichholtz† Nils Kok‡

October 2024

Abstract

Energy efficiency in the housing market is key in reducing energy consumption and carbon emissions, as well as to enhance national energy independence and protect consumer budgets. Insulation plays an important role in improving the energy efficiency of a home. However, studies of the impact of insulation measures on actual gas consumption are typically based on engineering predictions, and the efficacy of insulation measures is subject to debate. This study exploits a large sample of home insulation interventions, combined with detailed household data on actual gas consumption before and after these interventions, and information on the socio-economic characteristics of occupants. Using a difference-in-difference approach, we document that home insulation reduces gas consumption by about 19%, on average, both for owner-occupied and rental homes. For the latter, the treatment is plausibly exogenous. Importantly, we find strong evidence of persistence: the reduction in gas consumption is consistent up to ten years after the intervention. The average treatment effect translates into a €350 reduction in the annual gas bill, and an average rate of return of 19.9% on the initial investment.

JEL Codes: Q01, Q31, Q40

*Maastricht University, The Netherlands; l.kattenberg@maastrichtuniversity.nl

†Maastricht University, The Netherlands; p.eichholtz@maastrichtuniversity.nl

‡Maastricht University, The Netherlands; n.kok@maastrichtuniversity.nl

[1]We are grateful for the comments of Erdal Aydin, Dirk Brounen, and seminar participants at the AREUEA International Conference Dublin, the 2022 ODISSEI Conference, the MIT Climate and Real Estate Initiative Symposium, the ERES Annual Conference London, and the EAERE Annual Conference Limassol. We thank the Dutch Ministry of the Interior for funding this research project, Bameco B.V. for providing their data on insulation interventions, and Floor van Gulik for excellent research assistance. All errors pertain to the authors.

1 Introduction

Real estate plays a key role in the reduction of carbon emissions needed to mitigate climate change. Housing alone accounts for about 27% of final energy consumption in the EU (Eurostat, 2020). At the same time, many CO_2-reducing interventions are readily available in the housing sector – instruments to improve the energy efficiency of a home include, for example, solar photovoltaics, triple glazing, heat pumps, and cavity wall, roof, and basement insulation (Granade et al., 2009). In addition to reducing CO_2 emissions, the improved energy efficiency of homes can lead to a reduction in monthly energy expenses of households and improved living comfort. Indeed, popular belief holds that investing in home energy efficiency measures provides a relatively high return on investment, but that belief is typically not based on hard evidence. The exact returns on energy efficiency investments in homes are difficult to quantify, simply because realized energy savings are home-specific, difficult to observe, and ultimately prone to selection bias (and thus hard to generalize).

The uncertainty around actual energy savings from investments in energy efficiency is an important consideration in the discussion of the 'energy efficiency gap.' This term has been coined to explain the slow uptake of energy efficiency measures for society in general, and for homes specifically (Jaffe and Stavins, 1994). However, the size of the gap is up for debate (Allcott and Greenstone, 2012), as different factors could lead to overoptimistic predictions of profitability of energy efficiency investments. Numerous studies report a sizable disparity between expected and realized savings from energy efficiency retrofits in the housing market (Allcott and Greenstone, 2017; Fowlie et al., 2018; Christensen et al., 2023), with multiple explanations for the wedge. For instance, Christensen et al. (2023) find that heterogeneity in quality of installment, mistakes in engineering estimates, or a rebound effect are all factors that influence the gap between expected and realized energy savings. More accurate estimates of predicted savings and the profitability of energy efficiency investments can contribute to a better understanding of the true size of the energy efficiency gap.

The key contribution of this paper lies in the fact that we estimate the returns to home insulation across the whole housing market, covering homes in the rental and owner-occupied sector. Therefore, we are able to address the selection effect that is present in measuring in the energy consumption reduction after insulation by households that select themselves into treatment. Moreover, we know dwelling-level costs and energy savings for a large sample of homes to precisely estimate the return to a variety of insulation measures across heterogeneous households. We examine the effect of cavity wall, basement, and roof insulation on the actual energy consumption of households in the Netherlands, allowing for a detailed calculation of return on investment. Given that cavity walls are prevalent throughout Europe (Magrini et al., 2022), and that roof and basement insulation can be applied to most of the modern housing stock, we argue that the setting of this study can be generalized to a large part of the Western housing stock. Using hand-collected, proprietary information from a large insulation company, we identify homes where insulation measures were applied, including details on the type of insulation and the installation costs. We link this information to annual data on gas and electricity consumption, observable characteristics of the home, and extensive micro data on the household, including income, age, and the number of household members.

The data set includes both rental and owner-occupied homes. Furthermore, we can split the sample of rental homes into the homes that are owned by a social housing institution (i.e. affordable housing), and those that are rented out at market rate (by a for-profit owner). This distinction is important, given that the choice to install insulation is plausibly exogenous for tenants (i.e. the landlord decides on such measures), especially for social housing institution, where home maintenance is systematic rather than ideosyncratic, and thereby less likely to be influenced by the occupant. We empirically assess the energy consumption of treated homes, before and after the implementation of the insulation measures, constructing a control group of comparable homes where no insulation measures took place. We assess the causal effect of home insulation on household gas consumption using a difference-in-difference estimation method.

The results of the empirical analysis show that gas consumption decreases, on average, by just under 19% after insulation is installed. Cavity wall insulation leads to the largest decrease in gas consumption, followed by the effect of roof insulation, and basement insulation. The results are robust to a variety of robustness checks, including the construction of different control samples. Importantly, these results hold for both owner-occupied and rental homes. The reduction in gas consumption does not differ substantially across household and dwelling types, and gas consumption is not substituted by an increase in electricity consumption. In fact, we find some evidence of within-household spillover effects, with electricity consumption decreasing following an insulation intervention. Our 2010-2020 data set also allows for an analysis of the persistence of energy savings over time. Unlike Peñasco and Anadón (2023), we document that the effects of home insulation are persistent – the observed reduction in gas consumption remains stable for up to ten years after insulation measures have been taken.

A simple cost-benefit calculation on the economics of insulation investments indicates that, for an average household living in an average home in our sample, the average annual gas bill is reduced by €350, based on gas prices at the time of the investment. These savings translate into a payback period of 5 years. Assuming perpetuity of the savings, the return to insulation measures is 19.9% using gas prices at the time of investment (rather than using the assumption of perpetuity, in case of a home sale the capitalization of energy savings in the home price would represent part of this return, see Aydin et al. (2020)).

This paper adds to the broader literature on energy efficiency in the residential sector. The early literature focused primarily on understanding the cross-sectional and temporal variation in household energy consumption patterns (see, for example, Brounen et al. (2013)), while a more recent strand of literature attempts to identify the effect of behavioral interventions to reduce energy consumption (Allcott and Mullainathan, 2010; Aydin et al., 2018). There are some studies that empirically assess the effect of improved insulation (Metcalf and Hassett, 1999; Hong et al., 2006; Liang et al., 2018; Peñasco and Anadón, 2023), or weatherization more broadly (Schweitzer, 2005; Allcott and Greenstone, 2017; Fowlie et al., 2018; Christensen et al., 2023). Generally, these studies find a sizeable decrease in energy consumption after insulation measures, or from a combined package of measures that include insulation. However, studies comparing actual energy savings to projected savings based on engineering estimates find a large performance gap. For example, Allcott and Greenstone (2017), Fowlie et al. (2018), and Christensen et al. (2023) find realized energy savings of just 58%, 30%, and 51% of predictions, respectively. Importantly, rather than comparing engineering estimates with actual energy savings, this paper focuses on the initial investment versus ex-post monetary savings on the utility bill, allowing for the estimation of a rate of return on insulation measures [1]. However, these papers are considering a setting where households self-select into a subsidy program, often targeted at low-income households. Our paper will compare households that choose to invest in insulation measures, with those households where insulation improvements are plausibly randomly assigned (since the landlord decides on insulation improvements for rental tenants). Thus, our setting allows us to reflect on the potential issue of self-selection into treatment in measuring the returns to energy efficiency improvements.

In the literature, the size effect of energy efficiency interventions often varies, is context-dependent, and is typically heterogeneous over time. For instance, Peñasco and Anadón (2023) find a significant effect of insulation measures on energy consumption in the short term, but the effect disappears within 4 years. This may be related to other measures households when renovating their home, such as increasing the living space, which increases demand for energy. An alternative reason for the lack of persistence in savings from energy efficiency measures can be the "rebound effect" – an increase in energy efficiency of the home can lead to an *increase* in the use of energy-consuming appliances, because the unit cost of energy consumption decreases. In a meta-study of Sorrell et al. (2009), an average rebound effect of 20% is documented across 21 studies on household energy consumption in the OECD. In the Netherlands, for a sample of 560,000 Dutch dwellings, Aydin et al. (2017) document a rebound effect of 41.3% in rental homes, and of 26.7% in owner-occupied homes. Such an effect would be important to consider in order to make a reliable estimate of the true savings in gas consumption following home insulation measures. The results of our study cover a relatively long time period – we include treated observations where home insulation took place more than 10 years ago. As such, we also measure the long-run effects of home insulation, documenting strong persistence in the results.

The results in this paper have some implications for homeowners, (public) investors in residential homes, as well as policymakers. Given the paucity of reliable information on

[1]For instance, Christensen et al. (2023) examine unit-specific net benefits of home insulation measures and find that 42% of homes in their sample, while underperforming predictions, have a positive net benefit from investment through energy savings. The authors highlight that the presence of a performance "wedge" does not tell the full story and that the investment can still be profitable.

the efficacy of home insulation measures, it is often challenging for a homeowner, be it an owner-occupier or a landlord, to make well-informed decisions regarding the investments needed to improve the energy efficiency of the building. The return calculations in this paper may help to provide further insight into the real, monetary effects of insulation programs. The results also indicate that blanket subsidy programs should, in principle, not be necessary for home insulation, given the short payback period and high return on investment. Subsidy programs could be targeted at homeowners with limited net wealth, or rather be changed to loan programs, to overcome upfront financing constraints. In addition, government policy efforts may be directed to the housing rental market, where investors incur the capital cost and tenants typically benefit from energy savings.

Importantly, our return calculations include the dampening effect of a possible "rebound" effect, and thus reflect the true financial return to consumers or investors. Of course, we cannot observe the presence of an immediate rebound effect. That is, the difference between actual energy savings and energy savings based on engineering predictions that can be attributed to a behavior change *immediately after* the intervention (Fowlie et al., 2018; Christensen et al., 2023). In case of such immediate rebound, the total welfare effect would also include the consumer benefit of additional heating. In addition, we ignore the possible welfare effects from enhanced comfort through reducing cold and draft. As pointed out by Palacios et al. (2021), these effects may include reduced incidence of illness and frequency of doctor visits, which also has broader societal welfare effects.

The remainder of this paper proceeds as follows. We first provide a brief overview of the data sources used in the paper, including sample statistics and the results of the parallel trend analyses. Section 3 elaborates on the methodology. The regression results are presented in section 4. The final part of the paper includes a section on implications for homeowners and policymakers, based on a range of cost/benefit analyses of the insulation measures, and a conclusion.

2 Data

The two main sources of data used in this paper are Bameco BV, a private insulation company based in the Netherlands, and the Dutch Central Bureau of Statistics. Bameco is a large insulation provider in Limburg, the most southern province of the country. Its sole business is home insulation, with a focus on cavity wall, basement, and roof insulation.[2] The company maintains a (paper) archive for each home where an insulation intervention was carried out, including information on the cost, type, and date of the installation. We manually digitized these data over the full period of operation, which started in 2010. In the sample, we include all insulation measures up to 2019, such that we have at least one year for post-insulation measurement of energy consumption.

In total, we identify 2,240 households with a home insulation intervention in the period between 2010 and 2019. Figure 1 provides an overview of the insulation interventions for every year in the sample. Note that households can opt for just a single measure, or for multiple measures at the same time. Clearly, wall insulation is the most popular form of insulation, with 80.6% percent of households opting for that measure. Basement insulation is applied in some 9.7% of the sample, and roof insulation in 9.6% of cases. There is a clear upward trend in insulation interventions over the sample period – the 2016 dip likely represents an artifact of the data collection rather than a true decrease in interventions, given that some of the archive for that year was no longer retrievable due to a change in administrative systems. However, the proportions of the different insulation measures in that year look similar to the other years, such that the missing observations do not create a bias in our estimation.

—Insert Figure 1—

Table 1 provides further insight into the insulation measures. In our sample, 91% of households include just one measure, 8% include two measures, whereas three measures are

[2]Using data from a single company may lead to concerns about external validity, stemming from between-company differences in technology, human capital, and other unobserved characteristics. For example, Christensen et al. (2023) document that variation in ”craftmanship” accounts for large differences in subsequent energy savings. However, most insulation companies in the Netherlands use similar insulation techniques, and the application of insulation in cavity walls, floors and roofs is relatively straightforward, such that we do not expect large differences in outcomes between companies.

rarely taken at the same time. From an investment perspective, the average investment for wall insulation equals some €1,603 (in nominal terms), which is about 0.7% of the home value at the time of the intervention. Roof insulation is the most expensive intervention, whereas basement insulation is the cheapest form of home insulation[3]. We differentiate between homes that are owner-occupied (in the Netherlands, the homeowner rate is 57%), owned by social housing institutions (more than a third of the Dutch stock) and homes owned by private investors (around 5% of homes in the Netherlands). Homes owned by social housing institutions are regulated, with rents considered "affordable," while the latter are typically rented out at market prices. There is quite some difference between the insulation types installed in owner-occupied homes versus the insulation that is applied to homes owned by social housing institutions – private owners hardly opt for roof insulation, whereas social housing institutions are more likely to install roof (and basement) insulation. The difference in choice for the type of insulation per housing category may be related to the types of homes in each category. For instance, the share of apartments is much higher in the social housing sample, as compared to the rest of the sample. In the case of an apartment, wall insulation can be less beneficial from an energy and financial savings perspective as compared to other dwelling types.

—Insert Table 1—

The insulation interventions are matched to microdata on household and dwelling characteristics provided by the Dutch Central Bureau for Statistics (CBS). Each observation is matched to the CBS files, where we include all treated observations until 2019, which requires household data until 2020 (one year after the last intervention). Our unit of observation is a household living in a certain home. That is, when a household moves to a different home, the later observations are dropped from the sample. Out of the 2,240 observations, we have energy consumption and household data on 2,023 observations. The control group in our baseline analysis consists of a random 1% sample of all households that are based in the same province (Limburg), leading to 2,994 households in the control sample

[3]Note that the investment costs do not incorporate local, regional or national subsidies. Such subsidy programs have come in and out throughout the sample period, and while they may influence the propensity to insulate, such subsidies should not affect the gas consumption outcome of the intervention.

(we restrict the control group in the robustness checks of the analysis).

Table 2 provides the descriptive statistics of the treatment and the control group – the descriptive statistics are based on the year before any observation was treated, 2009. We distinguish treated observations that are owner-occupied from observations that are rental homes in the free market sector, or owned by a social housing institution. Quite clearly, in all sectors we observe that owners of homes with higher gas consumption are more likely to opt for insulation measures. Semi-detached homes, which have more exposed walls, are more likely to be insulated. We also observe that homes constructed between 1945 and 1980 have a higher propensity to be treated. Most homes in the Netherlands are constructed using two brick layers with a cavity wall in between, an innovation first introduced in the early 1900s, for insulation, health, and comfort purposes (Vekemans, 2016). In the 1970s, large-scale cavity wall insulation programs were introduced for new and existing homes, but with a type of insulation that turned out to last for just 15-20 years, rendering most cavity walls currently empty, or not properly insulated. Interestingly, we observe that homes constructed after 1980 are much less likely to be treated, even though the ”original” insulation in those homes may well have disappeared by now.

For the owner-occupied sample, there seems to be some selection bias in the type of homeowners that are in the treated sample, for instance through their observed higher income. They also have a larger average family size and more female occupants. Sorting into treatment is much less likely for tenants of rental units – the decision to renovate is less likely to be influenced by the tenants in the context of the Dutch housing market, where rent increases are restricted by law, and where more than 60% of the rental market consists of rent-controlled housing. Analyzing the household characteristics of treated observations in the free rental sector, we observe some significant differences with the control group – they have higher wealth and home value. Potentially, this has to do with the type of landlord and the segment in which they operate. As shown by the significantly higher home value in the treatment group, it could be the case that homes in a certain segment have had better maintenance. These concerns are not present in the sample of homes owned by social housing institutions. Here, we observe no significant difference between the characteristics of the inhabitants of treated and control homes. We only observe that homes presumably benefiting most from improved insulation, older homes, and homes with higher gas consumption, are more likely to be insulated during the sample period. These results can be related to the attributes of the social housing segment. For instance, their housing stock consists of a more homogeneous sample, which can be observed in our sample through the fact that there are no detached homes in our sample. Moreover, the standard error of the home value is smaller than in the owner-occupied and private rental market segment. The waiting lists for social housing are long, and assignment to a home takes more than 7 years in a quarter of the municipalities, with extreme outliers in highly urbanized areas. For instance, the top 10 of municipalities with the highest waiting list all have a waiting period of more than 14 years (NOS, 2021). This observation leads to low bargaining power of tenants in social housing to improve insulation, as their options to move are limited.

—Insert Table 2—

3 Methodology

To estimate the causal relationship between insulation improvement in a home and its subsequent gas consumption, we employ a difference-in-difference approach with time and household fixed effect[4]. We compare treated households with those that are never treated during our observation period. Equation 1 provides the empirical model:

$$ln(Gas\ use_{it}) = \beta_0 + \beta_1 Insulation_{it} + X_{it} + \lambda_i + \mu_t + \epsilon_{it} \tag{1}$$

where $ln(Gas\ use_{i,t}$ denotes the logarithm of yearly gas consumption in m^3 of household i in period t. $Insulation_{i,t}$ is a dummy variable that equals one when an observation is in the treatment period – the first year after the insulation year – and has been subject to insulation treatment. X_{it} is a vector of home and household characteristics that can vary over time, which are the household income, the amount of household members, and the surface of the home. λ_i and μ_t are household and year fixed effects, respectively. ϵ_{it} is the

[4]We only observe households in the home they live in when our sample period starts. When they move home, they leave the sample. That means that a household fixed effect also denotes a home fixed effect.

error term, assumed to be independent of treatment and normally distributed, and clustered at the household level.

To estimate the effect of insulation on gas consumption, we are taking into account insulation measures that have been installed at differed moments in time. In such staggered adoption settings, recent literature has uncovered potential identification issues when employing often used Two Way Fixed Effects models (De Chaisemartin and d'Haultfoeuille, 2020; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Goodman-Bacon, 2021; Baker et al., 2022; Roth et al., 2023). This is especially the case when ever-treated units are also used in the control group at some point in the sample, and treatment effects are heterogeneous over time, or across different cohorts. Then, the average estimator can be biased due to the treatment effect containing negative weights for some estimates (De Chaisemartin and d'Haultfoeuille, 2020).

In our setting, we limit this bias by only considering never treated homes as our control group. Additionally, we will provide a comparison of our estimation results with the proposed doubly-robust estimator of Callaway and Sant'Anna (2021), either using the never-treated or the not-yet-treated household as the control group, to show that the results under these alternative estimation methods are similar. We also examine the potential concern of heterogeneous treatment effects by cohorts, given that we aggregate 10 cohorts of homes that are insulated in different years. In Figure A.2 we plot the treatment effect over time for all cohorts separately, and see that patterns are rather consistent between different cohorts such that aggregating the effects can still give a reliable view of the treatment effect.

To get a first sense of the effect of the treatment (i.e. insulation) on subsequent gas consumption, Figure 2 plots the mean gas consumption for the treatment and control group for the 2010 insulation year.[5] Year 0 indicates the year of the insulation treatment. An important assumption for identification is that gas consumption of the treated homes in the sample follows a trend that is parallel to the gas consumption trend of the group of control homes. Indeed, we clearly observe that the two groups have different levels of gas consumption before insulation, where households in the treatment group consume more gas,

[5]Figures for other years are similar but omitted for brevity. Appendix A.2 provides estimation results for each individual year.

but the slope of the pre-trend across both samples is the same. There are some shocks that are visible, which occur simultaneously and in a similar magnitude for the treatment and the control group. These shocks can likely be attributed to weather conditions, such as colder or warmer winters. In the analysis, the comparison between treatment and control group will always be made within the same calendar year, such that factors that are attributed to a specific year are not affecting the estimation. After the installation of insulation, gas consumption drops in the treatment group and the consumption pattern of both groups becomes more similar. Note that there is a slightly downward long-term trend in gas consumption in the control group, perhaps due to unobservable energy-efficiency investments (e.g. new heating system, etc.) or perhaps consistently warmer winters. We also note that the possible presence of treatment in the control group could lead to an underestimation of the true treatment effect – there are multiple insulation companies active in the Netherlands. In the robustness checks, we address this issue by creating control samples that are more restrictive as compared to the general control sample.

—Insert Figure 2—

4 Results

4.1 Main Effects

Table 3 presents the results of the difference-in-difference analysis, where the dependent variable is the logarithm of annual gas consumption[6]. Column (1) does not include any control variables. We document an average treatment effect of 18.5% after the insulation intervention, as compared to the control group of non-treated homes in the same province. Column (2) includes time-varying control variables that could affect gas consumption, namely the dwelling surface, the number of household members, and household income. The treatment effect stays constant, with a decrease of 18.9% in annual gas consumption after

[6] All tables thus present the results from a log-linear estimation. In order to translate these effect sizes to percentage changes, we exponentiate the coefficients.

the application of insulation measures in the home[7].

4.1.1 Addressing Selection Bias

We may overestimate the effects of insulation on gas consumption given the potential selection bias of sorting into the treatment – environmentally-conscious consumers may be more likely to invest in insulation, and may also take other energy-saving measures. We therefore split the sample into owner-occupied homes and tenant-occupied homes, including homes in the free rental market and homes owned by not-for-profit social housing institutions. Presumably, the insulation treatment is exogenous for the sub-sample of rental homes, given that the landlord decides on investments in the energy efficiency of rental homes, while the tenant pays the energy bill (in the Netherlands, a landlord very rarely pays for energy costs when leasing out independent rental units).

We document that gas consumption decreases by 19.6% in the sample of owner-occupied homes. We find a similar effect for homes in the private rental segment, at 20.4%. The effect of 14.7% for social housing is smaller. However, we have to be cautious in interpreting these overall effects of the rental sample, as the number of observations is relatively small and the estimates for later periods get more noisy, especially for the private rental segment. Figure 3B shows the estimated treatment effect per year. Here we observe that the confidence intervals for the estimates in the rental market are relatively large. Especially in later periods, the estimated effect may be based on a relatively small sample of homes in the rental segment being observed for the extended time period, and this estimate may therefore be less reliable. We do observe that for all three segments, the yearly effect size is never statistically

[7] Table A1 provides the results of the same analysis using the doubly-robust estimator proposed by Callaway and Sant'Anna (2021). Slight differences could be the result of the way in which time-varying control variables are accounted for, where the TWFE approach includes time-varying controls, the Callaway and Sant'Anna estimator controls for the pre-treatment values of those control variables. Another reason would be the difference in the way in which control observations are paired with treated observations. Overall, the comparison of estimates show us that the effect is of comparable magnitude, such that we are not concerned about potential bias caused by our estimation method. Another possible concern could be that the never-treated units are not an appropriate counterfactual, as they could also be treated by another company during the sample period. It would also be possible to argue that the never treated and treated units are fundamentally different and gas consumption patterns could look differently because of e.g. different building quality, which could violate the parallel trend assumption. Table A2 provides the results of the same Callaway and Sant'Anna estimator, but takes the not-yet-treated units as the control group. The results remain very similar, such that we conclude that the never-treated units are an appropriate control group in this setting. Table 8 provides further robustness tests regarding our control group.

significantly different from the coefficient in the owner-occupied sample. These stratified estimates provide some comfort that the size of the treatment effect is consistent across owner-occupied and rental homes, and that a possible selection effect among homeowners is not driving our results.

—Insert Table 3—

The data set allows us to identify different types of insulation measures, including basement, wall and roof insulation. We estimate the treatment effects for each of these insulation types separately in Table 4. In Columns 1, 2, and 3 we include homes where just one insulation measure has been installed. Wall insulation yields average savings in gas consumption of 18.9%. For roof and basement insulation we find a slightly smaller effect of 16.3% and 14.2%, respectively. We also examine the interventions where two insulation types have been installed, for each of the different combinations. Columns 4, 5 and 6 of Table 4 provide the results. We find that in all three combinations, combining two types of insulation yields higher gas consumption reduction compared to using one type of insulation. The combination that shows the highest reduction in gas consumption is the combination of basement and roof insulation. However, we have to note that this estimate is based on a sample of only 7 homes.

—Insert Table 4—

4.2 Heterogeneity Analysis

As a first analysis of heterogeneity in the average treatment effect, we stratify the treated sample on different types of dwellings, as well as different income groups, to explore whether the average effect size varies across these groups. Table 5 divides the sample into different dwelling types. We observe effects that are of similar magnitude for corner homes, semi-detached homes, terraced homes, and detached homes – respectively 17.3%, 19.7%, 17.0%., and 18.6% For apartments, we do not observe a significant reduction in gas consumption after insulation treatment, and the point estimate is close to zero as well. Overall these results can be explained by the fact that insulation is most effective for homes that are (semi-)detached, since these homes have a large area of exposed walls, which enhances the effectiveness of insulation. However, we also note that the share of apartments in our sample is relatively small, which decreases the statistical power of the analysis. A larger sample of treated apartments could help in providing more conclusive results on the effect of insulation for this dwelling type.

—Insert Table 5—

We subsequently interact our treatment effect with a dummy variable that indicates whether a household is below or above the median income. We perform this analysis separately for owner-occupied homes, rental homes in the free sector, and homes owned by social housing institutions. Household income levels may affect the impact of energy efficiency improvements on energy consumption through a different baseline consumption level. If income constraints lead to a below-optimal consumption of energy in the baseline case, improvements in energy efficiency standards of the home may have smaller-than-anticipated effects due to a partial increase in energy consumption by the household (Saunders, 2013). Table 6 shows the effect of insulation on gas consumption for low- and high-income households, per housing segment. We find that in the owner-occupied sector, above-median income homeowners reduce their gas consumption less after insulation improvements compared to lower income households. A plausible explanation for this could be related to the age of the home that they live in. The data tells us that the average building age of the homes of below-median income households is 2,5 years below that of above median income households. They are also more likely to live in the oldest homes (with building years 1900-1929). Homes of different building ages can require specific strategies for insulation. The oldest homes might have a lower baseline energy efficiency level, and more to gain from improved insulation.

—Insert Table 6—

4.3 Persistence of Treatment Effect

The average treatment effect is not just an average effect across households, but also an average effect over the treatment period. The sample includes insulation interventions that took place between 2010 and 2019, and we observe post-intervention gas consumption up to a period of ten years after the insulation year. We therefore explore whether the reduction in gas consumption is persistent over time – households could decide to spend part of the energy savings on heating their home more, a rebound effect, or there could be other mechanisms for non-persistence of the documented effects, such as decay in the quality of insulation.

Figure 3A plots the coefficients of the difference-in-difference analysis as in Equation (1). Each dot represents the difference in gas consumption between the treatment and the control group, relative to the year before insulation, as well as the 95% confidence interval. The figure shows that the difference between treatment and control groups is stable before insulation. After installing the insulation there is a sharp drop in gas consumption of about 20%. Over time, the confidence interval widens. This is related to the fact that we have fewer observations at the start of the sample period, which subsequently leads to fewer observations with a long time span (Appendix Figure A.2 shows the results for each individual year). Moreover, there can be households in the sample that move, such that the observed time period is shorter for these observations. However, we observe that the estimated size of the treatment effect is quite consistent over time. As opposed to the attenuating effect of behavioral treatments, such as the Opower social comparison-based treatment (Allcott and Rogers, 2014), the effect of insulation is a structural change in the home that remains as time progresses. The persistence in energy savings following the intervention also provides some indication on the absence of a rebound effect, or the absence of other underlying mechanisms that could lead to non-persistence.

Similar to Figure 3A, Figure 3B plots the coefficient estimates of gas consumption over time, separately for owner-occupied homes and rental homes. Since the first three years of the sample period include insulation interventions in owner-occupied homes only, the analysis includes just seven years after insulation for private rental homes, and seven years for social housing. The figure shows that, over time, there is no significant difference in the reduction of gas consumption across owner-occupied and rental homes, with merely a widening confidence interval (likely due to fewer observations early in the sample period). Reductions in gas consumption are therefore largely persistent across types of homes, notwithstanding the tenure choice.

—Insert Figure 3—

4.4 Robustness Checks

4.4.1 Effect of Insulation on Electricity Consumption

The decrease in gas consumption following an insulation retrofit could also be due to a replacement effect. This could be the case if, for instance, households that insulate their home simultaneoulsy install a heat pump, and subsequently heat their home using electricity instead of natural gas. In such cases, we would overestimate total energy savings at the household level by just considering the effect of insulation on the consumption of natural gas. Therefore, we repeat the same analysis as our baseline difference-in-difference model, but replace gas consumption as the dependent variable by electricity consumption. Using this estimation strategy, we can observe whether natural gas consumption in homes with improved insulation is replaced by an increase in electricity consumption.

The results of the analysis are presented in Table 7. We document that, after controlling for home and household characteristics, there is an overall negative effect on electricity consumption. We therefore conclude that there is no replacement effect, but that implementing insulation is associated with a a *reduction* in home electricity consumption, rather than an increase. This reduction in electricity consumption might point to a within-household spillover effect of energy efficiency measures, consistent with recent findings of positive behavioral spillovers in the broader environmental domain (see, for example, Alacevich et al. (2021) and Jessoe et al. (2021)).

In addition, we analyze the distribution of savings in gas consumption through a non-parametric estimate, as shown in Appendix Figure A.1(a). Here we subtract the average gas consumption in the period after insulation, which is maximally 10 years, from the average gas consumption up to 5 years before insulation. We exclude the insulation year itself from this estimate. We observe that just a very small group of households realizes a 100% reduction in gas consumption after the insulation intervention. This finding suggests that completely halting the use of natural gas, for example through switching to electricity as the main energy source, is typically not observed in our sample. (We also note that the use of electric heat pumps was relatively rare during the sample period.)

—Insert Table 7—

4.4.2 Restricting the Control Group

In the baseline model, the control group consists of a random 1% of all homes in the same province. In this control group, there could also be homes in which insulation was improved during the sample period. If that is the case, our baseline estimate would underestimate the true treatment effect. In Table 8, we restrict the control group in a variety of ways, such that we can be more certain that insulation improvements in the control group are not influencing our results. In Column 1, we first exclude households where gas consumption decreased by more than the median change in gas consumption in the treatment sample (31.13%) between any two consecutive years during the treatment period. In these homes, we expect that the energy efficiency has been changed through a) insulation installed by a different provider, or b) different energy efficiency measures. In Column 2, we only include homes in the control sample that were constructed after the year 2000. In this case, we can be rather sure that there are no changes to the home insulation, as the existing insulation is of high quality due to building regulations that became effective in that year. In column 3, the sample is restricted in terms of geographical area. Rather than considering the full province of Limburg (some 1.1 million inhabitants), we consider just the city of Maastricht (some 120,000 inhabitants), where the company has its largest clientele. In this case, there is a lower likelihood that homes in the control group have improved insulation through another company. We observe that in Columns 1 and 2, the effect size increases as compared to the baseline model in Table 3, with an effect size of 27.1% for both analyses. This would provide an indication that our baseline estimates might underestimate the true treatment effect. In Column 3, the effect size is very similar to our baseline estimation. Ruling out unobserved treatment in the control group is likely not solved by applying a geographical restriction.

In Column 4, 5 and 6 we utilize information from the Energy Performance Certificate of a home, when this information is available. Firstly, in Column 4 we drop observations from the control group where the energy label of the house is improved during our sample period. Secondly, we drop households from the control group where the reported quality of the insulation is poor, which increases the likelihood that it would be upgraded during the sample period. Lastly, we remove homes from the treatment group where we know that the window quality is low, such that the chance that the windows are being replaced along with improving the insulation is limited. The coefficients in Columns 4, 5 and 6 are very similar to our main result in Table 3. Overall, the results in Table 8 suggest that we could underestimate the treatment effect, but also give us more confidence that the true effect size is of a similar magnitude as our baseline estimates.

—Insert Table 8—

5 Financial Considerations

Home insulation has a sizeable effect on household gas consumption. The question remains what the reduction in energy consumption implies for private individuals financially, as they face an upfront financial outlay to improve the energy efficiency of their home[8]. In the results section, we estimated the average treatment effect on the treated per insulation type. We use these estimates to perform a back-of-the-envelope calculation on the financial returns to different insulation types, exploiting invoice data of the insulation company to calculate the average investment costs in our sample. Importantly, this simple calculation ignores the possible presence of subsidies[9]. The possibility of subsidies implies that our return calculations are lower-bound estimates.

On the benefit side, we assume perpetuity of energy savings to calculate the return. While homes may be sold at some point, it is reasonable to assume the capitalization of energy efficiency into home prices (see, for example, Aydin et al. (2020)). In estimating yearly savings, we consider the gas consumption and gas prices in the year before the insulation was installed. These prices are inflation-adjusted to the year 2019.

[8]For landlords of rental homes, the return calculation is complicated due to tenants directly benefiting from the investment in energy efficiency by the landlord.

[9]Indeed, over the past decade, the Netherlands had a variety of subsidy programs to stimulate energy efficiency, for example for solar PV. At the time of writing, there was a government subsidy in place for home insulation measures, which required at least two forms of insulation. The level of the subsidy was at about 30% of the initial investment. See https://www.milieucentraal.nl/energie-besparen/isoleren-en-besparen.

Table 9 displays investment costs, yearly savings, annual return, and the payback period[10]. In Column 1, we consider all insulation types in the sample (including treatments with multiple measures), whereas in columns 2, 3, and 4 we only consider homes where just one type of insulation was used. The results show an annual average return of 19.9% given investment costs and energy prices prevailing at the moment of insulation, which corresponds to a payback period of 5 years. We observe that annual returns from wall insulation are particularly high, with an average of 21.8%. For basement and roof insulation, the annual return is 14.9% and 11.8%, respectively[11]. Considering the payback period, the average wall insulation investment of €1,656 will be earned back in about 4.6 years. For basement and roof insulation, the average payback period is 6.7 and 8.5 years, respectively. Considering that a household lives in a dwelling for around 10 years, all insulation types would be earned back within this period.

Of course, the consideration of energy efficiency measures hinges on more than financial returns alone. Upfront capital outlays (no matter the relatively small size of the investment), the "hassle" factor, energy illiteracy (Brounen et al., 2013) and the perceived risks of home insulation (e.g. an increase in the likelihood of mold) are all barriers that hold back private consumers from improving the energy efficiency of their homes. For landlords, an important (albeit solvable) consideration is the split incentive, where tenants reap the benefits of landlord-driven improvements in energy efficiency. Finally, an important but often ignored issue is the presence of supply-side constraints for energy efficiency improvements. Many of these measures are highly labor-intensive, and jobs can be hard to fill. Equally, more advanced energy efficiency improvements (e.g. heat pumps) require components that are in scarce supply, leading to long waiting times. Given the efficacy of investments in home energy efficiency, policies addressing supply-side issues, for example through workforce training, or targeted visa waivers, may help to more quickly improve the efficiency of the buildings stock, helping to reduce both energy dependence and global carbon emissions.

[10]The return is calculated by the formula: Yearly savings gas bill / Investment insulation * 100%. We assume perpetual savings, and keep the prices of energy constant at the rate of the moment of insulation.

[11]Appendix Figure A.1(c) displays the distribution of non-parametricly estimated annual returns in the sample. Here it becomes visible that there are more extreme cases present in the sample in terms of positive as well as negative annual returns.

—Insert Table 9—

6 Conclusion

Improving the energy efficiency of the building stock is important in decreasing household energy consumption and reduce the negative externality from carbon emissions. In addition, home energy efficiency may shield household budgets from negative price shocks such as those experienced by European consumers in 2022. The baseline measures to enhance the energy efficiency of a home are wall, roof, and basement insulation – other energy efficiency measures are dependent on the presence of insulation in the home. For instance, low-temperature heat pumps, which are slated to replace gas-fired furnaces and boilers in many parts of the world, are currently suitable just for homes that can be heated using low-temperature water – for that, proper insulation is needed.[12]. Using unique, hand-collected data on home insulation measures, this study examines the effect of roof, wall and basement insulation on gas consumption in a large sample of (rental and owner-occupied) residential homes.

The results of the difference-in-difference analysis show that home insulation measures significantly reduce gas consumption, with an average treatment effect of about 19%. We test for heterogeneous effects across types of homes, and across household characteristics. Not surprisingly, homes with the largest fraction of exposed walls (e.g. detached and semi-detached homes) benefit most from home insulation, while household income does not influence the yield from insulation significantly. Furthermore, we investigate long-run gas consumption for up to ten years after the energy efficiency improvements, to address potential concerns of persistence. Importantly, the point estimates remain stable in the long run, which provides some indication that the gas use reduction can be attributed to the changed physical characteristics of the home, rather than behavioral changes of the household.

Translating our findings to financial savings, we observe an average reduction in the energy bill of €350 per year. Compared to the investment to install insulation, this yield an annual return of 19.9%, translating into a payback period of 5 years. Wall insulation has the highest return, of 21.8%, while basement insulation returns 14.9% and roof insulation

[12]In fact, Austria will ban gas-fired heating boilers per 2023, and in the Netherlands, gas-fired boilers can be replaced just by heat pumps as of 2026.

returns 11.8% per year.

Of course, the generalizability of empirical findings hinges on the context, where the study at hand applies just to homes with cavity walls, in a climate with colder winters and limited need for air-conditioning in summer. With that caveat in mind, the results in this paper can inform homeowners, investors, and social housing institutions in their home retrofitting decisions, reducing investment uncertainty. First, the information in this paper can be used to get more realistic expectations of the energy savings from insulation. Second, given the financial rates of return documented in this paper, cavity wall insulation (as well as some other forms of home insulation) seems a prime example of a cost-effective carbon abatement strategy. Targeting homes with cavity walls could therefore increase the cost-effectiveness of large-scale weatherization programs (Christensen et al., 2023). Of course, this restricts the applicability of the findings in this study, given the way homes are constructed in, for example, the United States. Third, there seems to be limited necessity for subsidy programs aimed at stimulating home insulation. Targeted policy aimed at households who do not have the financial means to make the upfront insulation investment, preferably in the form of a loan, could be more suitable.

References

Alacevich, C., Bonev, P., and Söderberg, M. (2021). Pro-environmental interventions and behavioral spillovers: Evidence from organic waste sorting in Sweden. *Journal of Environmental Economics and Management*, 108:102470.

Allcott, H. and Greenstone, M. (2012). Is there an energy efficiency gap? *Journal of Economic Perspectives*, 26(1):3–28.

Allcott, H. and Greenstone, M. (2017). Measuring the welfare effects of residential energy efficiency programs. Technical report, National Bureau of Economic Research.

Allcott, H. and Mullainathan, S. (2010). Behavior and energy policy. *Science*, 327(5970):1204–1205.

Allcott, H. and Rogers, T. (2014). The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*, 104(10):3003–37.

Aydin, E., Brounen, D., and Kok, N. (2018). Information provision and energy consumption: Evidence from a field experiment. *Energy Economics*, 71:403–410.

Aydin, E., Brounen, D., and Kok, N. (2020). The capitalization of energy efficiency: Evidence from the housing market. *Journal of Urban Economics*, 117:103243.

Aydin, E., Kok, N., and Brounen, D. (2017). Energy efficiency and household behavior: the rebound effect in the residential sector. *The RAND Journal of Economics*, 48(3):749–782.

Baker, A. C., Larcker, D. F., and Wang, C. C. (2022). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*, 144(2):370–395.

Brounen, D., Kok, N., and Quigley, J. M. (2013). Energy literacy, awareness, and conservation behavior of residential households. *Energy Economics*, 38:42–50.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Christensen, P., Francisco, P., Myers, E., and Souza, M. (2023). Decomposing the wedge between projected and realized returns in energy efficiency programs. *Review of Economics and Statistics*, 105(4):798–817.

De Chaisemartin, C. and d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996.

Eurostat (2020). Energy consumption in households. `https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Energy_consumption_in_households`. Accessed: 10-02-2023.

Fowlie, M., Greenstone, M., and Wolfram, C. (2018). Do energy efficiency investments deliver? evidence from the weatherization assistance program. *The Quarterly Journal of Economics*, 133(3):1597–1644.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Granade, H. C., Creyts, J., Derkach, A., Farese, P., Nyquist, S., and Ostrowski, K. (2009). Unlocking energy efficiency in the US economy. *McKinsey & Company*.

Hong, S. H., Oreszczyn, T., Ridley, I., Group, W. F. S., et al. (2006). The impact of energy efficient refurbishment on the space heating fuel consumption in English dwellings. *Energy and Buildings*, 38(10):1171–1181.

Jaffe, A. B. and Stavins, R. N. (1994). The energy-efficiency gap: What does it mean? *Energy Policy*, 22(10):804–810.

Jessoe, K., Lade, G. E., Loge, F., and Spang, E. (2021). Spillovers from behavioral interventions: Experimental evidence from water and energy use. *Journal of the Association of Environmental and Resource Economists*, 8(2):315–346.

Liang, J., Qiu, Y., James, T., Ruddell, B. L., Dalrymple, M., Earl, S., and Castelazo, A. (2018). Do energy retrofits work? Evidence from commercial and residential buildings in Phoenix. *Journal of Environmental Economics and Management*, 92:726–743.

Magrini, A., Marenco, L., Leoni, V., and Gamba, R. (2022). Air cavity building walls: A discussion on the opportunity of filling insulation to support energy performance improvement strategies. *Energies*, 15(23):8916.

Metcalf, G. E. and Hassett, K. A. (1999). Measuring the energy savings from home improvement investments: evidence from monthly billing data. *Review of Economics and Statistics*, 81(3):516–528.

NOS (2021). Sociale huurwoning? in zeker een kwart van de gemeenten wacht je meer dan 7 jaar. `https://nos.nl/op3/artikel/2377995-sociale-huurwoning-in-zeker-een-kwart-van-de-gemeenten-wacht-je-meer-dan-7-jaar`. Accessed: 16-06-2023.

Palacios, J., Eichholtz, P., Kok, N., and Aydin, E. (2021). The impact of housing conditions on health outcomes. *Real Estate Economics*, 49(4):1172–1200.

Peñasco, C. and Anadón, L. D. (2023). Assessing the effectiveness of energy efficiency measures in the residential sector gas consumption through dynamic treatment effects: Evidence from England and Wales. *Energy Economics*, 117:106435.

Roth, J., Sant'Anna, P. H., Bilinski, A., and Poe, J. (2023). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2):2218–2244.

Saunders, H. D. (2013). Historical evidence for energy efficiency rebound in 30 US sectors and a toolkit for rebound analysts. *Technological Forecasting and Social Change*, 80(7):1317–1330.

Schweitzer, M. (2005). Estimating the national effects of the US department of energy's weatherization assistance program with state-level data: A metaevaluation using studies from 1993 to 2005. *US Department of Energy*.

Sorrell, S., Dimitropoulos, J., and Sommerville, M. (2009). Empirical estimates of the direct rebound effect: A review. *Energy Policy*, 37(4):1356–1371.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Vekemans, B. (2016). *Legislation and Brickwork Facades in a Historical Perspective*.

Table 1: Descriptive Statistics

	(1) Owner-Occupied	(2) Rental Private	(3) Rental Social
# of insulation measures	1.121	1.169	1.195
	(0.344)	(0.422)	(0.537)
Wall			
Percentage	0.878	0.831	0.654
	(0.327)	(0.378)	(0.477)
Total cost in €	1633.168	1891.067	970.886
	(650.4)	(2260.2)	(436.7)
Surface in m^2	106.701	125.229	64.734
	(45.23)	(157.4)	(32.74)
Basement			
Percentage	0.203	0.305	0.270
	(0.402)	(0.464)	(0.445)
Total cost in €	1355.773	1301.312	1189.812
	(504.0)	(411.2)	(628.9)
Surface in m^2	53.125	53.250	50.170
	(20.74)	(16.25)	(29.30)
Roof			
Percentage	0.034	0.034	0.103
	(0.180)	(0.183)	(0.304)
Total cost in €	1936.847	2047.000	2021.579
	(1296.5)	(1588.2)	(769.5)
Surface in m^2	61.593	71.500	52.053
	(22.47)	(54.45)	(14.48)
Other			
Percentage	0.006	.	0.168
	(0.0752)	.	(0.374)
Total cost in €	1750.482	.	2092.200
	(680.5)	.	(1633.6)
Surface in m^2	62.667	.	40.000
	(34.15)	.	(29.12)
Observations	1,806	218	216

Notes: Table 1 presents the insulation characteristics per type of insulation, separately for the full sample, owner-occupied dwellings and rental dwellings, where rental is reported separately for social housing and private homes. The "percentage" reports what share of households in the particular column installed that type of insulation. Standard errors are reported in parenthesis.

Table 2: Descriptive Statistics

	(1) Treatment Owner-Occupied	(2) Control Owner-Occupied		(3) Treatment Rental Private	(4) Control Rental Private		(5) Treatment Rental Social	(6) Control Rental Social	
Energy consumption									
Annual gas consumption	2384.549	2066.610	***	2435.060	1768.649	***	1567.436	1384.680	***
in m^3	(744.0)	(793.3)		(756.0)	(828.0)		(599.8)	(584.5)	
Annual electricity	3733.500	3613.118	*	3133.832	2746.543	**	2556.595	2432.949	
consumption in kWh	(1332.6)	(1400.0)		(1323.5)	(1214.8)		(1286.4)	(1236.1)	
Household characteristics									
# of household	2.288	2.162	**	1.790	1.615		1.756	1.662	
members	(1.053)	(1.076)		(0.917)	(0.809)		(0.960)	(0.922)	
# of children	0.395	0.341		0.124	0.096		0.198	0.190	
	(0.816)	(0.760)		(0.454)	(0.467)		(0.569)	(0.583)	
# of elderly (>65)	0.678	0.704		1.105	0.927		0.692	0.763	
	(0.909)	(0.931)		(0.940)	(0.865)		(0.854)	(0.854)	
# of females	1.041	0.962	**	0.924	0.773	*	0.895	0.816	
	(0.747)	(0.721)		(0.583)	(0.554)		(0.552)	(0.653)	
Household wealth	219.993	218.374		273.526	154.606	***	14.889	17.781	
(x €1000)	(181.8)	(181.4)		(237.5)	(223.2)		(24.53)	(40.51)	
Annual household	38.211	35.121	***	30.752	27.026	*	22.152	21.660	
income (x €1000)	(15.79)	(15.43)		(14.16)	(12.26)		(9.201)	(8.418)	
Dwelling characteristics									
Home value (x €1000)	217.599	221.538		218.533	186.225	**	125.192	125.171	
	(78.06)	(91.26)		(62.90)	(93.35)		(29.19)	(33.41)	
Dwelling surface in m^2	154.340	145.313	***	152.295	123.558	***	92.860	91.173	
	(47.20)	(55.68)		(42.73)	(57.39)		(20.70)	(22.73)	
Dwelling type									
Apartment	0.007	0.080	***	0.029	0.408	***	0.360	0.428	
	(0.0838)	(0.271)		(0.167)	(0.492)		(0.482)	(0.495)	
Corner	0.184	0.152	*	0.200	0.092	**	0.238	0.187	
	(0.387)	(0.359)		(0.402)	(0.290)		(0.427)	(0.390)	
Semi-detached	0.334	0.212	***	0.257	0.123	**	0.157	0.071	***
	(0.472)	(0.409)		(0.439)	(0.329)		(0.365)	(0.256)	
Row	0.252	0.331	***	0.257	0.262		0.244	0.314	
	(0.434)	(0.471)		(0.439)	(0.440)		(0.431)	(0.464)	
Detached	0.224	0.225		0.257	0.115	***	0.000	0.000	
	(0.417)	(0.418)		(0.439)	(0.320)		(0)	(0)	
Building period									
1900-1929	0.033	0.101	***	0.049	0.095		0.006	0.034	*
	(0.178)	(0.302)		(0.216)	(0.294)		(0.0765)	(0.183)	
1930-1944	0.074	0.059		0.010	0.063	*	0.012	0.017	
	(0.261)	(0.235)		(0.0985)	(0.244)		(0.108)	(0.130)	
1945-1959	0.240	0.125	***	0.194	0.119		0.187	0.184	
	(0.427)	(0.331)		(0.397)	(0.324)		(0.391)	(0.388)	
1960-1969	0.282	0.177	***	0.388	0.198	***	0.503	0.233	***
	(0.450)	(0.382)		(0.490)	(0.400)		(0.501)	(0.423)	
1970-1979	0.336	0.230	***	0.311	0.222		0.251	0.183	*
	(0.472)	(0.421)		(0.465)	(0.417)		(0.435)	(0.387)	
1980-1989	0.028	0.153	***	0.019	0.127	**	0.029	0.227	***
	(0.164)	(0.360)		(0.139)	(0.334)		(0.169)	(0.419)	
1990-1999	0.006	0.130	***	0.029	0.123	**	0.012	0.108	***
	(0.0786)	(0.337)		(0.169)	(0.329)		(0.108)	(0.311)	
>2000	0.003	0.024	***	0.000	0.052	*	0.000	0.013	
	(0.0516)	(0.154)		(0)	(0.222)		(0)	(0.115)	
Observations	1249	1829	3078	161	398	559	177	1072	

Notes: Table 2 presents the descriptive statistics. The control group consists of a random selection of 1% of households in the same region. The table splits between owner-occupied homes and rental homes. The table displays the statistics for the year 2009, before any of the households in the treatment group installed insulation. Standard errors are reported in parenthesis. * P<0.05. ** P<0.01. *** P<0.001

Figure 1: Insulation Measures Over Time

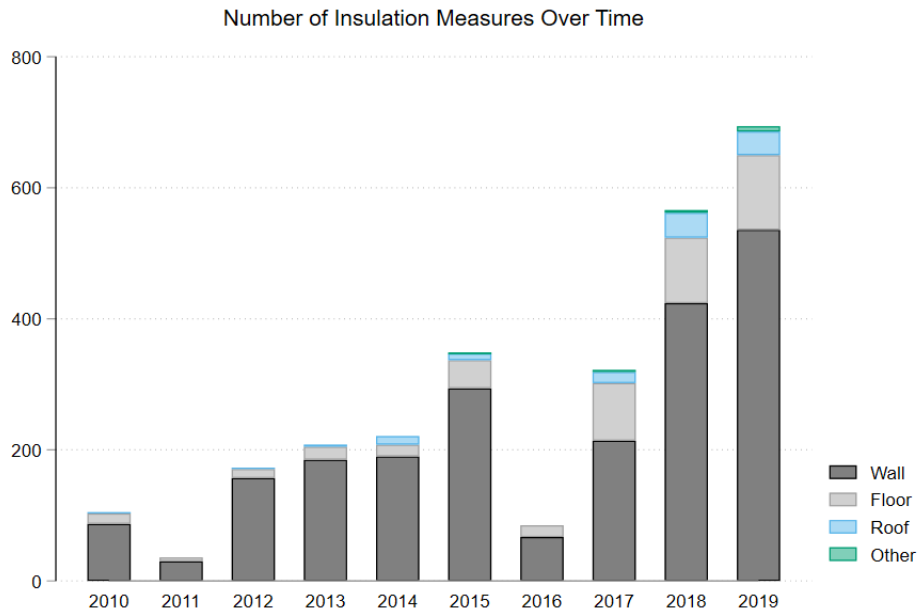

Notes: Figure 1 presents the number of recorded insulation measures in our sample over the sample period, split up per type of insulation.

Figure 2: Gas Consumption in Treated versus Non-Treated Homes

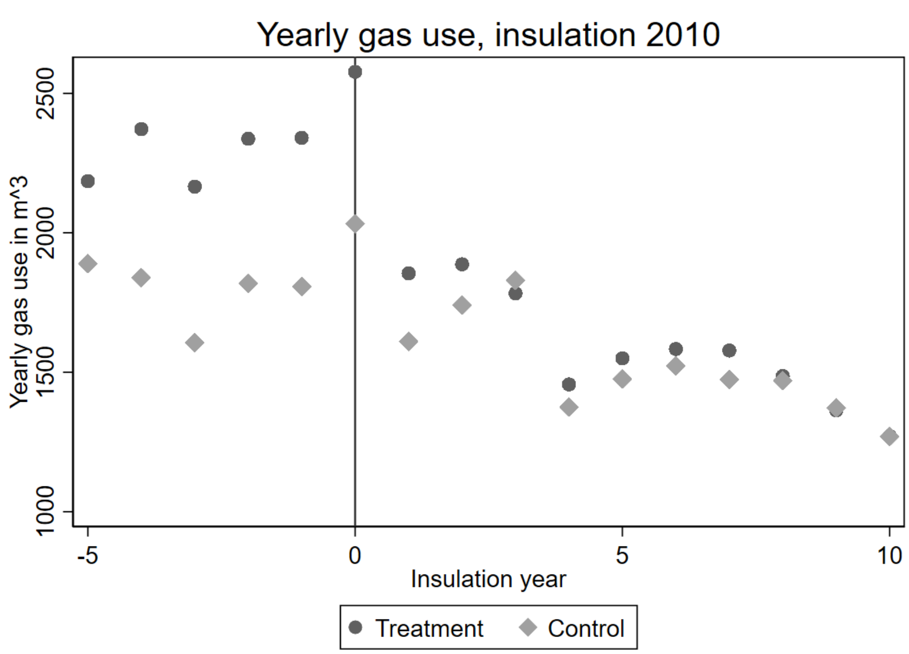

Notes: Figure 2 presents the mean of yearly gas use in the treatment and control group. Year 0 is the year of insulation.

Table 3: Insulation and Gas Consumption

	(1) Full Sample	(2) Full Sample	(3) Owner- Occupied	(4) Rental Private	(5) Rental Social
Insulation	-0.204***	-0.209***	-0.218***	-0.228***	-0.159***
	(0.0113)	(0.0111)	(0.0115)	(0.0375)	(0.0400)
Observations	78,090	77,970	57,721	8,980	18,822
Number of treated households	2,023	1,986	1,702	211	216
Year FE	YES	YES	YES	YES	YES
Household FE	YES	YES	YES	YES	YES
Controls	NO	YES	YES	YES	YES

Notes: Dependent variable: log annual gas consumption. Standard errors are clustered at the household level. Standard errors are reported in parenthesis. * P<0.05. ** P<0.01. *** P<0.001

Table 4: Heterogeneity by Insulation Type

	(1) Wall	(2) Basement	(3) Roof	(4) Wall & Basement	(5) Wall & Roof	(6) Basement & Roof
Insulation	-0.210***	-0.153***	-0.178***	-0.270***	-0.235***	-0.411***
	(0.0124)	(0.0252)	(0.0672)	(0.0318)	(0.0925)	(0.0782)
Observations	69,154	54,222	52,027	56,465	54,494	54,343
Number of treated households	1,482	215	39	170	22	7
Year FE	YES	YES	YES	YES	YES	YES
Household FE	YES	YES	YES	YES	YES	YES
Controls	YES	YES	YES	YES	YES	YES

Notes: Dependent variable: log annual gas consumption. Columns 1, 2, and 3 only include households where one insulation measure is installed. Columns 4, 5, and 6 only include households where only two insulation measures are installed. Households where more than two insulation are installed are excluded from the table, since the sample only has 4 of these observations. Standard errors are clustered at the household level. Standard errors are reported in parenthesis. * P<0.05. ** P<0.01. *** P<0.001

Table 5: Heterogeneity by Dwelling Type

	(1) Apartment	(2) Corner	(3) Semi-detached	(4) Between	(5) Detached
Insulation	0.0208	-0.190***	-0.220***	-0.186***	-0.207***
	(0.0782)	(0.0258)	(0.0191)	(0.0209)	(0.0198)
Observations	10,221	14,116	16,732	23,911	12,395
Number of treated households	42	307	492	379	320
Year FE	YES	YES	YES	YES	YES
Household FE	YES	YES	YES	YES	YES
Controls	YES	YES	YES	YES	YES

Notes: Dependent variable: log annual gas consumption. Standard errors are clustered at the household level. Standard errors are reported in parenthesis. * P<0.05. ** P<0.01. *** P<0.001

Table 6: Heterogeneity by Income Levels

	(1) Owner-Occupied	(2) Rental Private	(3) Rental Social
Insulation	-0.240***	-0.231***	-0.154***
	(0.0171)	(0.0487)	(0.0613)
Insulation* High Income (>50%)	0.0424**	0.000475	-0.0117
	(0.0198)	(0.0551)	(0.0677)
Observations	57,721	8,980	18,822
Number of treated households	466	1236	33
Year FE	YES	YES	YES
Household FE	YES	YES	YES
Controls	YES	YES	YES

Notes: Dependent variable: log annual gas consumption. Standard errors are clustered at the household level. The table presents the results of a triple difference regression, where a dummy for Treatment (Insulation), a dummy for Post (After Insulation), and a dummy for households with above median income are interacted. Standard errors are reported in parenthesis. * P<0.05. ** P<0.01. *** P<0.001.

Figure 3: Insulation Effect Over Time

(a) Full Sample

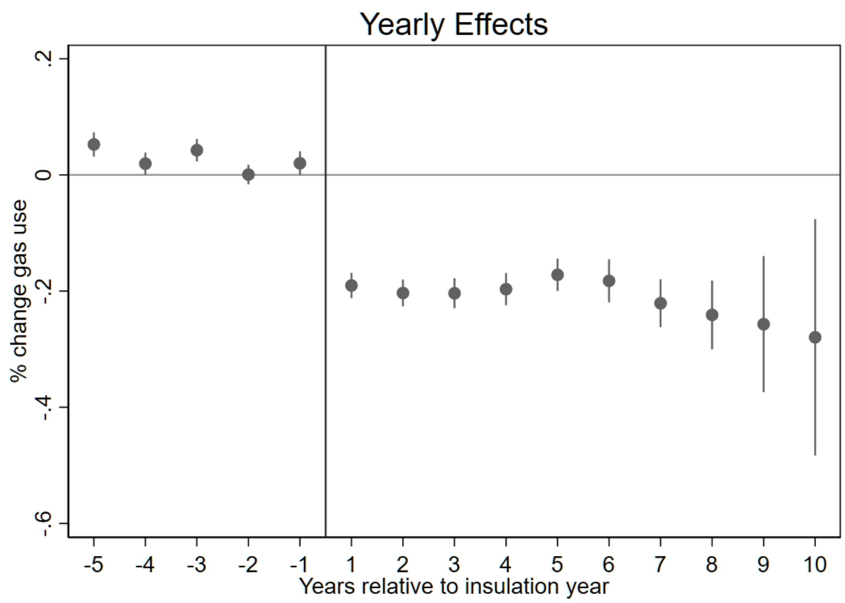

(b) Owner-occupied and Rental Homes

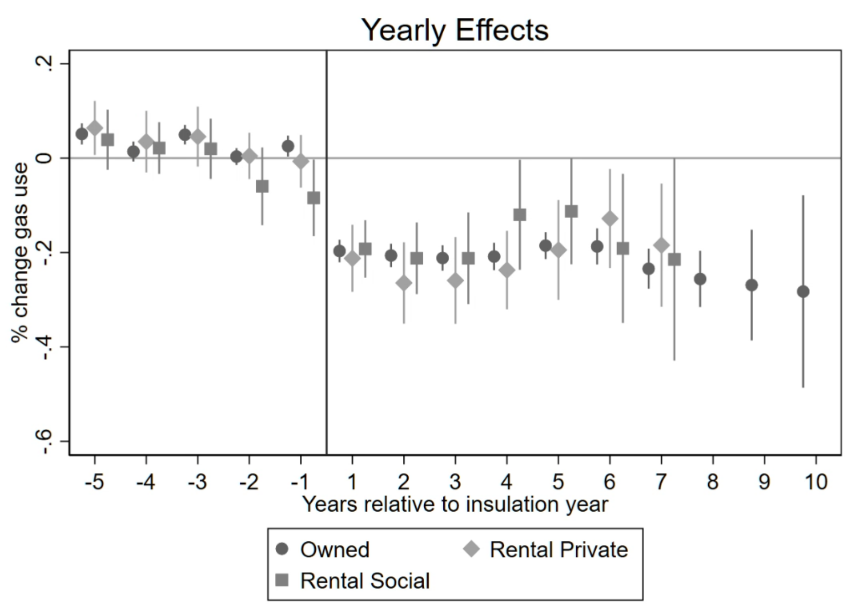

Notes: Figure displays annual gas consumption relative to year 0, the last year before insulation. The figure shows the point estimates, with their 95% confidence interval.

Table 7: Replacement Effects: Insulation and Electricity Consumption

	(1) Full Sample	(2) Full Sample	(3) Owner-Occupied	(4) Rental Private	(5) Rental Social
Insulation	-0.0234*	-0.0492***	-0.0349***	-0.0375	-0.0623*
	(0.0123)	(0.0104)	(0.0110)	(0.0360)	(0.0366)
Observations	76,415	76,301	56,365	8,797	18,566
Number of treated households	2,023	1,986	1,702	211	216
Year FE	YES	YES	YES	YES	YES
Household FE	YES	YES	YES	YES	YES
Controls	NO	YES	YES	YES	YES

Notes: Dependent variable: log annual electricity consumption. Standard errors are clustered at the household level. Standard errors are reported in parenthesis. * P<0.05. ** P<0.01. *** P<0.001

Table 8: Robustness Checks: Restricting the Sample

	(1) Control: Gas Consumption Change	(2) Control: Building year >2000	(3) Treatment & Control: Same City	(4) Control: Label Upgrade	(5) Control: Improved Insulation	(6) Treatment: Window Quality
Insulation	-0.316***	-0.316***	-0.215***	-0.216***	-0.216***	-0.208***
	(0.0106)	(0.0298)	(0.0231)	(0.0112)	(0.0114)	(0.0113)
Observations	55,303	28,308	15,727	72,899	69,850	76,091
Number of treated households	1,986	1,986	653	1,986	1,986	509
Year FE	YES	YES	YES	YES	YES	YES
Household FE	YES	YES	YES	YES	YES	YES
Controls	YES	YES	YES	YES	YES	YES

Notes: In Table 8 shows we restrict the sample in different ways. In column 1, we remove households from the control group where gas consumption dropped with more that the median gas consumption reduction in the treatment group (31.13%). In column 2, we only include homes built after 2000 in the control group. Column 3 only includes homes that are located in the city of Maastricht, both in the treatment and the control group. Dependent variable: log annual gas consumption. Standard errors are clustered at the household level. Column 4 exludes homes in the control group where the energy label was improved during the sample period. Column 5 removes households from the control group where a poor insulation quality was reported. Lastly, Column 6 does not include homes in the treatment group where a poor window quality is reported. Standard errors are reported in parenthesis. * P<0.05. ** P<0.01. *** P<0.001

Table 9: Returns to Insulation Measures

	Actual prices			
	(1)	(2)	(3)	(4)
	All	Wall	Basement	Roof
Yearly savings	€350	€361	€211	€278
Investment	€1,759	€1,656	€1,416	€2,359
Annual return	19.9%	21.8%	14.9%	11.8%
Payback period	5.0	4.6	6.7	8.5
Number of treated households	2,023	1,507	219	40

Notes: Table 9 displays average yearly savings and investment costs. Savings are calculated based on the average estimated effect size per insulation measure. The investment costs are obtained from the invoices of the insulation company. In column 1 to 4, we multiply the average savings by the gas price in the year before installation.

A Appendix

Table A1: Insulation and Gas Consumption: Callaway and Sant'Anna DiD Estimator: Never Treated Control Group

	(1) Full Sample	(2) Full Sample	(3) Owner- Occupied	(4) Rental Private	(5) Rental Social
Insulation	-0.179***	-0.181***	-0.185***	-0.161***	-0.124***
	(0.00767)	(0.00764)	(0.00805)	(0.0253)	(0.0284)
Observations	79,026	77,842	57,298	8,946	19,071
Number of treated households	2,023	1,986	1,702	211	216
Controls	NO	YES	YES	YES	YES

Notes: Dependent variable: log annual gas consumption. Standard errors are clustered at the household level. Standard errors are reported in parenthesis. * P<0.05. ** P<0.01. *** P<0.001. The results in this table are estimated by using the Difference-in-Difference doubly-robust estimator proposed by Callaway and Sant'Anna (2021). Column 2 to 5 control for pre-treatment values of the control variables. The control group consists only of households that are not identified as treated in our sample period.

Table A2: Insulation and Gas Consumption: Callaway and Sant'Anna DiD Estimator: Not Yet Treated Control Group

	(1) Full Sample	(2) Full Sample	(3) Owner- Occupied	(4) Rental Private	(5) Rental Social
Insulation	-0.176***	-0.178***	-0.182***	-0.157***	-0.123***
	(0.00757)	(0.00755)	(0.00795)	(0.0250)	(0.0284)
Observations	79,390	78,233	57,624	9,011	19,121
Number of treated households	2,023	1,986	1,702	211	216
Controls	NO	YES	YES	YES	YES

Notes: Dependent variable: log annual gas consumption. Standard errors are clustered at the household level. Standard errors are reported in parenthesis. * P<0.05. ** P<0.01. *** P<0.001. The results in this table are estimated by using the Difference-in-Difference doubly-robust estimator proposed by Callaway and Sant'Anna (2021). Column 2 to 5 control for pre-treatment values of the control variables. The control group consists of households that are not yet treated in our sample period.

Figure A.1: Insulation Effect Over Time

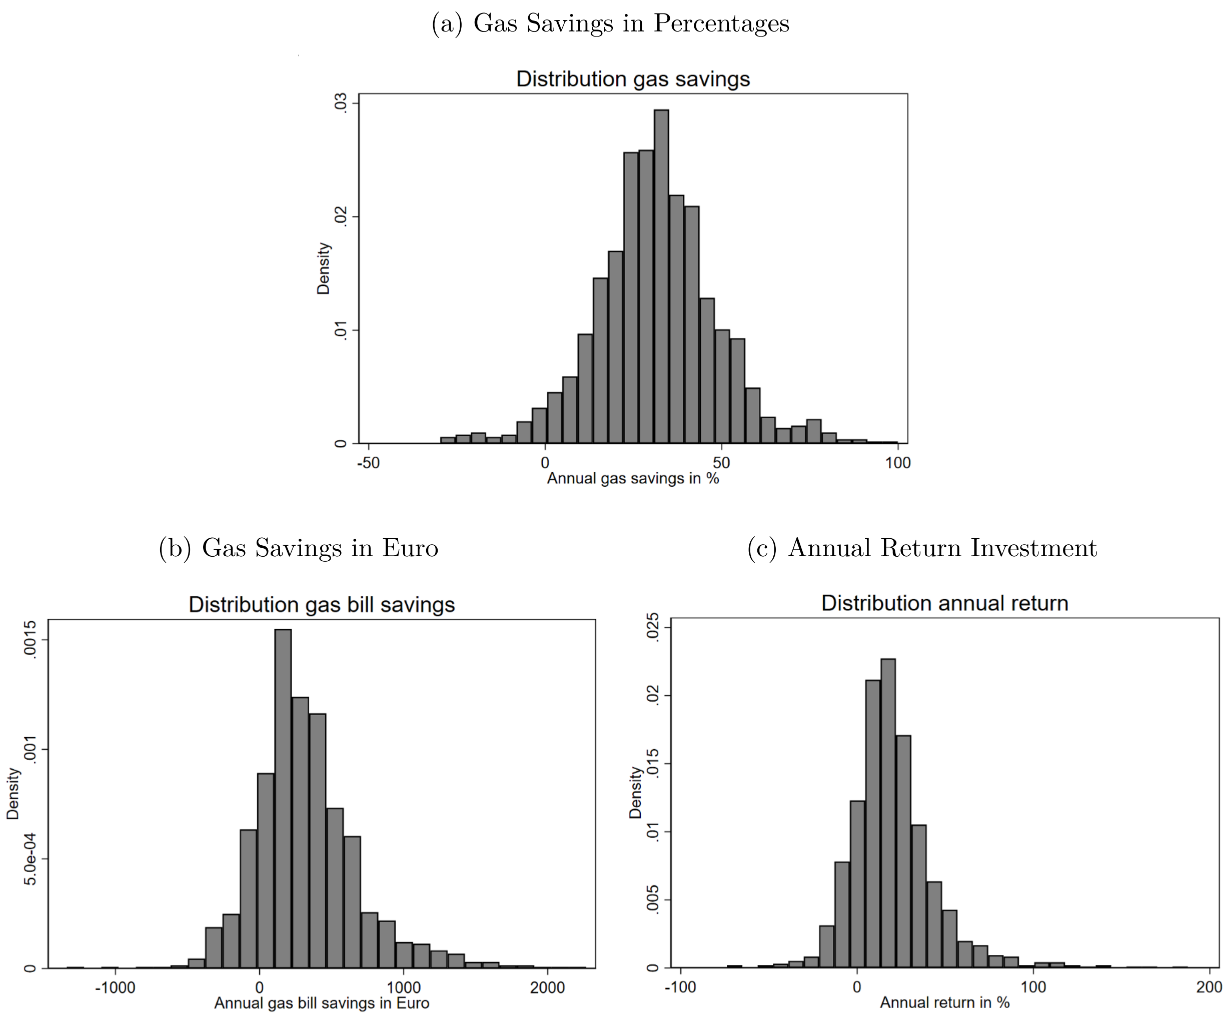

Notes: Figure A.1 displays the non-parametric distribution of gas savings based on the 5 years before, until maximally 10 years after improved insulation. Gas use in the year of insulation is not included in this calculation. Firstly, A.1(a) shows the annual gas savings in percentage terms. A.1(b) shows the annual gas bill savings in Euro, while A.1(c) displays the annual return through gas bill savings on the total investment in insulation in percentages.

Figure A.2: Treatment effect per cohort

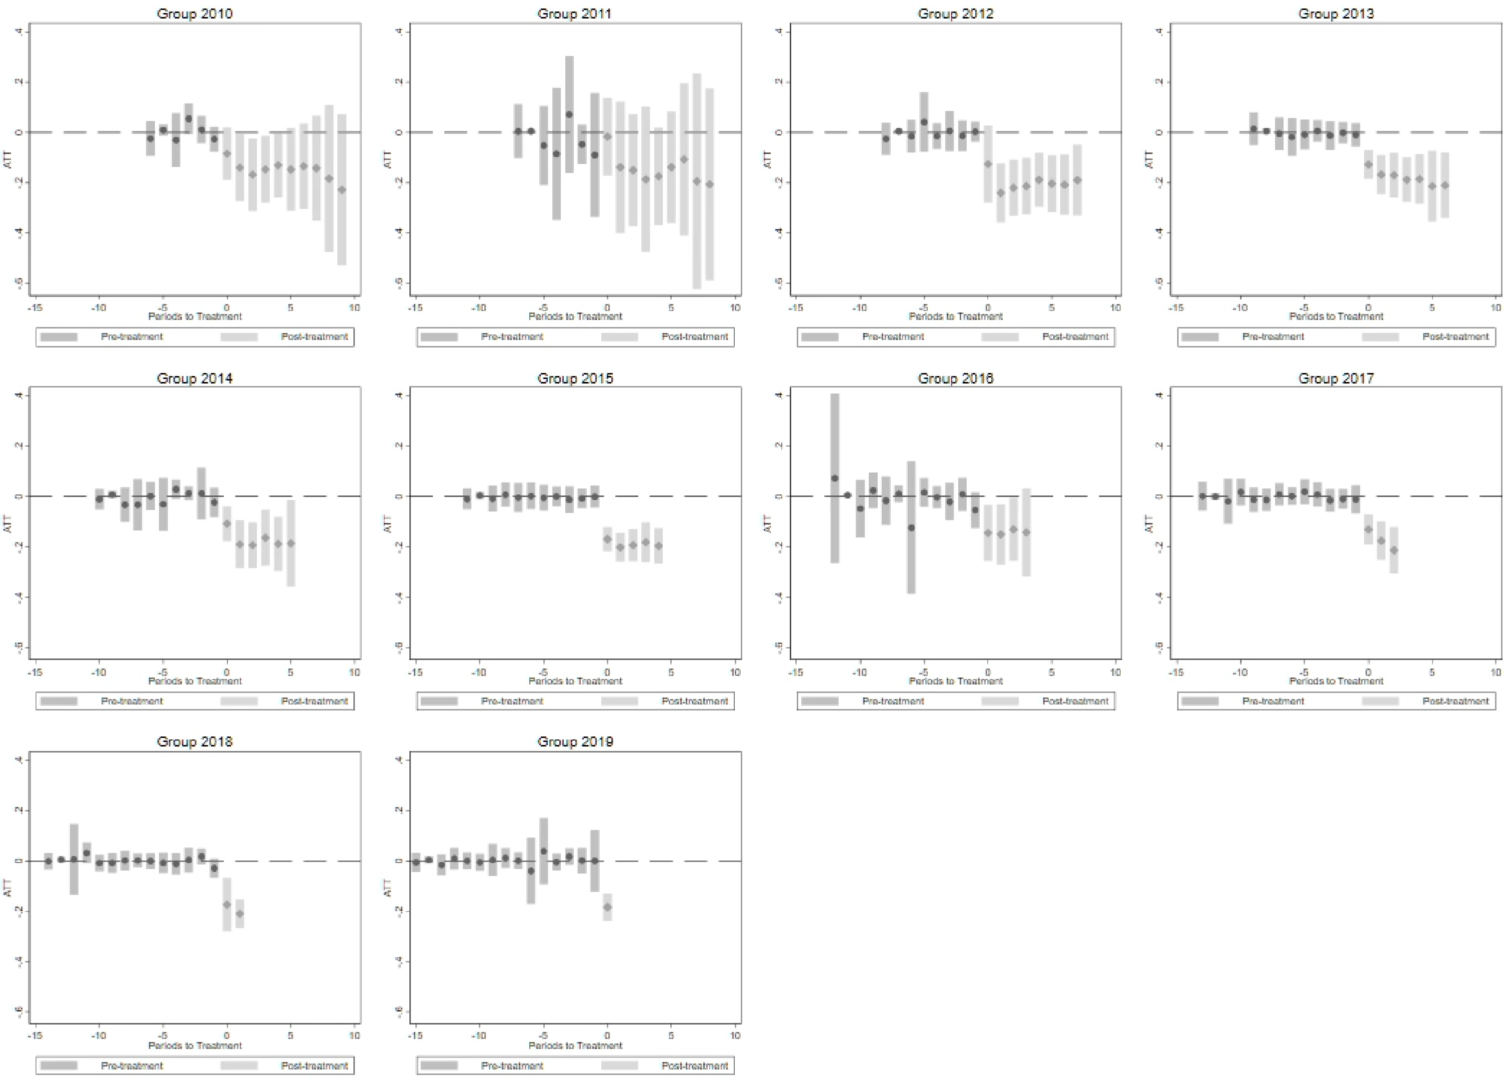

Notes: Figure A.2 presents the plotted coefficients of the main effect for different cohort, i.e. per year of insulation.