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

Linde Kattenberg* Piet Eichholtz† Nils Kok‡

June 2023

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, the impact of insulation measures on actual gas consumption is typically based on engineering predictions, and the efficacy of insulation measures is subject to debate. This study exploits a unique home insulation sample, 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 20%, on average, both for owner-occupied and rental homes. For the latter, the treatment is plausibly exogenous. We find no evidence of a temporal rebound effect: the reduction in gas consumption is consistent up to ten years after the intervention. The average treatment effect translates into a €273 reduction in the annual gas bill, and an average rate of return of 15.5% on the initial investment.

*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 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 the 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. Indeed, 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., 2021), with multiple explanations for the wedge. For instance, Christensen et al. (2021) find that heterogeneity in quality of installment, engineering mistakes, or a rebound effect are all factors that drive 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 size of the energy efficiency gap. This paper exploits unit-level costs and energy savings to precisely estimate the return to a variety of insulation measures across heterogeneous households.

Gaining a better understanding of the economics of insulation investments is also important given the dependence of other energy efficiency measures on the presence of insulation in the home. For instance, 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.[1]

In this study, 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 that 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 – this is important, given that the choice to install insulation is plausibly exogenous for tenants (i.e. the landlord decides on such measures). 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 in the free market (by a for-profit owner). 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 the difference-in-difference estimation method proposed by Callaway and Sant'Anna (2021).

The results of the empirical analysis show that gas consumption decreases, on average, by 19.84% after insulation is installed. Cavity wall insulation leads to the largest decrease in gas consumption, whereas the effects of basement insulation is smaller, and the effect of roof insulation insignificant. The results are robust to a variety of robustness checks, including the construction of different control samples. The reduction in gas consumption does not differ substantially across household types and dwelling types. Importantly, our 2010-2020

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

data set also allows for an analysis of the persistence of energy savings over time. We do not find evidence of a temporal rebound effect: the observed reduction in gas consumption remains around 20% 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, the yearly gas bill is reduced by €273, based on gas prices at the time of the investment. Using average gas prices observed in 2022 (during a period of war-induced spikes in energy prices), the annual savings equal €688, on average. These savings translate into payback periods of 6.4 years and 3.2 years, respectively. Assuming perpetuity of the savings, the return to insulation measures is 15.5% using gas prices at the time of investment, and 30.8% at 2022 prices (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 focuses 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., 2021). Generally, these studies find a sizeable decrease in energy consumption after insulation measures, or from a combined package of measures that include insulation. However, the size of the effect varies, is context-dependent, and could be heterogeneous over time. For instance, Peñasco and Anadón (2023) find a significant effect of insulation measures on energy consumption, that disappears within 4 years. Likely related to other measures that households undertake in their home, such as increasing the living space, which increases the demand for energy. Moreover, studies comparing actual energy savings to projected savings based on engineering estimates find large disparities between those two. For example, Allcott and Greenstone (2017), Fowlie et al. (2018), and Christensen et al. (2021) 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 [2].

Another unique factor of this study is that the data covers a relatively long time period – we include treated observations where home insulation took place more than 10 years ago. As such, we can also measure the long-run effects of home insulation. An increase in energy efficiency can lead to an *increase* in the use of energy-consuming appliances, because the unit costs of energy consumption decrease. This concept is defined in the literature as the "rebound effect." 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. We investigate the long-term gas consumption after insulation, and can therefore check for the presence of a temporal rebound effect, which does not seem to be present in our results.

The results in this paper have implications for homeowners, (public) investors in residential homes, as well as policymakers. Given the paucity of reliable information on 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 provide an indication that blanket subsidy programs should, in principle, not be necessary for home insulation, given the short payback period and high return on investment. Such subsidy programs should 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 home rental market, where

[2] For instance, Christensen et al. (2021) 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.

investors incur the capital cost and tenants typically benefit from energy savings. Another focus point of subsidy could be to provide subsidy for other energy efficiency measures that currently have a lower financial return, such as home battery storage systems, or heat pumps.

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., 2021). 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 main source of data used in this paper is from Bameco BV, a private insulation company based in the Netherlands. This company is a large insulation provider in Limburg, the most southern province of the country. The sole business of Bameco is home insulation, with a focus on basement insulation, cavity wall insulation, and roof insulation. The company maintains a (paper) archive for each home where an insulation intervention was carried out, including information on the cost of the installation, the type of insulation, and the date of the installation. We manually digitized the invoice 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 2,351 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 88% percent of households opting for that measure. Basement insulation is applied in some 17% of the sample, while roof insulation is the least popular, at 3%. 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.

—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 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 (slightly less than a third of the Dutch stock) and homes owned by private investors (15% 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

[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 outcome of the intervention.

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 they move to a different home, the later observations are dropped from the sample. Out of the 2,351 observations, we have energy consumption and household data on 2023 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 2994 households in the control sample (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 that are 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 should be orthogonal to the characteristics of the tenants, and rather be based on the quality and/or vintage of the rental home. 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 deviation of the home value is smaller than in the owner-occupied and private rental market segment. The wait lists for social housing are long, and assignment to a house 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 wait 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 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). Especially when effects are heterogeneous over time, or across different cohorts, 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 analysis, we are using the difference-in-difference estimator proposed by Callaway and Sant'Anna (2021).

Our main estimator of interest is the Average Treatment effect on the Treated as defined in Equation 1, where we compare treated households to never treated households:

$$ATT(g,t) = \mathbb{E}[Y_{i,t} - Y_{i,g-1}|G_i = g] - \mathbb{E}[Y_{i,t} - Y_{i,g-1}|G_i = NT] \tag{1}$$

where we measure the ATT of a group treated in time period g, in time period t. The effect is estimated relative to the outcome variable in the period before treatment, that is period g-1. The control group consists of never treated units. Gas consumption in time t for a group i is determined by Equation 2:

$$Y_{i,t} = \beta_0 + \beta_1 Insulation_{i,t} + X_{it} + \epsilon_{it} \tag{2}$$

where $Y_{i,t}$ 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 falls in the treatment period – one year after insulation – and has been subject to insulation treatment. X_{it} is a vector of home and household characteristics that can vary over time. ϵ_{it} is the error term, assumed to be independent of treatment and normally distributed, and clustered at the household level.

In our results, we will present the aggregated ATT. Following Callaway and Sant'Anna (2021), the weighted average of all cohort's ATT's is based on the amount of observations per cohort. We present the weighted averages either per n periods after treatment, t or on average across all post-treatment periods.

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. 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, but the slope of the pre-trend across both samples is exactly 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 the year specifically 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. Column (1) does not include any control variables. We document an average treatment effect of 19.6% after the insulation intervention, as compared to the control group of non-treated homes in the same province. Column (2) includes control variables that could affect gas consumption and that can vary over time, such as the dwelling surface, the number of household members, and household income. The treatment effect stays constant, with a decrease of 19.8% in annual gas consumption after the application of insulation measures in the home.

Of course, we may overestimate the effects of insulation on gas consumption given the 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 those homes in the free market and those owned by 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 units).

We document that gas consumption decreases by 20.3% in the sample of owner-occupied homes. We find a smaller effect for homes in the private rental segment and social housing, namely 17.5% and 13.2%, respectively. 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 private rental segment does not show a significant treatment effect in most of the post-treatment periods due to the relatively large standard errors. Especially in later periods, the estimated effect may be based on a relatively small sample of homes in the private 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 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. Roof insulation does not lead to a significant reduction in gas consumption in our analysis, but the low number of observations could contribute to a relatively high standard deviation and insignificant results. Wall insulation yields average savings in gas consumption of 20.8%. For basement insulation, we find a smaller, but still significant effect of 6.2% reduction in gas consumption. 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 combining wall insulation with basement and roof insulation yield higher gas savings than undertaking the measures in isolation. The combination of basement and roof insulation does not show any significant result, but we have to note that in our sample only 7 homes observed an improvement in insulation on these two aspects.

—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, and detached homes – respectively 18.9%, 22.8%, and 20.7%.

For dwellings that are wedged between other homes, so-called ”row homes”, the effect is somewhat smaller, at 15.8%. 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. 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. This reasoning could also explain why the effect is smaller for row homes. 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 split the sample according to the lower and upper 50% of the income distribution, separately for owner-occupied homes and rental homes in the free sector and those 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. In all three segments, we find that the confidence intervals of the estimates are overlapping, such that the point estimates are not statistically different. Thus, we conclude that we do not find heterogeneous effects across income levels of households.

—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. Figure 3 plots the coefficients of the difference-in-difference analysis as in Equation (1), estimating the ATT per period after treatment. 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. 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 – the effect of insulation as a structural change in the home remains as time progresses. As opposed to the attenuating effect of behavioral treatments, such as the Opower social comparison-based treatment (Allcott and Rogers, 2014), we do not find a change in gas consumption over time due to adjustments in household behavior.

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 Substitution Effect

The decrease in gas consumption following an insulation retrofit could also be due to a substitution effect. That is, gas consumption could be decreasing, while electricity consumption would be increasing. This could be the case if, for instance, households that insulate their home install a heat pump and heat their home with electricity instead of gas. In such cases, we would overestimate total energy savings by only considering the effect of insulation on gas consumption. Therefore, we perform the same analysis as our baseline difference-in-difference model, but substitute gas consumption with electricity consumption as the dependent variable. Using this estimation strategy, we can observe whether gas consumption in homes with improved insulation is substituted for electricity consumption.

The results of the analysis are presented in Table 7. We document that after controlling for home and household characteristics, we do not find a statistically significant decrease in electricity consumption after insulation. We can conclude that there is no substitution effect – installing insulation leads, on average, to a *reduction* in home electricity consumption, not an increase. In addition, we analyze the distribution of savings in gas consumption through a non-parametric estimate in Figure A.1(a). Here, we observe that just a very small group of households realizes a 100% reduction in gas consumption after the insulation intervention. This means that completely substituting natural gas, for example using electricity as an energy source, is not typically observed in our sample. We also note that the use of electric heat pumps was rare in the recent past – in the future, due to more stringent policies regarding gas consumption, and due to increased financial attractiveness of heat pumps, the adoption of heat pumps could be quite different.

—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 current insulation is of high quality due to building regulation that is effective since 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 home insulation through another company. We observe that in Columns 1 and 2, the effect size increases slightly as compared to the baseline model in Table 3, with an effect size of 20.2% and 23.6%, respectively. In Column 3, the effect size is slightly smaller than in our baseline estimation, which is counter to our expectation. 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, 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 show 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 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[4]. 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[5]. 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. In addition, we substitute the gas prices that the households (and landlords used at the time of their decision-making process, with 2022 prices (the period of a resource shock caused by the war in Ukraine). In this scenario, we also adjust the investment costs to 2022 levels. We do this by taking into account the average increase in insulation costs provided by the insulation company.

Table 9 displays investment costs, yearly savings, annual return, and the payback period [6]. 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 average results show an annual return of 15.5% at investment costs and energy prices prevailing at the moment of insulation, which corresponds

[4]For investors, the return calculation is complicated due to tenants directly benefiting from the investment in energy efficiency by the investor.

[5]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.

[6]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.

to a payback period of 6.4 years. We observe that annual returns from wall insulation are particularly high, with an average of 17.6%. For basement and roof insulation, the annual return is 5.4% and 7.5%, respectively[7]. However, we should be cautious when interpreting the returns to roof insulation, as our estimated effect size for this insulation type is not statistically significant. Considering the payback period, the average wall insulation investment of €1,656 will be earned back in about 5.7 years. For basement and roof insulation, the average payback period is 18.4 and 13.3 years, respectively.

Using average 2022 gas prices changes the investment decision considerably, with significantly shortened payback periods. The average annual return increases to 30.8%, a return that will be challenging to find for many other investments. For wall insulation, an average investment can be earned back already within 2.8 years. Basement insulation has a payback period of 9.5 years, and roof insulation has a payback period of 6.8 years. Considering that a household lives in a dwelling for around 10 years, all insulation types would be earned back within this period in the second scenario. That is, in the financial decision, the extent of capitalization of the insulation investment into the selling price is not relevant anymore.

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 that 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

[7]Figure A.1(c) displays the distribution of non-parametric 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.

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.

—Insert Table 9—

6 Conclusion

Improving the energy efficiency of the building stock is important to decrease 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 measure to enhance the energy efficiency of a home is wall, roof, or basement insulation. Such insulation also provides the basis for the subsequent installation of a heat pump, which would allow for the home to be taken off natural gas. 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 20%. 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 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 a longitudinal rebound effect. 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 €273 per year. Compared to the investment to install insulation, this yield a yearly return of 15.5%, translating into a payback period of 6.4 years. Wall insulation has the highest return, of 17.6%, while basement insulation returns 5.4% and roof insulation returns 7.5% per year. If the gas prices that households in our sample pay would be substituted by the 2022 gas price, an average annual return on investment of 30.8% would be realized, and the payback period of investing in insulation would be just 3.2 years.

Insights from this study contribute to better estimates of the returns to home insulation. Most of the literature focuses on the dearth of energy efficiency adoption, and explanations for the difference between predicted savings and realized savings (Allcott and Greenstone, 2017; Fowlie et al., 2018; Christensen et al., 2021). There are no papers that study the actual *returns* to energy efficiency investment at scale in a setting where selection effects can be overcome. The results in this paper can inform homeowners, investors, and social housing institutions in their home retrofitting decisions, reducing investment uncertainty. The results can also inform policymakers on the efficacy of energy efficiency in the housing market, which represents an important pillar in reducing carbon emissions. First, the information in this paper can be used to make more realistic expectations of the energy savings from insulation. Second, given the financial rates of return documented in this paper, there seems to be limited necessity for subsidy programs aimed at stimulating home energy efficiency measures in general, and home insulation in particular. 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, would be more suitable. Subsidy could be reallocated to energy efficiency measures that have high potential for mitigating CO_2 emissions, but currently have less favorable financial performance.

References

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. (2021). Decomposing the wedge between projected and realized returns in energy efficiency programs. *The Review of Economics and Statistics*, pages 1–46.

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.

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.

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-ja
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*.

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. *ORNL/CON-493. http://weatherization. ornl. gov/pdf/CON-493FINAL10-10-05. pdf. Washington, DC: US Department of Energy, Office of the Weatherization and Intergovernmental Program.*

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) Full Sample	(2) Owner-Occupied	(3) Rental Private	(4) Rental Social
# of insulation measures	1.129	1.121	1.169	1.195
	(0.369)	(0.344)	(0.422)	(0.537)
Wall				
Percentage	0.856	0.878	0.831	0.654
	(0.351)	(0.327)	(0.378)	(0.477)
Total cost in €	1593.810	1633.168	1891.067	970.886
	(754.2)	(650.4)	(2260.2)	(436.7)
Surface in m^2	104.269	106.701	125.229	64.734
	(52.34)	(45.23)	(157.4)	(32.74)
Basement				
Percentage	0.212	0.203	0.305	0.270
	(0.409)	(0.402)	(0.464)	(0.445)
Total cost in €	1333.941	1355.773	1301.312	1189.812
	(518.4)	(504.0)	(411.2)	(628.9)
Surface in m^2	52.783	53.125	53.250	50.170
	(21.72)	(20.74)	(16.25)	(29.30)
Roof				
Percentage	0.04	0.034	0.034	0.103
	(0.196)	(0.180)	(0.183)	(0.304)
Total cost in €	1959.725	1936.847	2047.000	2021.579
	(1184.3)	(1296.5)	(1588.2)	(769.5)
Surface in m^2	59.575	61.593	71.500	52.053
	(21.82)	(22.47)	(54.45)	(14.48)
Other				
Percentage	0.020	0.006	.	0.168
	(0.142)	(0.0752)	.	(0.374)
Total cost in €	1872.524	1750.482	.	2092.200
	(1065.4)	(680.5)	.	(1633.6)
Surface in m^2	54.571	62.667	.	40.000
	(33.25)	(34.15)	.	(29.12)
Observations	2023	1806	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 deviations 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 all non-treated 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 deviations are reported in parenthesis. * P<0.05. ** P<0.01. *** P<0.001

Figure 1: Insulation Measures Over Time

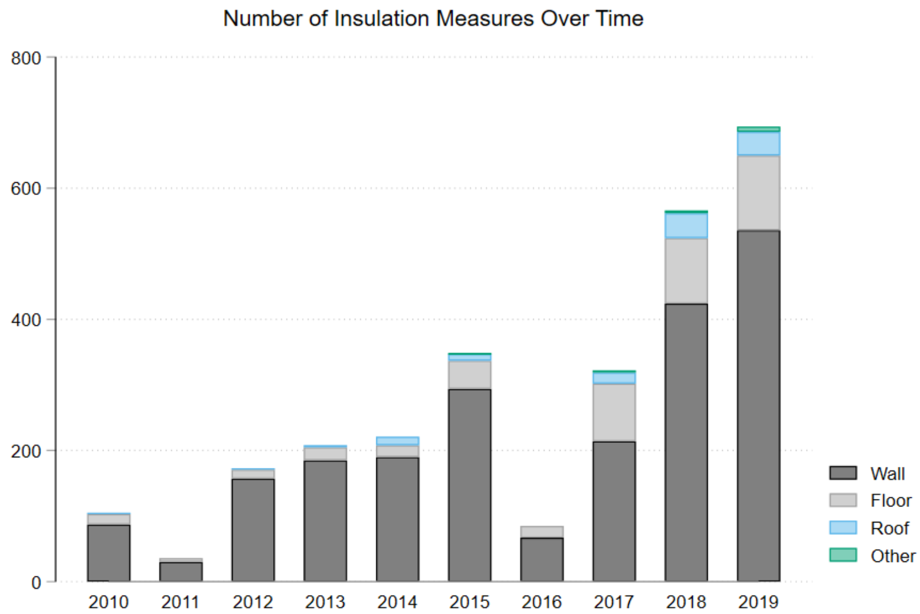

Notes: Figure 1 presents the number of recorded insulation measures in our sample over the sample period, split up per type of insulation.

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

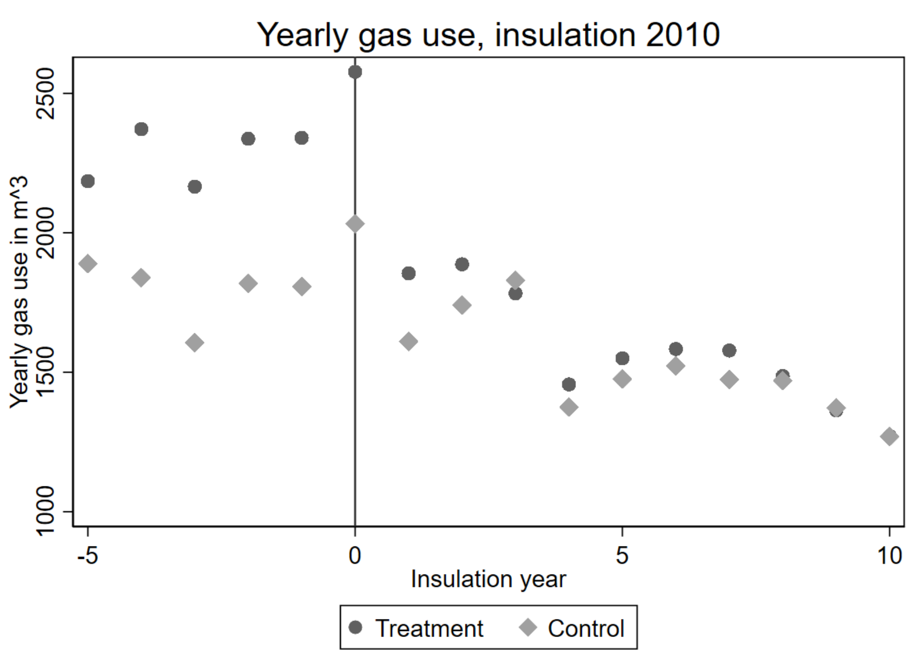

Notes: Figure 2 presents the mean of yearly gas use in the treatment and control group. Year 0 is the year of insulation.

Table 3: Insulation and Gas Consumption

	(1) Full Sample	(2) Full Sample	(3) Owner-Occupied	(4) Rental Private	(5) Rental Social
Insulation	-0.179***	-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 deviations 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	-0.189***	-0.0603***	-0.124	-0.225***	-0.374**	-0.100
	(0.00840)	(0.0189)	(0.0850)	(0.0276)	(0.181)	(0.0714)
Observations	70,380	55,105	52,435	54,468	52,279	51,567
Number of treated households	1,482	215	39	170	22	7
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 deviations 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	-0.0529	-0.173***	-0.205***	-0.147***	-0.188***
	(0.0577)	(0.0215)	(0.0138)	(0.0172)	(0.0150)
Observations	9,442	13,420	15,633	22,916	11,273
Number of treated households	42	307	492	379	320
Controls	YES	YES	YES	YES	YES

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

Table 6: Heterogeneity by Income Levels

	(1) Owner- Occupied Low Income	(2) Owner- Occupied High Income	(3) Rental Private Low Income	(4) Rental Private High Income	(5) Rental Social Low Income	(6) Rental Social High Income
Insulation	-0.198***	-0.188***	-0.141***	-0.185**	-0.136*	-0.141***
	(0.0166)	(0.00970)	(0.0626)	(0.0292)	(0.0700)	(0.0245)
Observations	27,699	27,023	4,206	4,288	8,819	9,425
Number of treated households	466	1236	33	178	80	136
Controls	YES	YES	YES	YES	YES	YES

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

Figure 3: Insulation Effect Over Time

(a) Full Sample

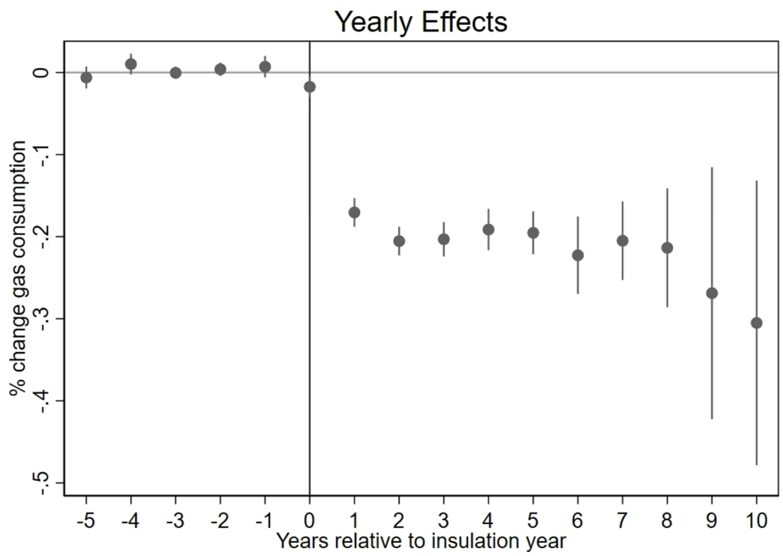

(b) Owner-occupied and Rental Homes

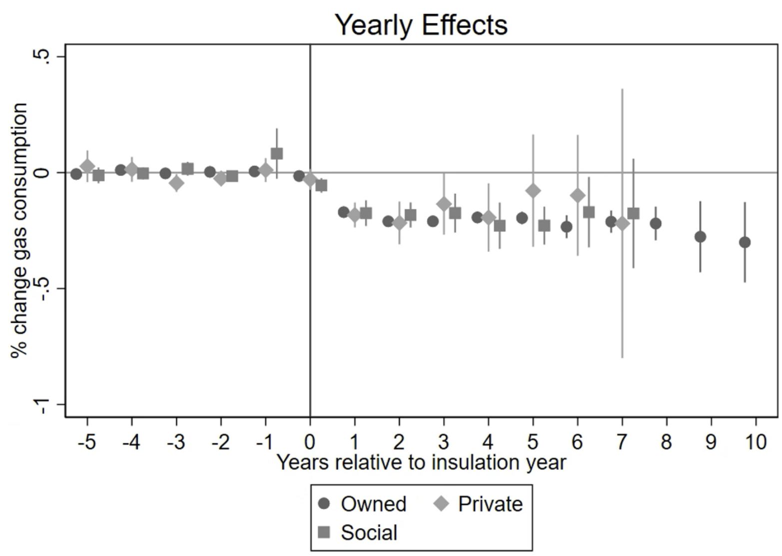

Notes: Figure displays annual gas consumption relative to year 0, the last year before insulation. The figure shows the point estimates, with the 95% confidence interval.

Table 7: Substitution Effects: Insulation and Electricity Consumption

	(1) Full Sample	(2) Full Sample	(3) Owner-Occupied	(4) Rental Private	(5) Rental Social
Insulation	-0.0188**	-0.00850	-0.0101	0.0440	-0.0113
	(0.00858)	(0.00870)	(0.00926)	(0.0297)	(0.0237)
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 electricity consumption. Standard errors are clustered at the household level. Standard deviations are reported in parenthesis. * P<0.05. ** P<0.01. *** P<0.001

Table 8: Robustness Checks: Restricting the Control Sample

	(1) Gas Consumption Change	(2) Building year >2000	(3) Same City	(4) Label Upgrade	(5) Improved Insulation	(6) Window Quality
Insulation	-0.202***	-0.212***	-0.171***	-0.183***	-0.186***	-0.182***
	(0.00757)	(0.0156)	(0.0150)	(0.00766)	(0.00768)	(0.00765)
Observations	53,028	26,591	15,512	72,661	69,488	73,802
Number of treated households	1,986	1,986	653	1,986	1,986	509
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 electricity consumption. Standard errors are clustered at the household level. Standard deviations are reported in parenthesis. * P<0.05. ** P<0.01. *** P<0.001

Table 9: Returns to Insulation Measures

	Actual prices				2022 prices			
	(1) All	(2) Wall	(3) Basement	(4) Roof	(5) All	(6) Wall	(7) Basement	(8) Roof
Yearly savings	€273	€291	€77	€177	€688	€739	€189	€444
Investment	€1,759	€1,656	€1,416	€2,359	€2,233	€2,103	€1,799	€2,997
Annual return	15.5%	17.6%	5.4%	7.5%	30.8%	35.1%	10.5%	14.8%
Payback period	6.4	5.7	18.4	13.3	3.2	2.8	9.5	6.8
Number of treated households	2,023	1,507	219	40	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. In column 5 to 8, we use the average 2022 consumer gas price. Investment costs are adjusted to 2022 prices, based on the average price development of the insulation company.

A Appendix

Figure A.1: Insulation Effect Over Time

(a) Gas Savings in Percentages

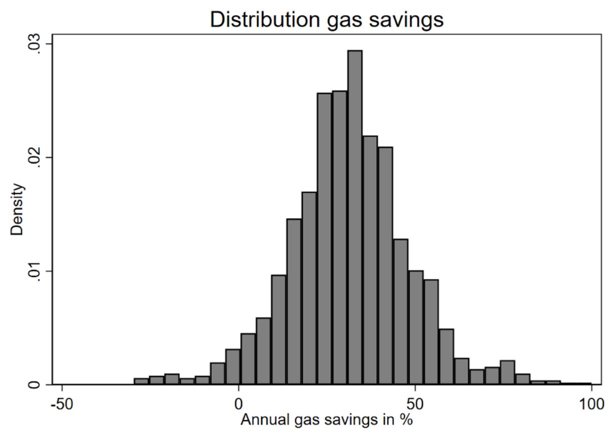

(b) Gas Savings in Euro

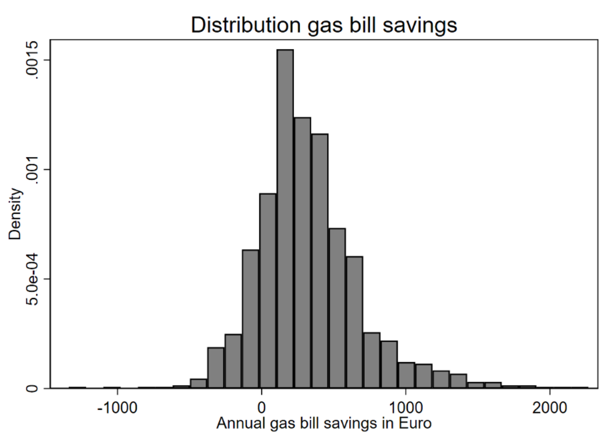

(c) Annual Return Investment

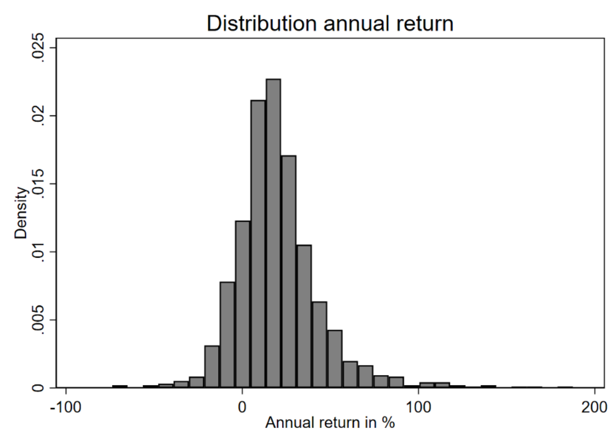

Notes: Figure A.1 displays the non-parametric distribution of gas savings based on the 5 years before, until maximally 10 years after improved insulation. Gas use in the year of insulation is not included in this calculation. Firstly, A.1(a) shows the annual gas savings in percentage terms. A.1(b) shows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

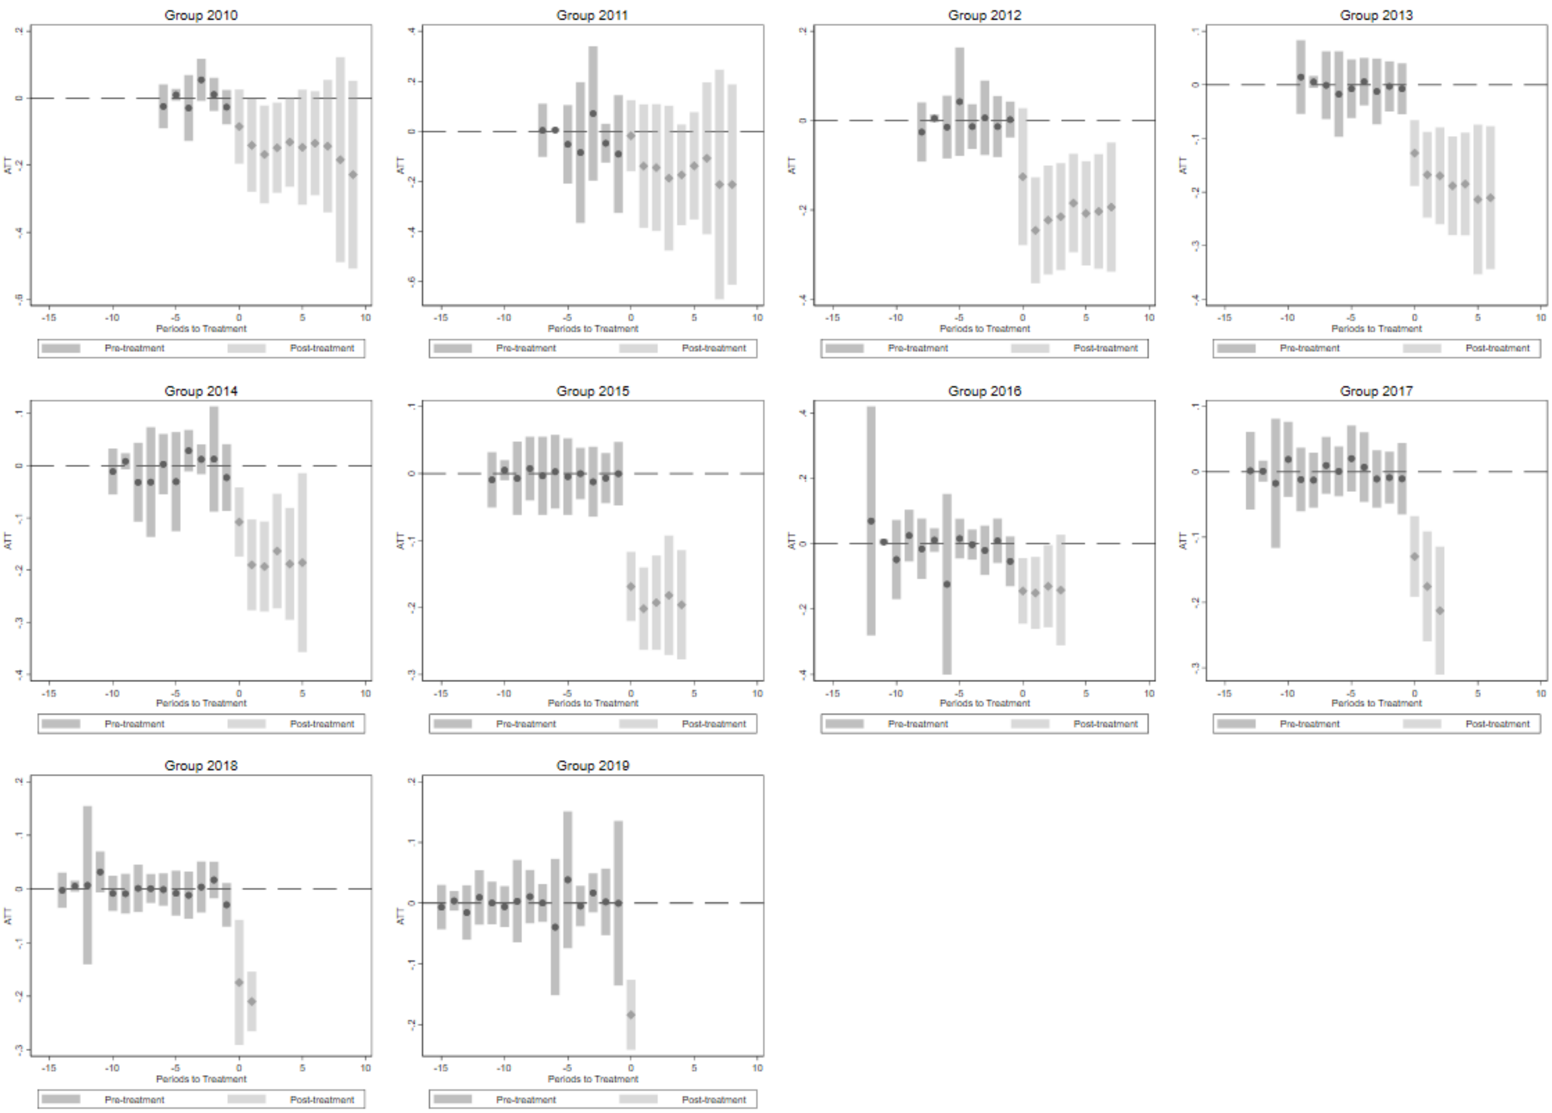

Notes: Figure A.2 presents the plotted coefficients of the main effect for different cohort, i.e. per year of insulation.

B Data and Code Availability

The data that support the findings of this study are available from the Dutch Central Bureau for Statistics (CBS) but restrictions apply to the availability of these data, which were used under license for the current study, and are not publicly available due to privacy concerns. Data are however available at CBS from the authors upon reasonable request.

The code to reproduce the analysis in the study is available from the authors upon request.