Do Homeowners Care About Sustainability?*

Milind Goel

London Business School, mgoel@london.edu.

Abstract

Although housing is the world's largest asset class and consumes over one-fifth of global energy, it is not well understood whether homeowners care about the environment. We compile a comprehensive dataset containing 55% of all residential property transactions in the UK from 2010 to 2020, and provide large-scale evidence that homeowners derive both pecuniary and non-pecuniary benefits from the energy efficiency of their dwellings. We show that homeowners behave in a rational manner, and price energy efficiency of a dwelling based on expected utility savings and their ability to recoup their investments. On aggregate, homeowners use a social discount rate of 4.83% to value investments in energy efficiency. Homeowners who purchase greener dwellings pay a premium in excess of the present value of future energy savings. We observe a commensurate increase in proportion of energy upgrades across market segments impacted and not impacted by regulation, but absence of a price-impact. This suggests that government interventions facilitated sustainable development through an indirect channel.

Keywords: Sustainability, Real Estate, Climate Finance, Household Finance

JEL Classification: R1, R3, R5, Q4, Q5

*I am especially grateful to my PhD advisor, Professor Victor DeMiguel, for his supervision, mentorship, and support during this project. I would also like to thank Professor Nitish Jain and Professor Derek Bunn for their helpful comments and feedback.

1 Introduction

Sustainability and *Climate Finance* have become increasingly important themes in financial economics and investment management. Baldauf, Garlappi, and Yannelis (2020) eloquently underpin the suitability of real estate as an asset class in addressing questions around these themes. First, the long-duration nature of buildings exposes them to the type of risks that arise from climate change. Second, for a majority of homeowners, real estate is the most important asset class. In fact, housing is the world's largest asset class (Eichholtz, Korevaar, Lindenthal, and Tallec, 2021). Third, real estate is an important source of household debt, adding to its relevance in the overall economy. Eurostat (2021) further shows that households represented 27.4% of Europe's final energy consumption in 2020. Therefore, understanding the impact of climate risk and sustainability in residential real estate markets assumes central importance. Surprisingly, however, it is not well understood whether homeowners care about the environment. This information gap can often act as a barrier to sustainable development, as it makes is hard for market participants, such as developers and real estate private equity firms, to ascertain whether their investments in energy efficiency will be priced by the market. It also makes it harder for policymakers to understand the potential incentives and barriers to improving the energy profile of the housing stock.

To investigate whether homeowners care about sustainability, we compile a comprehensive dataset containing 5.45 million transactions that span 55.56% of property sales recorded in England and Wales between 2010 and 2020, and over 20% of the total number of dwellings in the United Kingdom. We show that homeowners derive *both* pecuniary *and* non-pecuniary benefits from the energy efficiency of their dwellings.

We make five contributions to the existing literature. Our first contribution is to provide definitive large-scale evidence that energy efficiency is priced in the residential real estate market. We find that a 10-point increase in the energy efficiency rating (from 1 to 100) of a property is associated with a 1.8% increase in price per unit area. Although earlier studies have also attempted to document this *energy premium*, their analysis has been limited by the size and duration of their samples, or the availability of a sufficient number of relevant covariates to control for the heterogeneity inherent in the asset class. In constrast, our results are stable over time, robust to subsamples, and survive alternative econometric specifications.

Our second contribution is to show that homeowners factor expected utility savings into their investments in energy efficiency in a rational manner. Homeowners pay a higher energy premium in colder climates, where each unit increase in energy ratings is expected to lead to higher utility savings. A four-fold increase in heating requirements of a dwelling is associated with a 2.5 times increase in premium. We further show that buy-to-let landlords are willing to pay half the premium paid by buy-to-live homeowners. This finding empirically substantiates prior works that theorise several market imperfections which make it difficult for buy-to-let landlords to recoup their investments in energy efficiency (Jaffe, Stavins, and Cleveland, 2004; Rehdanz, 2007; Iwata and Yamaga, 2008; Gillingham, Newell, and Palmer, 2009; Gerarden, Newell, and Stavins, 2017).

Our third contribution is the empirical quantification of the *social discount rate* (SDR) that homeowners use to value reductions in energy expenditures.[1] We find that, on aggregate, the market discounts future energy savings at a rate of 4.83% or higher. In comparison, the 3.5% SDR adopted by the regulators (UK Government, 2022a) computed using the Ramsay Rule (Ramsey, 1928) is 27% lower. Because this approach does not take into account project-specific risks and costs of raising capital, the SDR used by homeowners may be higher. There is also widespread disagreement between economists on what values should be assigned to the parameters that go into computing the SDR (Groom and Maddison, 2019). Regulators may strategically select values that yield lower SDRs to make public policies look more attractive or to balance intergenerational social welfare. However, setting the SDR too low can misalign incentives of the policymaker and homeowners. For instance, consider a program that offers subsidies to homeowners to cover the difference between cost of upgrading the energy profile of a property and the present value of future energy savings. Regulators will overestimate the latter, and thus offer lower subsidies than what are required to incentivise homeowners to undertake energy efficiency improvements.

[1] An SDR is the discount rate used to appraise future costs and benefits in economic evaluations of public interventions.

Our fourth contribution is to show that homeowners who purchase more energy efficient properties pay a premium in *excess* of the present value of future energy savings. We define this difference as the *green premium*, which can be thought of as the value of the utility that homeowners derive from the sustainability of their dwellings. The market is willing to accept a 2.1% lower return (measured in terms of future utility savings) for dwellings labelled `C` than those labelled `E`. We show that green premium is increasing in the initial level of energy efficiency of a dwelling, evidenced by a declining SDR. Consequently, homeowners who intend to sell their properties are more likely to have the energy profiles of their properties upgraded, and to invest in larger energy improvements.

Our fifth contribution is to study how homeowners' investments in energy efficiency are impacted by the costs of upgrading the energy ratings of a property and government regulations aimed at improving the energy profile of the housing stock in the UK. We find that subsequent improvements in energy efficiency ratings become progressively more expensive. As a result, properties with higher initial energy efficiency ratings are less likely to have their energy profiles improved. An analysis of dwellings that had their energy profiles reappraised at the time of their second transaction shows that properties impacted by regulation are more likely to have their energy efficiency ratings upgraded after the regulation was approved than before. However, our analysis reveals that properties that were not impacted by regulation also saw a comparable increase in the proportion of energy upgrades. Thus, we argue that regulatory interventions helped improve the energy profile of the housing stock in the UK through an *indirect* channel, such as raising the market sentiment or awareness towards energy efficiency. This reasoning is supported by our finding that dwellings that are negatively impacted by the regulations *do not* transact at a discount post-policy.

Recent literature investigating impact of climate risk and sustainability in residential real estate asset markets has broadly focused on two areas. The first area is modelling climate and disaster risks into property valuations, contingent investors' beliefs about environmental risk. Bernstein, Gustafson, and Lewis (2019) show that coastal properties exposed to projected sea-level rise sell at an approximately 7% discount relative to otherwise similar properties. In contrast, Murfin and Spiegel (2020) find no significant price effects when they match residential real estate transactions with property-elevation and tidal data to compare prices of otherwise similar homes but for which the time to inundation will differ depending on the pace of sea level rise. Baldauf, Garlappi, and Yannelis (2020) show that homes projected to be underwater because of climate change and located in climate change *denier* neighbourhoods sell for about 7% more than homes in *believer* neighbourhoods. We extend this area of study by demonstrating that homeowners respond to climate risk and sustainability more broadly, and not only when their dwellings are subjected to more "immediate" disaster risks. This is important, because disaster risks can influence property prices through channels other than environmental concern, such as changes in insurance premia or threat of physical destruction.

Concurrent and independent work to ours by Clara, Cocco, Naaraayanan, and Sharma (2022) shows that for private-rental properties that fall into the bottom tercile when sorted by energy efficiency scores, the regulation triggered an increase in the issuance of subsequent energy certificates and the magnitude of improvement in energy ratings between these subsequent issuances. In contrast, our analysis reveals that the regulatory impact was not limited to the properties or market segments that were directly impacted by regulation. Whereas Clara, Cocco, Naaraayanan, and Sharma (2022) study the frequency, timing, and magnitude of investments in energy efficiency for lower-rated properties in response to policy, our main focus is on whether homeowners derive non-pecuniary benefits from the energy efficiency of their dwellings, and how these attitudes are impacted by the regulation. Furthermore, having been drawn from a wide array of sources, our dataset contains a richer set of covariates.

The second strand of literature studies the mechanisms through which energy efficiency is capitalised into property prices. Several countries in Europe assign energy efficiency ratings to buildings in order to measure and certify their *level* of sustainability (European Commission, 2002). The scales on which properties are rated and the laws governing the adoption of the rating system vary across countries. In general, dwellings with better energy ratings are expected to expend less energy, demonstrate higher resilience to climate risks, and increase the utility of homeowners and renters who *care* about environmental impact. Hence, more energy efficient properties should command a premium. Consequently, there have been several works that measure this *energy premium* across countries in Europe (e.g., Brounen and Kok, 2011; Amecke, 2012; Cajias and Piazolo, 2013; Högberg, 2013; Hyland, Lyons, and Lyons, 2013; Fuerst, McAllister, Nanda, and Wyatt, 2015a; Ramos, Pérez-Alonso, and Silva, 2015; Fuerst, McAllister, Nanda, and Wyatt, 2016a; Jensen, Hansen, and Kragh, 2016).

However, the typical barrier to estimation of energy premium in real estate finance has been the lack of availability of a dataset that not only includes enough *relevant* covariates to capture the heterogeneity in the asset class, but also contains enough observations over a sufficiently long duration in order to produce estimates that are representative of the population parameters, and robust across subsamples and alternative econometric specifications. For instance, Cerin, Hassel, and Semenova (2014) analyse 67,599 transactions from 2009 to 2010 and note that while there is a positive relationship between energy ratings and property prices over their full sample, this relationship reverses in their subsamples. Similarly, several studies highlight that their estimates for energy premia are specific to the regions (typically, metropolitan areas) from which their samples are drawn (e.g., Davis, McCord, McCord, and Haran, 2015; Ayala, Galarraga, and Spadaro, 2016; Fuerst, Oikarinen, and Harjunen, 2016b) and may not represent the overall market sentiment towards energy efficiency.

In contrast, size of our dataset is more than five times that of the second largest study on energy premium (Cajias, Fuerst, and Bienert, 2019), and the duration of our dataset is two times that of the second longest study in the literature (Fuerst, Haddad, and Adan, 2020). Our dataset samples transactions proportionately over time and across regions, and the characteristics of our compiled dataset closely match those of the population databases. We consistently obtain an R^2 of over 78% in our linear regressions. In addition, the coefficients of all our covariates are not only statistically significant at a 0.1% level, but also make economic sense. In comparison to the existing studies, our estimates for energy premium are modest by up to a factor of four. This paper further extends this stream of literature by exploring the heterogeneity in energy premium across various channels, such as tenure, property type, ambient temperature, and level of urbanisation.

Lastly, our paper contributes to the active and ongoing debate in the economics of climate change on the appropriate discount rate to be used when valuing investments in sustainability (Stern, 2007; Nordhaus, 2007; Groom and Maddison, 2019; Kaplow, Moyer, and Weisbach, 2010; Schneider, Traeger, and Winkler, 2012). Giglio, Maggiori, Rao, Stroebel, and Weber (2021) argue that discount rates observed in real estate markets are informative about those for investments in climate change abatement, as real estate is specifically exposed to climate risk, which is reflected in house prices.[2] Therefore, the empirically observed rate that homeowners use to discount future reductions in energy expenditures can provide a useful benchmark for SDRs.

We conclude this section with an outline of the paper. Section 2 details data construction and sample properties. Section 3 computes energy premium and the implied SDR. Section 4 explores heterogeneity in premia. Section 5 provides evidence for existence of green premium. Section 6 studies impact of upgrade costs and regulations on homeowners' investment decisions. Section 7 concludes. The internet appendix supplements the manuscript with technical details, additional results, follow-up discussions, and robustness checks.

2 Data

We compile official datasets on the energy performance of buildings, price paid for residential property transactions, socio-economic indices, gridded land surface temperature records, and urban classifications, published by various departments of the UK Government to compile a comprehensive dataset that contains 5,451,054 transactions, which constitute 55.56% of the 9,808,400 property sales registered with the HM Land Registry between 01 January 2010 and 31 December 2020. Our sample spans 4,473,099 unique properties, which constitute roughly 20% of the total number of dwellings in the UK. Because these databases are produced by different agencies, they are not linked together through unique keys, and are often aggregated at different levels of spatial granularities. Therefore, we leverage a wide array of databases on regional mappings and geographic boundaries published by the Office of National Statistics to aggregate each database at an appropriate level of spatial granularity and to map entries from one database to another. Table 1 enumerates and summarises the databases deployed

[2] Note that Giglio, Maggiori, Rao, Stroebel, and Weber (2021) study the term structure of discount rate for real estate. In this paper, we focus only on the aggregate discount rates that the market uses to value investments in energy efficiency.

in this paper. Section 2.1 introduces each database and how it was compiled in a sequential manner; technical details are deferred to Section IA.1 of the Internet Appendix. Section 2.2 illustrates that the properties of the compiled sample match closely of the population data, indicating that our findings can be generalised.

2.1 Data Construction

2.1.1 Price Paid Data

The Price Paid Data is published by the HM Land Registry and provides information about residential property transactions recorded starting January 1995. For each transaction, the database records the transaction value, transaction date, select building characteristics such as property type, tenure, and an indicator variable for whether the property is new. Of the 27,359,802 property sales recorded since 01 January 1995 until 27 June 2022[3], 9,808,400 transactions occur between 01 January 2010 and 31 December 2020.[4] Our goal is to maximise the number of transactions retained in final sample.

2.1.2 Energy Performance of Buildings Data

The Energy Performance of Buildings Data is published by the Department for Levelling Up, Housing & Communities and provides data on Energy Performance Certificates (EPC) issued for residential properties starting 01 October 2008 until the month of September of the most recent calendar year, grouped by 341 *Local Authorities* (or administrative units) in the UK.[5] Each EPC provides information about the *current* and *potential* measurements for the energy efficiency of the properties for which the certificate was issued, together with property characteristics (e..g, floor area, built form, construction period), measurements of environmental impact (e.g., carbon emissions) and utility costs (e.g., heating).

[3]The most recent date on which the data was downloaded.

[4]We restrict our sample between 01 January 2010 and 31 December 2020 because Energy Performance of Building Data contains starts in October 2008 and does not contain sufficient data for 2008 and 2009.

[5]We refer to local authorities as districts, regions, or boroughs interchangeably.

A. Mapping the Energy Performance of Buildings Data to Price Paid Data

In order to investigate the relationship between property values and energy profiles, and to control for building and transaction characteristics, we must link each transaction recorded in the HM Land Registry with a *valid* EPC.[6] However, there does not exist a unique key or locational identifier that can provide a one-to-one mapping between the two datasets. Therefore, the only method to link the two datasets is through *address matching*.

Unfortunately, addresses are not entered consistently within and between datasets. For example, "`FLAT 42, 16A BROADWAY, 413`" may also be recorded as "`42 BROADWAY, 16A 413`". One method to link addresses is to use *fuzzy* matching techniques such as the *Levenshtein* distance, which computes the minimum number of single-character edits required to change one word into the other. However, such methods suffer from several drawbacks, as illustrated in Section IA.1.1 of the Internet Appendix. Furthermore, given the heterogenous nature of real estate, inexact matches may distort results substantially. Therefore, we design a custom algorithm that produces *exact* matches, which is detailed in Section IA.1.2 of the Internet Appendix. The trade-off is a smaller dataset post-compilation. We are able to uniquely map 7,239,549 transaction entries.[7]

B. Feature Selection and Formatting

Due to missing values, we must trade-off the number of covariates with the total number of entries in the dataset.[8] There is no fixed rule on how to accomplish this. Nonetheless we are able to retain all features that are of first order importance; and are enumerated in Table 2. Thereafter we format (or clean) the mapped dataset feature-by-feature. We defer to Section IA.1.3 of the Internet Appendix for the implementation details. Table 3 provides a quick summary of the key operations in the order in which they are carried out, together with the number of entries lost at each step. The resultant sample contains 7,022,645 entries.

[6] A *valid* EPC is defined as the most recent certificate for a property lodged no earlier than 10 years before the transaction date, as mandated by the law.

[7] Of the 9,808,400 transactions recorded in HM Land Registry between 01 January 2010 and 31 December 2020, only 9,692,971 transactions take place in postcodes for which we have entries in the Energy Performance of Buildings Data.

[8] See Section IA.1.3 of the Internet Appendix for a more detailed discussion on this trade-off.

C. Upgrade Costs

Each EPC is complemented by a *recommendations* document, uniquely identifiable using a `LODGEMENT KEY`. Each document provides a list of improvements and their expected range of costs (i.e., minimum to maximum) to *upgrade* the property from its current to its potential energy efficiency rating. We extract the cost metrics from each document. Then, for each EPC, we take the average of the suggested range of cost of upgrade for each line item, and then sum over these costs. Doing so provides us with a measure for the expected *Upgrade Cost* associated with each energy certificate.

2.1.3 Multiple Deprivation Indices

The Multiple Deprivation Indices (MDI), or the English Indices of Deprivation, are published by the Ministry of Housing, Communities & Local Government for the years 2007, 2010, 2015 and 2019, and provide a measure of multiple deprivation experienced by people living in an area. An assortment of indicators is weighted to produce seven component indices (Income, Employment, Health Deprivation, Education, Crime, Housing Barrier, and Living Environment) and an overall composite Index of Multiple Deprivation (IMD) for every Lower layer Super Output Area (LSOA). LSOA are a geographic hierarchy designed to improve the reporting of small area statistics in the UK.

There are four considerations in compiling the indices for analysis. First, the format in which these indices are recorded is inconsistent across reports. Second, the LSOA classification used in 2007 and 2010 is different from that used in 2015 and 2019. Third, we must select one of two formats in which the indices are reported: *scores* or *ranks*. Fourth, we must interpolate indices for those years between 2010 and 2020 for which we do not have a MDI report. Section IA.1.4 of the Internet Appendix explains how we address each consideration. We use the Postcode to LSOA 2011 Lookup published by the Office of National Statistics to assign each postcode in our compiled dataset to its corresponding LSOA 2011. We then use LSOA 2011 and transaction year to link MDI to previously compiled data.

2.1.4 Degree Days

Heating *degree days* is a measure derived from the historical temperature observations of a region, and is directly proportional to the heating requirements of buildings in that region. To construct degree days, we use temperature averages (TAS) published by the Met Office from 1862 to 2020 derived from a network of land surface observations. The data is available at various frequencies (daily, monthly, annual) and at various resolutions (5×5km, 12×12km, 25×25km). We use monthly TAS recorded over 10,432 5×5km grid points (each represented by a coordinate). Section IA.1.5 of the Internet Appendix provides a detailed documentation how year-wise degree days measures for each region are constructed. We use the LSOA 2011 Boundaries database published by the Office of National Statistics to extract representative coordinates for each of the 32,844 LSOA 2011, and assign to them the degree days values for years 2008 through 2021 corresponding to the closest 5×5km grid. As in the case of MDI, we use LSOA 2011 and transaction year to link degree days to previously compiled data.

2.2 Sample Properties

The dataset compiled in Section 2.1 contains 7,022,645 transactions. We refer to this dataset as the *merged sample*. We remove entries with missing values for any of the features used in regressions in Section 3. The resultant dataset consists of 5,451,054 transactions. We refer to this dataset as the *regression sample*.

This section compares the properties of the regression sample to those of the parent datasets. If the properties of our sample closely matches that of the population data, then we can be reassured that our analysis in the subsequent sections is generalisable, and that the estimates obtained from our regressions are representative of the population parameters.[9]

For each quarter, Figure 1 plots the number of transactions retained in the sample against the number of transactions in the Price Paid Data. We observe that the horizontal dashed

[9] Although we refer to the parent databases as the *population data*, it is not strictly true. Properties that may not have had their energy profiles appraised would not present in the Energy Performance of Buildings Data, as an EPC was not issued. Similarly, Price Paid Data may exclude domestic transactions that were not registered with the HM Land Registry at the time of a sale.

Figure 1: Entries Sampled per Quarter

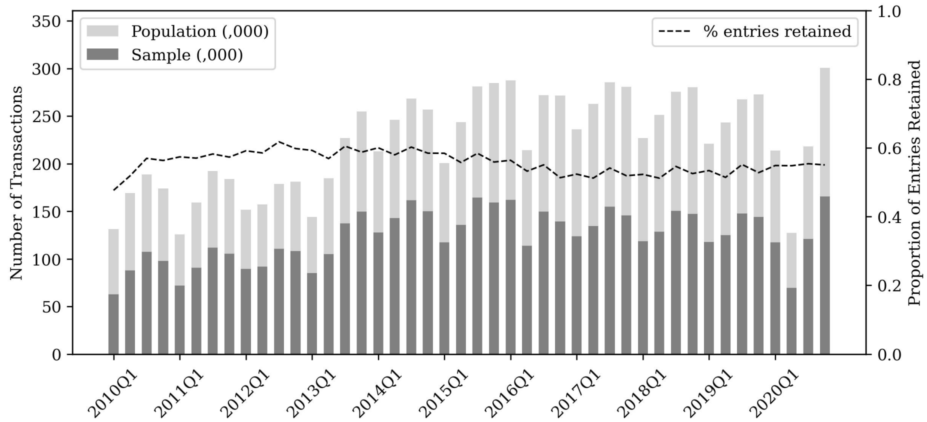

This figure shows how transactions are sampled over time. The bars correspond to the primary (left) y-axis, and the horizontal dashed line corresponds to the secondary (right) y-axis. For each quarter, the light grey bars mark the total number of transactions recorded in the Price Paid Data, while the dark grey bars mark the number of transactions retained in the regression sample. The dashed line illustrates the proportion of transactions sampled.

line, which shows the proportion of transactions retained, is stable over the duration of our sample. Figure IA.4 in Section IA.1.6 of the Internet Appendix shows that for each quarter, the proportion of *new builds* in our sample at the time of transaction is the same as that observed in the Price Paid Data. Figure 2 shows that the proportion of entries present in our sample across regions belonging to different *levels* of urbanisation is comparable to those observed in the population data. Figure IA.6 in Section IA.1.6 of the Internet Appendix further shows that the number of transactions sampled for each local authority is approximately proportional its population. Finally, the *Quantile-Quantile* (QQ) plots in Figure 5 reveal that when we chart the quantiles of the degree day measure and the composite multiple deprivation index corresponding to transactions in the Price Paid Data (x-axes) against those in our regression sample (y-axes), we obtain near perfect fits.

We now focus on comparing distributions of some key features of interest. The QQ plot in Panel (a) of Figure 4 charts quantiles of the logarithm of transaction prices in the Price Paid Data (x-axis) against those retained in the sample (y-axis). The grey circular markers closely line up along the 45-degree dashed line, indicating that the distribution of property prices in the sample is comparable to that in the population. The four QQ plots in Panel (b) of Figure 4 chart the quantiles of property sizes in the Energy Performance of Buildings Data

Figure 2: Proportion Entries by Rural Urban Code

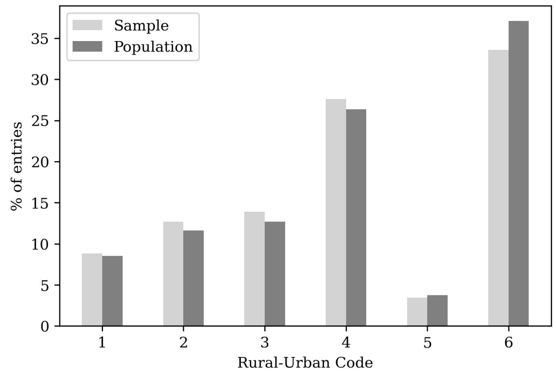

The light grey bars illustrate the proportion of entries present in the regression sample for regions belonging to different *levels of urbanisation*. The dark grey bars illustrate a similar breakdown for the Price Paid Data. The Rural Urban Classification (RUC) codes, published by Department for Environment, Food & Rural Affairs, categorise local authority districts in the UK from most rural (1) to most urban (6).

Figure 3: Distribution of Environmental Controls

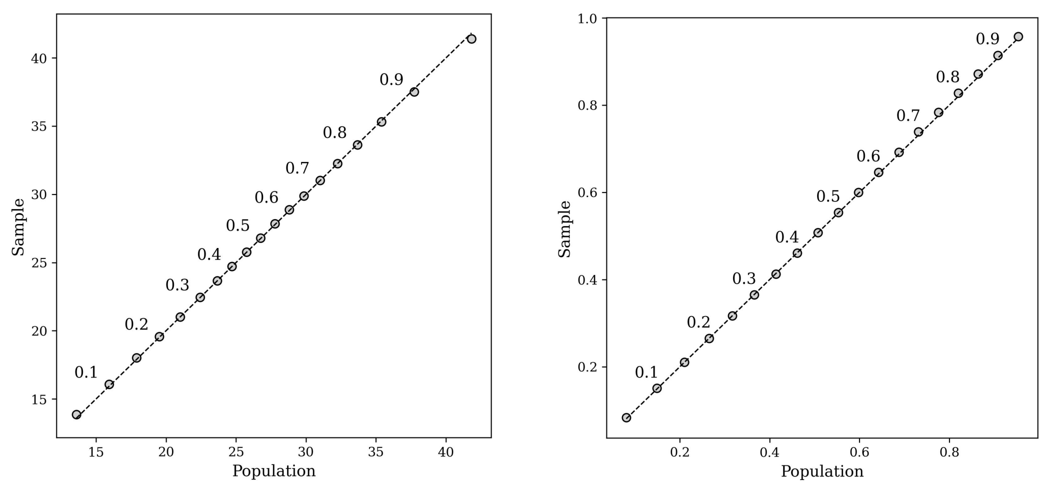

(a) Degree Days

(b) Multiple Deprivation Index

The Quantile-Quantile (QQ) plot in Panel (a) charts quantiles of the degree days measure corresponding to transactions in the Price Paid data (x-axis) against those in regression sample (y-axis). Panel (b) charts quantiles of the composite multiple deprivation index corresponding to transactions in the Price Paid data (x-axis) against those in regression sample (y-axis).

Figure 4: Distribution of Features Used in Construction of Dependent Variable

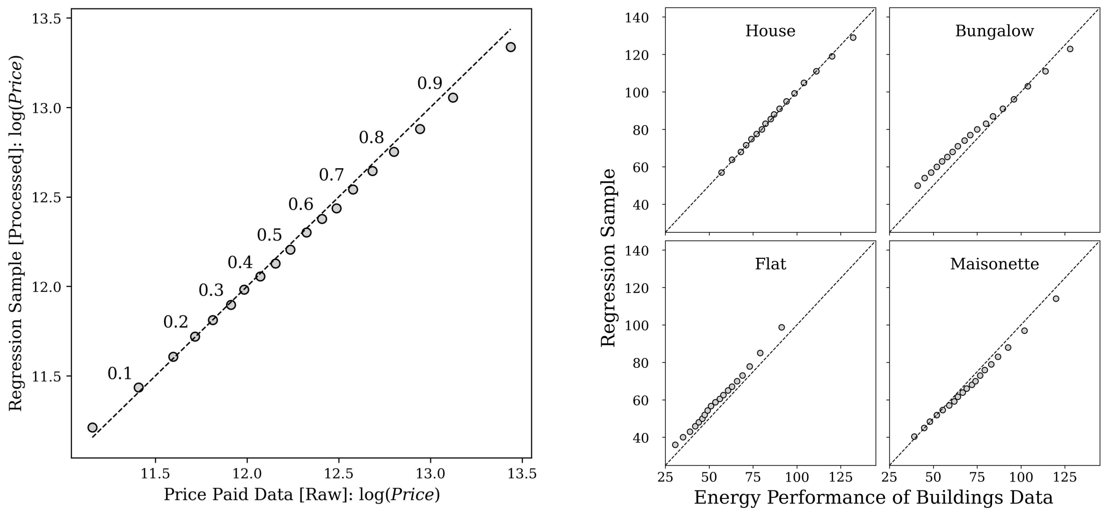

(a) log(Price) **(b)** Total Floor Area

The Quantile-Quantile (QQ) plot in Panel (a) charts quantiles of the logarithm of property prices in the Price Paid Data (x-axis) against those in regression sample (y-axis). Panel (b) shows QQ plots for total floor area for each property type present in the Energy Performance of Buildings Data. The grey circular markers closely line up along the 45-degree dashed line, suggesting that distributions of property prices in sample and in population are comparable. Note that the comparison is not direct, as the regression sample was subject to a range of operations (see Table 3); in particular, records with price per unit area outside the 0.5% and 99.5% quantiles were removed.

Figure 5: Distribution of Energy Ratings

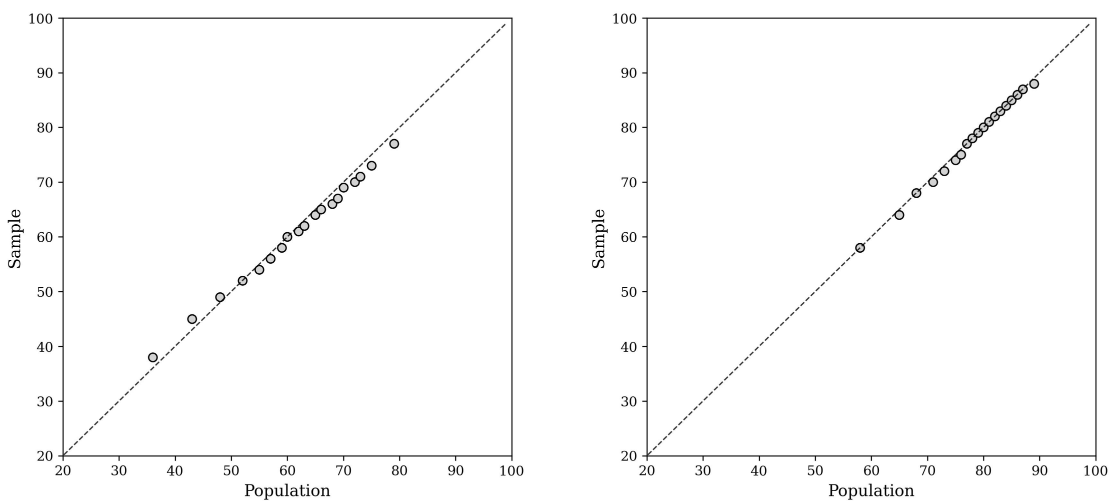

(a) Current Energy Score **(b)** Potential Energy Score

The Quantile-Quantile (QQ) plot in Panel (a) charts quantiles of current energy efficiency scores in the Energy Performance of Buildings Data (x-axis) against those in regression sample (y-axis). Panel (b) shows a QQ plot for potential energy efficiency ratings.

(x-axis) against those retained in the sample (y-axis). Modest deviations from the 45-degree line are expected because we restrict our sample to properties with a total floor area between 20 and 400 squared meters.[10]

Note here that we map energy certificates for only those properties that correspond to a sale in the HM Land Registry over the duration of our sample. Because the rates at which properties are sold are heterogenous across different market segments, the composition of tenures (e.g., owner occupied, private rental) and property types (e.g., flats, houses) in our sample is different from that of the Energy Performance of Buildings Data.[11] We account for the changes in composition by performing subsample analysis for different market segments in Section 4. Reassuringly, we learn that the properties of our sample match those of the the Energy Performance of Buildings Data for all other features enumerated in Table 2. Panels (a) and (b) of Figure 5 show QQ plots for the numerical current and potential energy efficiency ratings respectively, revealing that the energy profile of the housing stock in our sample is representative of the population distribution. Similarly, Figure IA.5 in Section IA.1.6 of the Internet Appendix illustrates that the composition of the properties belonging to different construction age bands in the sample reflects that observed in the Energy Performance of Buildings Data. In summary, the results in this section indicate that our analysis in the subsequent sections is indeed generalisable.

[10]The distribution of total floor area is heterogenous across property types. For instance, the median size of a flat in the Energy Performance of Buildings Data is $54m^2$ while a house is $88m^2$. Therefore, the size restriction eliminates 1.79% of flats but only 0.05% of houses from the data.

[11]In fact, our *merged sample* produced in Section 2.1 provides reasonable benchmark for the rate at which properties across different market segments have been sold over time. We also note here that our exact-matching approach (see Section IA.1.2 of the Internet Appendix) slightly undersamples flats and maisonettes because when multiple housing units in the same building omit their SAON, it is not possible to uniquely identify them.

Table 1: Datasets

Dataset	Downloaded	Description
Price Paid Data	14 Feb 2022	Published by Her Majesty's Land Registry, the dataset provides transaction values for all domestic properties sold and registered with the registry, starting from 01 January 1995 to 28 February 2022 (most current month).
Energy Performance of Buildings Data	14 Feb 2022	Published by Department for Levelling Up, Housing & Communities, the EPC dataset tabulates all domestic energy efficiency certificates issued between October 2008 and September 2021 for each local authority in the United Kingdom.
2007 Multiple Deprivation Indices 2010 Multiple Deprivation Indices 2015 Multiple Deprivation Indices 2019 Multiple Deprivation Indices	20 Oct 2021	Published by Ministry of Housing, Communities & Local Government, the English indices of deprivation use 30+ indicators organised across seven categories (e.g., income, employment,...) that are then weighted appropriately to produce an overall deprivation index (IMD). The 2007 and 2010 indices are provided for each of the 32,482 LSOA 2001, while the 2015 and 2019 indices are provided for each of the 32,844 LSOA 2011. The number of indicators and the exact weighting scheme slighly varies across reports.
Gridded Monthly Average Temperature	13 Oct 2021	Provided by Met Office, the dataset contains monthly temperature averages (TAS) form 1862 to 2020 derived from the network of land surface observations for 10,432 5x5km grid points.
Postcode to LSOA 2011 Lookup	14 Nov 2021	Published by Office of National Statistics, the dataset provides mapping between LSOA 2011 and Postcodes.
LSOA 2001 to LSOA 2011 Lookup	14 Nov 2021	Published by Office of National Statistics, the dataset provides mapping between LSOA 2011 and LSOA 2001.
Local Authority District Boundaries	19 Feb 2022	Published by Office of National Statistics, the dataset provides centroid coordinates and the digital vector boundaries for Local Authority Districts in the United Kingdom as at May 2021.
LSOA 2011 Boundaries	19 Feb 2022	Published by Office of National Statistics, the dataset provides centroid coordinates and the digital vector boundaries for the 32,844 LSOA 2011 in the United Kingdom as at December 2011.
Rural Urban Classification Lookup	19 Feb 2022	Published by Department for Environment, Food & Rural Affairs, the RUC lookup classifies local authority districts in the UK from most rural (1) to most urban (6).

Each entry in the *Datasets* column is hyperlinked to the website from which it was downloaded; the links were last checked on 01 March 2022. The entries in the *Downloaded* column provide the most recent date of download.

Table 2: Feature Selection

Category	Feature	Measurement
General	Price[1]	GBP, e.g., 102500, 45000
	Transaction Year[1]	Year, e.g., 2008, 2019
Energy Ratings	Current Energy Label[2]	Alphabetical label from G to A
	Potential Energy Label[2]	
	Current Energy Score[2]	Numerical score from 0 to 100
	Potential Energy Score[2]	
Building Properties	Total Floor Area	Squared meters, e.g., 60, 85
	Property Type	Categorical, e.g., House, Flat
	Built Form	Categorical, e.g., Detached, Mid-Terrace
	Habitable Rooms	Integer, e.g., 2, 3, 4
	New[1]	Categorical, either Y or N
	Construction Age Band	Categorical, e.g., 1900-1929, 1930-1949
	Glazed Area	Categorical, e.g., Normal, More than Typ.
	Multi-Glaze Proportion	Percentage value from 0 to 100
	Low-Energy Lighting	Percentage value from 0 to 100
Transaction Characteristics	Tenure	Categorical, e.g., Owner-occupied, Rental
	Transaction Type	Categorical, e.g., Marketed Sale
Environmental Metrics	Current Environmental Impact	Numerical score from 0-100 based on
	Potential Environmental Impact	carbon emissions (higher is better)
	Current Energy Consumption	Annual energy consumption measured in
	Potential Energy Consumption	kWh per squared meter
	Current Carbon Emissions	Tonnes per year
	Potential Carbon Emissions	
Utility Costs	Current Lighting Cost	Annual cost in GBP
	Potential Lighting Cost	
	Current Heating Cost	
	Potential Heating Cost	
	Current Hot Water Cost	
	Potential Hot Water Cost	
Locational Identifiers	Postcode	Alphanumeric, e.g., NW1 4SA, HA9 0QE
	Local Authority	Alphabetical, e.g., Camden, Oxford
	Local Authority Code	Alphanumeric, e.g., E06000042
	Constituency Code	Alphanumeric, e.g., E14000822

1. Features that belong to the Price Paid Data.
2. Current (Potential) Energy Rating (Efficiency) have been renamed to Current (Potential) Energy Label (Score) to make it clearer to the reader whether the energy rating being referred to is a numerical "score" or an alphabetical "label".

Table 3: Feature Formatting

Category	Operation	Entries
Mapped Data		**7,239,549**
Energy Ratings	Ensure that current energy ratings are less than potential.	(890)
	Cap current and potential energy scores to 100.	—
Building Characteristics	Total floor area is between 20 and 400 square meters.	(23,637)
	Property type is not `Park Home`.	(22)
	Property has 1 to 12 habitable rooms.	(3,023)
	Collapse construction age bands post-2007 into one class.	—
	Subsume `Much More (Less) Than Typical` glazed area into `More (Less) Than Typical` categories.	—
	Collapse multi-glaze area into a categorical variables with three classes: High (≥ 66.6), Low(≤ 33.3), and Medium.	—
	Floor (Cap) low energy lighting proportion at 0 (100).	—
Transaction Characteristics	Remove entries for which tenure is unknown.	(10,669)
	Cluster categories with few or related entries.	—
Price	Restrict entries with price per unit area within the 0.25% and 99.75% quantiles.	(35,975)
Environmental Metrics	For current (potential) environmental impact, restrict entries within the 0.01% and 99.99% quantiles. Divide current (potential) carbon emission and energy consumption metrics by the total floor area and restrict the sample to 99.98% quantile range. Floor (or cap) aforementioned metrics between 0 and 100.	(2,249)
	Remove entries for which potential environmental impact is greater than current, and for which current carbon emissions and energy consumption are greater than potential.	(41,179)
Utility Costs	Restrict current (potential) heating, lighting and hot water costs to the 99.98% quantile range.	(3,480)
	Remove entries for which current utility costs are less than potential costs with a 2.5% tolerance. Format entries inside the tolerance range so that current equals potential.	(95,780)
Formatted Data		**7,022,645**

Operations are listed in the order in which they are carried out. Because formatting Price, Environmental Metrics, and Utility Costs involve operations that eliminate entries outside a given quantile range, they are carried out at the end to minimise the loss of entries.

3 Estimation of Energy Premium

In this section, we explore whether properties with higher energy efficiency ratings command a premium in the UK residential real estate market. Starting 01 October 2008, it became a legal requirement in the UK for homeowners to hold a valid Energy Performance Certificate (EPC) when selling or leasing out a property. Each property receives a numeric energy efficiency *score* from 1 (least sustainable) to 100 (most sustainable). These scores are then sorted into seven alphabetical *labels*. Government policies in European countries aimed at improving the energy profile of the housing stock, such as Clean Growth Strategy (CGS) (UK Government, 2017a) and Minimum Energy Efficiency Standards (MEES) (UK Government, 2017b) in the UK, are typically based on a threshold or a group of these labels. Therefore, literature on estimation of energy premium further aggregates these labels into *groups* or *classifications*, such as "green" or "brown". Table 4 shows four common levels of aggregation deployed in the literature.

Table 4: Energy Rating Aggregations

Energy Score	Energy Label	Energy Label Group	Energy Classification
92+	A		Green
81-91	B	BC	
69-80	C		
55-68	D	DE	Brown
39-54	E		
21-38	F	FG	
01-20	G		

This table shows the four common levels of aggregation deployed in the literature. The first two columns represent the official ratings and labels used by the UK Department for Levelling Up, Housing & Communities. The second two columns represent classifications typically used in the literature. Green (Brown) labels are often referred to Safe (Unsafe) or Sustainable (Unsustainable).

3.1 Methodology

We use *hedonic regression models* to estimate energy premium. A hedonic pricing model is a revealed preference method that assumes that (i) the value of a composite object can be decomposed into its constituent components and the external factors that affect the value of the object, and that (ii) the market values these individual components and factors. These models are widely applied in real estate finance to analyse property values, where the market price is determined by a combination of structural characteristics of the dwelling (e.g., floor area, number of habitable rooms, age of property, etc.) and the socio-economic and environmental characteristics of the surrounding area (e.g., ambient temperature, proximity to green spaces, quality of schools, access to transportation hubs, etc.). Thus, a hedonic pricing model can be used to determine to which extent each structural characteristic or external factor impacts property prices. See Baranzini, Ramirez, Schaerer, and Thalmann (2008) for an overview on hedonic methods in residential real estate markets.

While hedonic models can be fairly general and nonlinear, we *assume* that the marginal contribution of each constituent component and external factor to the overall property price is *linear and additive*. This assumption enables us to use a *linear regression* for estimation, which yields several advantages. First, ordinary least squares (OLS) is the standard method of estimating energy premium in the literature (e.g., Brounen and Kok, 2011; Högberg, 2013; Hyland, Lyons, and Lyons, 2013; Ramos, Pérez-Alonso, and Silva, 2015; Jensen, Hansen, and Kragh, 2016; Cajias, Fuerst, and Bienert, 2019; Fuerst, Haddad, and Adan, 2020) which allows us to directly relate our results to those obtained in prior studies. Second, the statistical properties of the estimates obtained are well understood. Third, augmenting the regression model to include additional covariates, interaction effects, or time-varying components is relatively straightforward. This makes it easier to study how the estimates evolve over time, how they vary with climatic conditions, different intended uses of property, etc. Furthermore, linear regressions can be naturally extended into difference in difference and regression discontinuity methods that are used in this paper to study policy impact. Lastly, Figure 6 reveals that price per unit area is approximately log-normal. Therefore, using logarithm of price per unit area as the target variable helps us assume that the error term follows a conditional normal distribution (in addition to being zero-mean and homoskedastic) and that we obtain the most precise unbiased estimates. This allows us to compare our estimates to those obtained from all unbiased estimators, and not only linear ones.

Figure 6: Distribution of Dependent Variable

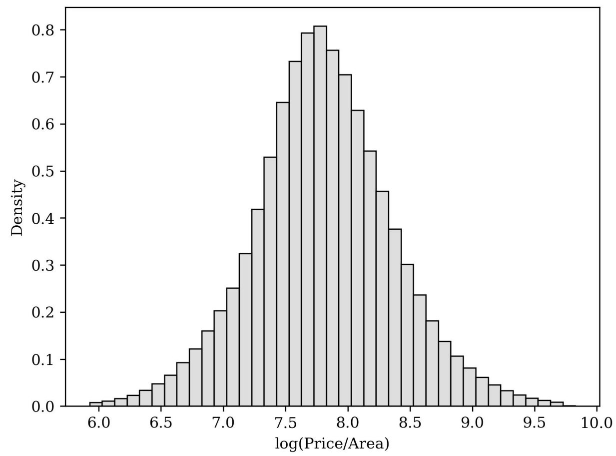

Our notation is adapted from Eichholtz, Korevaar, Lindenthal, and Tallec (2021). The unit of observation in our hedonic regression model is transaction $i \in I$ (where $|I|$ is the number of entries in the dataset) of property h, in region r (also referred to as district or local authority), at time t. We denote the numerical energy efficiency *score* of dwelling h associated with transaction i as S_{ih}, and represent the corresponding *label* as $\texttt{Label}(i,h)$, the corresponding *group* as $\texttt{Group}(i,h)$, and the corresponding *classification* as $\texttt{Class}(i,h)$. Therefore, for a property with $\mathrm{S}_{ih} = 73$, we have $\texttt{Label}(i,h) =$ "C", $\texttt{Group}(i,h) =$ "BC", and $\texttt{Class}(i,h) =$ "Green", as per Table 4. The dependent variable in our model is the logarithm of transaction price per unit area of the underlying property, denoted by $\log{(P/A)}_{ihrt}$. We run four hedonic regression specifications, one for each aggregation level of energy efficiency ratings:

$$\log{(P/A)}_{ihrt} = \alpha_r + \delta_t + \mathrm{S}_{ih}\beta_{\texttt{Score}} + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt}, \tag{1a}$$

$$\log{(P/A)}_{ihrt} = \alpha_r + \delta_t + \beta_{\texttt{Label}(i,h)} + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt}, \tag{1b}$$

$$\log{(P/A)}_{ihrt} = \alpha_r + \delta_t + \beta_{\texttt{Group}(i,h)} + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt}, \tag{1c}$$

$$\log{(P/A)}_{ihrt} = \alpha_r + \delta_t + \beta_{\texttt{Class}(i,h)} + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt}. \tag{1d}$$

The term α_r represents the region r specific fixed effect, and δ_t represents the time t (or year) specific fixed effect. We control for the physical attributes of dwelling h by introducing a 9×1 vector of building properties denoted by B_h, and control for the characteristics of transaction i by including a 2×1 vector of controls denoted by T_i; see Table 2 for an enumeration and description of building properties and transaction characteristics. The 7×1 vector of multiple deprivation indices associated with region r in year t is denoted by MDI_{rt}, and the associated degree days measurement is represented by the scalar DD_{rt}. In a hedonic model, structural characteristics (B_h and T_i) and external factors (MDI_{rt} and DD_{rt}) are known as *hedonic controls*. The parameters associated with the hedonic controls are denoted by $\theta \in \mathbb{R}^9$, $\gamma \in \mathbb{R}^2$, $\nu \in \mathbb{R}^7$, and $\omega \in \mathbb{R}$.

The parameters of interest are $\beta_{\texttt{Score}}$, $\beta_{\texttt{Label}(i,h)}$, $\beta_{\texttt{Group}(i,h)}$, and $\beta_{\texttt{Class}(i,h)}$, where $\beta_{\texttt{Score}} \in \mathbb{R}$ is the coefficient of S_{ih} in Equation (1a), whereas $\beta_{\texttt{Label}(i,h)}$, $\beta_{\texttt{Group}(i,h)}$, and $\beta_{\texttt{Class}(i,h)}$ can be interpreted as the energy efficiency label, group, and classification specific fixed effects in Equations (1b), (1c), and (1d) respectively. Lastly, we denote the residual error term as ε_{ihrt}. We are most interested in estimate of $\beta_{\texttt{Score}}$, as all other levels of aggregation of energy ratings are derived from the numerical energy efficiency score S_{ih}.

3.2 Results

3.2.1 Energy Premium

Table 5 presents the parameter estimates obtained for different levels of aggregations of energy ratings across the four hedonic regression specifications, with heteroskedasticity-robust standard errors in parenthesis.[12] Columns (a_1) and (a_2) provide estimates for $\beta_{\texttt{Score}}$ in Equation (1a). Column (a_1) excludes properties with energy efficiency labels A (score ≥ 92) and

[12] *Note on Implementation.* We can estimate the parameters in Equation (1) by introducing a constant (intercept term) and dropping one *level* (dummy variable) for each fixed effect and categorical variable. We interpret α_r as a separate intercept for each region and therefore drop the intercept and retain all levels for each categorical variable; see Angrist and Pischke (2008, Chapter 5.1). In practice, we find that both specifications lead to exactly the same conclusions, as the theory suggests. However, since Equation (1) has multiple categorical variables and fixed effect terms, it is hard to assign the "distributed intercept" interpretation to any one set of categorical variables; for instance, we think of δ_t as time-specific intercepts. Essentially, opting to drop the constant term would "distribute" the effect of the intercept across categorical variables, changing the magnitude of estimates of categorical variables but will preserve the *difference in levels*, which is what we are interested in. Parameter estimates corresponding numerical features remain unchanged. Therefore, the difference in parameter estimate corresponding to energy efficiency labels A and D remains the same. We choose to not use the constant term in our hedonic regression models as otherwise, the resulting parameter estimates of categorical variables become harder to read and compare visually.

G (score $\leq$ 20) while column (a_2) includes them.[13] Columns (b), (c), and (d) provide estimates for $\beta_{\texttt{Label}(i,h)}$, $\beta_{\texttt{Group}(i,h)}$, and $\beta_{\texttt{Class}(i,h)}$ corresponding to Equations (1b), (1c), and (1d) respectively. We obtain an R^2 upwards of 78%, and estimates that are significant at a 99.9% confidence level, for each specification. It may appear at first that the estimates for energy premia are inconsistent across different model specifications. In Section IA.2 of the Internet Appendix, we discuss how the estimates are consistent across specifications and with those observed in the existing body of literature. In this section, we focus our discussion on $\beta_{\texttt{Score}}$ in column (a_1) for Equation (1a).

We find that a unit increase in the numerical energy efficiency score is associated with a 0.18% increase in price per unit area. We define *range of energy premium* as the difference between the price of a property with the highest energy efficiency rating minus that of the lowest rating, holding all else equal. The range of energy premium that we observe (17.82%) is similar to that reported by Hyland, Lyons, and Lyons (2013) (19.5%), and modest when compared to Cajias and Piazolo (2013) (45%). Notwithstanding, the high magnitude of energy premium observed raises the concern that $\beta_{\texttt{Score}}$ is capturing the effect on an unobserved variable, such as *quality* or *condition* of a dwelling at the time of transaction.

If energy efficiency scores are exogenous, then the expectation of residual ε_{ihrt} conditional on S_{ih} should be zero for all S_{ih}. Figure IA.7 in Appendix IA.2 reveals that $\mathbb{E}[\varepsilon_{ihrt}|S_{ih}] \neq 0$ for properties with energy labels A (score $\geq$ 92) and G (score $\leq$ 20). Therefore, restricting our sample to properties with labels F (score $\geq$ 21) through B (score $\leq$ 91) yields an unbiased estimate for $\beta_{\texttt{Score}}$, which is reported in column (a_1) in Table 5. Indeed, when we include properties with all labels, we find that our estimate for $\beta_{\texttt{Score}}$ reported in column (a_2) is biased upwards from 0.18% to 0.21%. Furthermore, we find that energy premia is not only robust to alternative model specifications, but also *persists* in subsamples sorted by energy efficiency labels (Section IA.2 of the Internet Appendix), period in which the transaction

[13] We exclude properties with labels A (score $\geq$ 92) and G (score $\leq$ 20) while estimating $\beta_{\texttt{Score}}$ in Equation (1a) due to the following reason. Figure IA.1 in Section IA.1.3 of the Internet Appendix shows that a very small proportion of entries have labels A or G in the data. These entries are not uniformly distributed across years. Column (b) in Table 5 shows that properties with labels A (G) command a significant premium (discount). Therefore, when we run regressions on year-wise subsamples to estimate $\beta_{\texttt{Score},t}$ in Equation (3), the estimates obtained for years with higher proportion of entries with labels A or G are biased upwards; see Section IA.2 of the Internet Appendix for more discussion.

Table 5: Energy Premium Estimates

	(a1)	(a2)	(b)	(c)	(d)
Current Energy Score	0.0018	0.0021			
	(0.0000)	(0.0000)			
Current Energy Label : A			0.5402		
			(0.0074)		
Current Energy Label : B			0.4978		
			(0.0014)		
Current Energy Label : C			0.4956		
			(0.0012)		
Current Energy Label : D			0.4794		
			(0.0012)		
Current Energy Label : E			0.4536		
			(0.0012)		
Current Energy Label : F			0.4240		
			(0.0013)		
Current Energy Label : G			0.3271		
			(0.0015)		
Current Energy Label Group : BC				1.0265	
				(0.0004)	
Current Energy Label Group : DE				1.0056	
				(0.0004)	
Current Energy Label Group : FG				0.9434	
				(0.0005)	
Current Energy Classification : C+ (Green)					1.4108
					(0.0005)
Current Energy Classification : D- (Brown)					1.3883
					(0.0004)
N	5397985	5451054	5451054	5450155	5451054
Adj. R^2	0.7875	0.7870	0.7869	0.7864	0.7858

This table reports the estimates for energy premia with heteroskedasticity robust standard errors in parenthesis; all estimates are significant at a 99.9% confidence level. Columns (a_1) and (a_2) provide estimates for $\beta_{\texttt{Score}}$ in Equation (1a). Column (a_1) excludes properties with energy efficiency labels A (score $\geq$ 92) and G (score $\leq$ 20) while column (a_2) includes them. Columns (b), (c), and (d) provide estimates for $\beta_{\texttt{Label}(i,h)}$, $\beta_{\texttt{Group}(i,h)}$, and $\beta_{\texttt{Class}(i,h)}$ corresponding to Equations (1b), (1c), and (1d) respectively.

occurred (Section 4.1), level of urbanisation (Section IA.3 of the Internet Appendix), and property types (Section IA.3 of the Internet Appendix). Section 3.2.2 and Section 3.2.3 show that coefficients of various covariates and fixed effects in Equation (1) make economic sense, which lends credibility to our results. In addition, we find that the discount rate implied by our estimate for energy premium is economically plausible (Section 3.2.4) and that the premium is higher in regions where an increase in energy ratings is expected to lead to higher utility savings. Finally, our analysis in Section 6.2 reveals that the marginal cost (0.36%) associated with undertaking a unit improvement in numerical energy efficiency score is two times the observed premium (0.18%). Therefore, if anything, the observed energy premium is too low.

3.2.2 Hedonic Controls

Structural characteristics (B_h and T_i) and external factors (MDI_{rt} and DD_{rt}) in Equation (1) are known as *hedonic controls*. We report the estimates for select building properties and transaction controls in Table 6, and those for degree days and the seven multiple deprivation indices in Table 7.

The results act as a robustness check across specifications, and we see that the coefficients of numerical features (e.g., Total Floor Area), and the differences in levels of categorical features (e.g., New, Tenure), are consistent across specifications (a) through (d). It is reassuring to find that our estimates make economic sense. For a $10\mathrm{m}^2$ increase in the floor area, we see that the price per unit area declines by 2.8%, which suggests economies of scale. New properties sell at a 3.57% premium relative to old ones, and properties that are purchased with the intent of occupation command a 5.48% premium over those with an intent to lease out. We observe that the value of a property declines with age, except for those constructed before the 20^{th} century, which sell at a premium. Indeed, period properties may command higher market valuations due to their unique character and history. Unsurprisingly, bungalows (+28.2%) and houses (+13.6%) are more expensive than flats.

For each unit increase in the degree days measure, the price paid per unit area declines by 0.08%, ceteris paribus. Therefore, properties in colder climates are more economical, perhaps due to a higher demand for housing in regions with more temperate ambient conditions. A lower (normalised) deprivation index rank implies higher deprivation. We find that estimates corresponding to income, employment, and education and positive, suggesting that affluent neighbourhoods command a premium. The estimate for crime is positive, whereas that

Table 6: Select Building and Transaction Control Estimates

	(a1)	(a2)	(b)	(c)	(d)
Total Floor Area	-0.0028	-0.0028	-0.0028	-0.0028	-0.0028
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Habitable Rooms	0.0193	0.0192	0.0191	0.0192	0.0193
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Property Type : Bungalow	1.0275	1.0251	0.9844	0.9228	0.8782
	(0.0005)	(0.0005)	(0.0005)	(0.0004)	(0.0004)
Property Type : Flat	0.7456	0.7430	0.7069	0.6472	0.6033
	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Property Type : House	0.8814	0.8792	0.8390	0.7780	0.7338
	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Property Type : Maisonette	0.7279	0.7254	0.6874	0.6276	0.5837
	(0.0007)	(0.0007)	(0.0007)	(0.0007)	(0.0007)
New : No	1.6733	1.6688	1.5905	1.4691	1.3810
	(0.0010)	(0.0010)	(0.0010)	(0.0010)	(0.0010)
New : Yes	1.7090	1.7039	1.6272	1.5064	1.4181
	(0.0014)	(0.0014)	(0.0014)	(0.0013)	(0.0013)
Construction Age Band : 1900 prior	0.3263	0.3284	0.3084	0.2818	0.2614
	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Construction Age Band : 1900-1929	0.2475	0.2488	0.2287	0.2027	0.1837
	(0.0003)	(0.0003)	(0.0004)	(0.0003)	(0.0003)
Construction Age Band : 1930-1949	0.2794	0.2804	0.2619	0.2374	0.2196
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Construction Age Band : 1950-1966	0.2625	0.2629	0.2457	0.2230	0.2063
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Construction Age Band : 1967-1975	0.2612	0.2612	0.2444	0.2223	0.2066
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Construction Age Band : 1976-1982	0.2868	0.2859	0.2708	0.2502	0.2358
	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Construction Age Band : 1983-1990	0.3222	0.3208	0.3065	0.2870	0.2731
	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Construction Age Band : 1991-1995	0.3493	0.3478	0.3335	0.3143	0.3003
	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0005)
Construction Age Band : 1996-2002	0.3575	0.3549	0.3437	0.3244	0.3102
	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0005)
Construction Age Band : 2003-2006	0.3456	0.3418	0.3361	0.3153	0.3002
	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0005)
Construction Age Band : 2007 onwards	0.3441	0.3397	0.3381	0.3171	0.3017
	(0.0008)	(0.0008)	(0.0008)	(0.0008)	(0.0008)
Tenure : Owner Occupied	1.1910	1.1873	1.1357	1.0550	0.9963
	(0.0007)	(0.0007)	(0.0007)	(0.0007)	(0.0007)
Tenure : Rental (Private)	1.1362	1.1332	1.0811	1.0000	0.9410
	(0.0006)	(0.0006)	(0.0006)	(0.0006)	(0.0006)
Tenure : Rental (Social)	1.0552	1.0522	1.0010	0.9205	0.8617
	(0.0010)	(0.0010)	(0.0010)	(0.0010)	(0.0010)

for health deprivation is negative.[14] One explanation is that more affluent neighbourhoods have better access to hospitals and law enforcement, and therefore, a higher number of cases are reported. An alternative explanation could be that the population in inexpensive neighbourhoods is younger, and therefore, healthier. Similarly, crime could be more prevalent in urban areas, which also tend to be more expensive. Consecutively, a negative coefficient for living environment indicates poorer air quality and higher road accidents in more congested (but expensive) areas. The negative estimate for housing barrier can be explained away by the fact that lower housing affordability is itself a *barrier* to housing.

Table 7: Deprivation Indices and Degree Days Estimates

	(a1)	(a2)	(b)	(c)	(d)
Degree Days	-0.0008	-0.0008	-0.0008	-0.0008	-0.0009
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Income Index	0.1681	0.1678	0.1677	0.1678	0.1680
	(0.0015)	(0.0015)	(0.0015)	(0.0015)	(0.0015)
Employment Index	0.0109	0.0121	0.0121	0.0121	0.0124
	(0.0014)	(0.0014)	(0.0014)	(0.0014)	(0.0014)
Health Deprivation Index	-0.0483	-0.0478	-0.0481	-0.0488	-0.0501
	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0011)
Education Index	0.4293	0.4304	0.4298	0.4295	0.4298
	(0.0008)	(0.0008)	(0.0008)	(0.0008)	(0.0008)
Crime Index	0.0585	0.0591	0.0590	0.0590	0.0584
	(0.0006)	(0.0006)	(0.0006)	(0.0006)	(0.0006)
Housing Barrier Index	-0.0263	-0.0279	-0.0276	-0.0269	-0.0250
	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0005)
Living Environment Index	-0.0151	-0.0156	-0.0145	-0.0124	-0.0096
	(0.0006)	(0.0006)	(0.0006)	(0.0006)	(0.0006)

This table reports the estimates for degree days and multiple deprivation indices in Equation (1) with heteroskedasticity robust standard errors in parenthesis; all estimates are significant at a 99.9% confidence level.

3.2.3 Fixed Effects

Similar to estimates for hedonic controls, the estimates for region and time fixed effects are consistent across model specifications and realistic from an economic standpoint. Figure 7 plots the evolution of prices implied by the time (or year) fixed effects against the property

[14] The health deprivation index is not based on access to healthcare, but rather, derived from statistics on mortality, morbidity, disabilities, and mood or anxiety disorders.

Figure 7: Time Fixed Effects

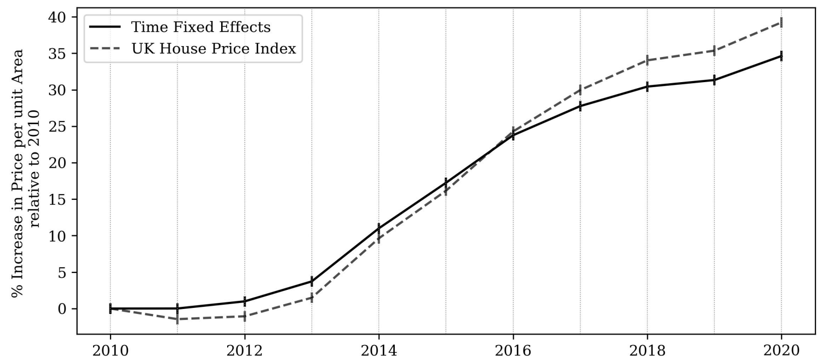

The solid black line tracks the evolution of property prices, relative to 2008, implied by the time fixed effects estimates obtained from Equation (1). For example, property values in 2014 were 18.95% higher than those in 2008, on average. The dashed line plots the evolution of property prices, relative to 2008, implied by the property price index published by the HM Land Registry.

price index published by the HM Land Registry, and it is reassuring to observe that the former closely tracks the latter. We observe that property prices increased by 34.63% over the duration of our sample. Figure 8 illustrates the region-specific fixed effects; we observe that properties in central London are three times as expensive as the national average.

Figure 8: Region Fixed Effects

This heatmap illustrates the region (local authority) fixed effect estimates for Equation (1). The levels marked on the legend are reported in log(Price/Area). Therefore, properties in regions corresponding to a value of 1.5 (e.g., boroughs in London) are three times as expensive as those in regions corresponding to a value of 0.5, on average.

3.2.4 Implied Social Discount Rate (SDR)

If we assume that homeowners are *rational*, the energy premium computed in Section 3.2.1 should be equal to the present value of expected annual energy savings[15], that is:

$$\beta_{\texttt{Score}} = \frac{m^u}{r - g}, \tag{2}$$

where $\beta_{\texttt{Score}}$ is the energy premium obtained from Equation (1a), m^u is the marginal utility savings metric (computed in Section 5.1), r represents the discount rate, and g represents the rate of growth in utility expenditures.[16] We find that the marginal annual energy savings in utility bills from each unit improvement in energy efficiency score (m^u) is 0.0087%. Given that the average growth rate of energy expenditures over the duration of our sample is 3.33%, an energy premium of 0.18% implies that the market discounts future energy savings at a rate of 8.16%, or a *net discount rate* ($r - g$) of 4.83%.[17]

We can think of this metric as the empirically observed *social discount rate* (SDR) that homeowners use to appraise investments in energy efficiency. An SDR is the discount rate used to value future costs and benefits in economic evaluations of public interventions. In comparison, the SDR adopted by the Office of National Statistics, as per The Green Book published by HM Treasury (UK Government, 2022a), is 3.5%, which is 27% lower than what we have estimated. What the correct SDR should be is an ongoing and contentious area of debate in the economics of climate change. Small changes in SDR can dramatically alter the present value of climate abatement investments that pay off over long horizons (Giglio, Maggiori, Rao, Stroebel, and Weber, 2021).

[15] We think of expected marginal annual utility savings as recurring dividend payments and then apply the perpetuity formula to obtain the present value of marginal savings.

[16] *Computing Growth Rate of Energy Bills.* We download Consumer Price Indices (CPI) for gas, electricity, and water supply published by the Office of National Statistics. We then compute the proportion of expenditure that heating (74.5%), lighting (9.3%), and hot water (16.2%) costs constitute out of the total utility bills for an average household in our sample. We *weight* the gas, electricity, and water supply indices by heating, lighting, and hot water proportions respectively to create a composite energy expenditure index. Finally, we compute the year-on-year growth rate for the composite index from 2008 to 2020, and then take an average.

[17] We use *median* marginal energy savings (0.0087%) to compute the net discount rates as it is more robust to outliers. Using *mean* marginal energy savings (0.0108%) results in a net discount rate of 5.55%, which is quite close to that computed by Nordhaus (2007) (5.5%).

On one hand, setting the SDR too high could preclude socially desirable public projects from being undertaken (Zhuang, Liang, Lin, and De Guzman, 2007).[18] On the other hand, setting the SDR too low can result in a suboptimal policy design. If our empirically observed net discount rate of 4.83% is correct, then government policies based on an SDR of 3.5% will risk being overly optimistic. For instance, a program which offers subsidies to homeowners to cover the difference between the cost of energy efficiency improvements and the present value of future energy savings will overestimate the latter, and therefore offer lower subsidies than what are required to incentivise homeowners to undertake energy improvements. This might explain why the 2014 Green Deal, a policy where homeowners could claim up to £7,600 for improving the energy profile of their properties, failed to gain traction.[19]

Why is the SDR employed by the regulator lower than that observed empirically? Perhaps, we can generate some insight by considering the potential difference in the attitudes of the regulator and homeowners towards computation of the SDR. Starting 2003, the HM Treasury adopts the Social Rate of Time Preference (SRTP) as the SDR , which is calibrated according to the Ramsey Rule (Ramsey, 1928).[20] SRTP is driven by three parameters: (i) pure rate of time preference, (ii) consumption growth, and (iii) the elasticity of marginal utility. Groom and Maddison (2019) remark that economists disagree on the interpretation these parameters; the resultant variation in estimation procedures can lead to very different SDRs. The Green Book (UK Government, 2022a) argues that the parameter values used to calibrate SRTP fall within the range of those found in the literature. However, the regulator may display a bias in selecting parameter values that yield lower SDRs, as it makes public undertakings look more attractive. More importantly, the regulator may strategically opt for a lower pure rate of time preference to balance the welfare of present and future generations.[21] For example, Stern (2007) uses a pure rate of time preference of 0.1% to arrive at an

[18] Recently, the a 17.5% set by European Commission sparked intense debate because it made energy efficiency policies in the building sector look very unattractive.

[19] *The 2014 Green Deal*: On 01 May 2014, the UK Government announced the new Green Deal Home Improvement Fund (GDHIF). Starting 9 June 2014, households in England and Wales were eligible to claim up to 7,600 for improving the energy efficiency of their homes. We do not analyse the 2014 Green Deal in this paper.

[20] SRTP can be interpreted as the rate at which society is willing to postpone a marginal unit of current consumption in exchange for more future consumption; see Zhuang, Liang, Lin, and De Guzman (2007).

[21] This is an ethical argument. A pure rate of time preference equal to zero would imply that we value the social welfare of future generations exactly the same as that of the current generation.

SRTP of 1.4%. In contrast, Nordhaus (2007) and Groom and Maddison (2019) empirically argue for higher SRTPs – 5.5% and 4.2% respectively – that are closer to our estimate.

Unlike the regulator who wishes to optimise social welfare, homeowners can be expected to assume a myopic perspective and view energy upgrades as private investments. Therefore, the SDR used by homeowners should reflect their opportunity cost (i.e., the marginal rate of return on investment in the private sector), which is typically higher than SRTP. Zhuang, Liang, Lin, and De Guzman (2007) remark that in the absence of market distortions, this is equivalent to the marginal social rate of return on private investment, also termed Social Opportunity Cost of Capital (SOC). Prior to 2003, HM Treasury used SOC as their measure of SDR, which was indeed higher than the 3.5% SRTP, ranging from 5% to 10%.

An alternative reasoning is that homeowners are exposed to project- or property-specific idiosyncratic risks when upgrading the energy profile of their homes, which cannot be diversified away. This results in a higher SDR, as SRTP does not does factor in project-specific risks and costs of raising capital. One can also call for industry adjusted SDRs, as opposed to a blanket 3.5% discount rate for all public undertakings. Giglio, Maggiori, Rao, Stroebel, and Weber (2021) show that returns to real estate are positively correlated with consumption, thus increasing the riskiness of the asset class. Thus, the SDR used for public interventions in the building sector could benefit from an upwards revision. For public undertakings that do not rely on market participation, one can argue that the choice of SDR only impacts the fiscal attractiveness of the project. However, when market participation is necessitated – such as in the case of 2014 Green Deal described above and improvement of energy profiles of the housing stock in general – it becomes important for the regulator to take into account the empirically observed SDR employed by the homeowners, notwithstanding the source of the discrepancy.

We conclude this section by defining *green premium* as the difference between energy premium and the present value of marginal energy savings, which can be thought of as the value of the utility that homeowners derive from the *sustainability* of their dwelling. If the rate at which future energy savings should be discounted is greater than 4.83%, then the present value of future savings will be lower than the observed energy premium, indicating a positive green premium. Conversely, if we believe the correct discount rate is lower than 4.83%, then green premium is negative, the market is irrational, and prospective buyers can arbitrage by purchasing dwellings with higher energy ratings.

4 Heterogeneity in Energy Premium

In this section, we investigate heterogeneity in energy premium along five channels: (i) time, (ii) ambient temperature (as measured by degree days), (iii) tenure, (iv) level of urbanisation, and (v) property type. Section 4.1, 4.2, and 4.3 discuss the first three channels respectively, while the remainder of the two are presented in Section IA.3 of the Internet Appendix. Our main motivation to study these channels is to better understand to what extent homeowners factor utility savings into their investment decisions. Are variations in the magnitude of expected utility savings, or the heterogeneity in the ability of homeowners to recoup their investments in energy efficiency reflected in the observed energy premia? If so, what are the implications of heterogeneity in energy premia for policy makers and market participants?

4.1 Time

4.1.1 Methodology

Estimating Equation (1a) provides a static estimate for $\beta_{\texttt{Score}}$ over the entire duration of the sample. However, we are also interested in learning how the energy premium changes over time. We use two methods to track the evolution of $\beta_{\texttt{Score}}$. First, we construct subsamples for each of the 44 quarters starting 2010-Q1 until 2020-Q4, and run a separate hedonic regression for each subsample; that is, for each t = 2010-Q1, 2010-Q2,..., 2020-Q3, 2020-Q4, we estimate:

$$\log{(P/A)_{ihrt}} = \alpha_r + \mathrm{S}_{ih}\beta_{\texttt{Score},t} + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt}, \tag{3}$$

where the unit of observation is transaction $i \in I_t$ (where $|I_t|$ is the number of entries in the dataset for quarter t) of property h, in region r, at time t; all other symbols have exactly the same meaning as those in Equation (1a); with two differences. First, we drop the time fixed-effects term δ_t, as the t is now fixed in each of the 44 individual regressions.[22] Second, we include the subscript t in parameter associated with energy efficiency score, $\beta_{\texttt{Score},t}$, as it is now specific to the period for which the regression is run.

Our second approach is to use a single hedonic regression over the full sample, but allow energy premium to vary over time by introducing an interaction term between energy efficiency score and an indicator variable for time, as follows:

$$\begin{aligned} \log{(P/A)_{ihrt}} = \alpha_r + \delta_t + \mathrm{S}_{ih}\beta_{\texttt{Score}} + (\mathbb{1}_t.\mathrm{S}_{ih})\beta_{\texttt{Score},t} \\ + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt}, \end{aligned} \tag{4}$$

where the unit of observation is transaction $i \in I$ (where $|I|$ is the number of entries in the dataset) of property h, in region r, at time t; all other symbols have exactly the same meaning as those in Equation (1a). The indicator variable for quarter t is denoted by $\mathbb{1}_t$. Therefore, the *time invariant* component of the effect of energy efficiency score on property prices is captured by $\beta_{\texttt{Score}}$, and $\beta_{\texttt{Score},t}$ captures the *time varying* component. As an example, the estimate of energy premium for quarter t =2015-Q2 can be computed as $\beta_{\texttt{Score}} + \beta_{\texttt{Score},t=2015Q2}$. In contrast to the quarter-wise subsample regressions, Equation (4) accounts for the quarter fixed effects δ_t, but forces the estimates for the hedonic covariates to remain constant over the duration of the sample.

4.1.2 Results

Figure 9 plots the estimates obtained for energy premium from the quarter-wise subsample regressions (Equation (3)) and the time-interacted effects model (Equation (4)) over the duration of our sample. For both models, Figure 10 tracks the net discount rates implied by the energy premium estimates and the period-wise marginal energy savings (m_t^u) computed in Section 5.[23] We observe that the energy premium increases from 0.15% to 0.22% from 2010 to 2014, but remains fairly close to the sample average (0.18%) starting 2015. With the exception of first two quarters of 2012 and last quarter of 2020, we find that the implied

[22]Note that we cannot drop the subscript t in Equation (3) because a property h may be transacted twice in the year t; in which case the transactions are uniquely determined by the index i.

[23]We report quarterly net discount rates from 2012 to 2020 as we do not have sufficient information in our data to compute marginal energy savings for 2010 and 2011.

Figure 9: Evolution of Energy Premium

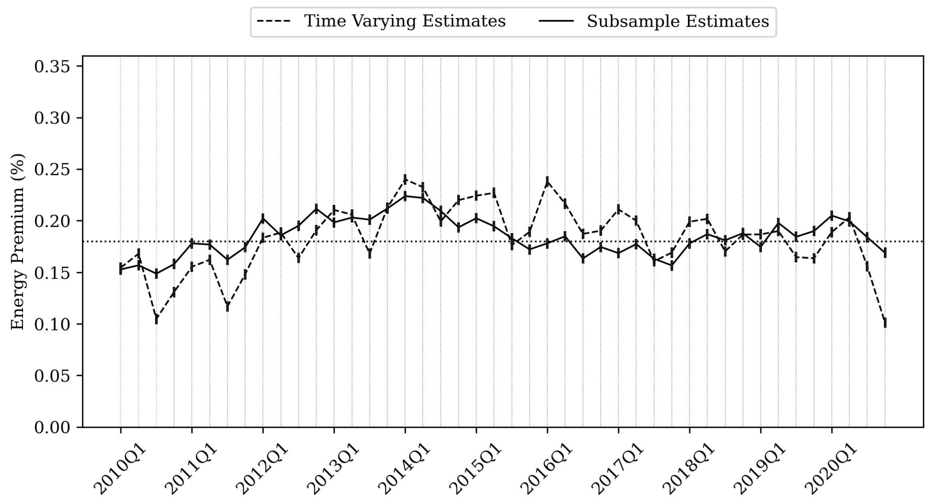

This figure tracks the evolution of energy premium over the duration of our sample. The solid line plots estimates obtained from year-wise subsample regressions in Equation (3). The dashed line plots the estimates obtained from the time-interacted effects model in Equation (4).

Figure 10: Evolution of Implied Discount Rates

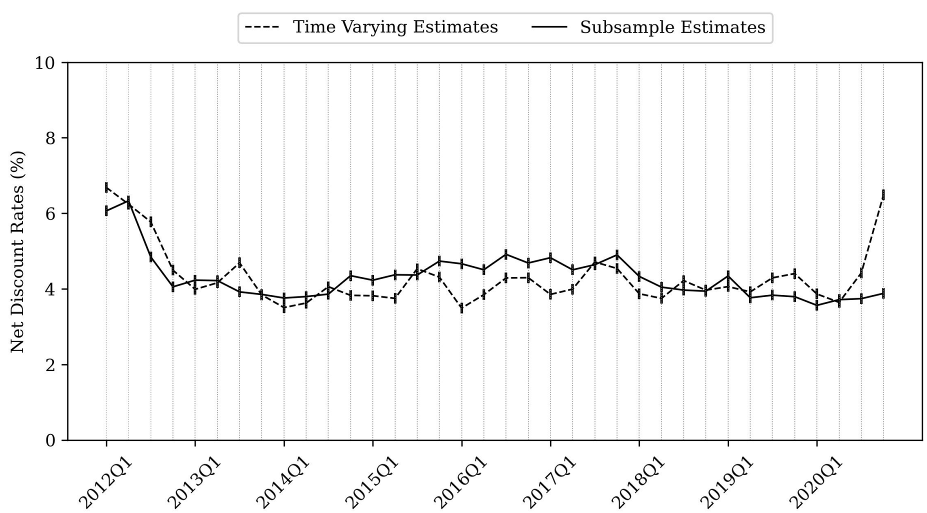

The net discount rate $r_t - g_t$ for quarter t is computed using the perpetuity formula, i.e., $r_t = (m_t^u/\beta_{\text{Score},t}) + g_t$ where g_t is assumed fixed and equal to the mean growth rate (3.33%) over the duration of the sample, . The dashed line uses energy premia estimates from subsample year-wise regressions (Equation (3)) and the solid line uses estimates for energy premium estimates from time-interacted effects model (Equation (4)).

net discount rates remain stable over the duration of our sample. A consistent discount rate suggests that the market is rational and homeowners factor energy expenditures into their investment decisions. Our findings in this section further reinforce our remark in Section 3.2.4 that the net discount rate specified by the HM Treasury (3.5%) runs the risk of being overly optimistic, and could potentially be revised upwards to inform policy design and assessment. Surprisingly, the introduction of MEES in 2015 or CGS in 2017 (see Section 6.1) did not lead to a noticeable increase in energy premium. However, when we track the evolution of energy premium by market segments in Section 4.3, we observe that the premia *weakly* appreciates for the private rental market segment, which is impacted by the regulation.

4.2 Degree Days

4.2.1 Methodology

Our findings in Section 3.2.4 and Section 4.1 indicate that homeowners factor energy expenditures into their investment decisions. If the market is indeed rational, then homeowners should be willing to pay a higher premium in regions where a marginal improvement in the energy efficiency of the dwelling is expected to result in larger reductions in energy bills. We find that for an average household in our sample, heating costs constitute 74.5% of the total energy expenditures. This motivates us to augment our hedonic regression model (Equation (1a)) in Section 3.1 by introducing an interaction term between the energy efficiency score of a property S and the degree days measure DD. Because heating requirements of a property are proportional to the degree days measure corresponding to its location, we expect properties in regions with higher degree days to benefit from higher energy savings for each unit improvement in energy efficiency, resulting in an increase in energy premium. This leads to the following regression:

$$\begin{aligned} \log{(P/A)_{ihrt}} = \alpha_r + \delta_t + \mathrm{S}_{ih}\beta_{\texttt{Score}} + \mathrm{DD}_{rt}\omega + (\mathrm{S}_{ih} \times \mathrm{DD}_{rt})\mu \\ + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \varepsilon_{ihrt}, \end{aligned} \tag{5}$$

where the unit of observation and the meaning of symbols are same as that in Equation (1a). The interaction term between energy efficiency score (of property h associated with transaction i) and the degree days (for region r in year t) is represented by the scalar $S_{ih} \times \mathrm{DD}_{rt}$, and the corresponding parameter is denoted by $\mu \in \mathbb{R}$. Therefore, the estimate for energy premium in a region with degree days measure d can be computed as $\beta_{\texttt{Score}} + d\mu$.

We corroborate our findings by deploying an alternative methodology. For each district in our sample with at least 10,000 transactions, we run separate hedonic regressions (as in Equation (1a) but without the region specific fixed effect term as r is now fixed for each subsample) and collect the region-specific estimates for energy premium. We refer to these estiamtes as *regional betas*. Then, we round the values for degree days in each region to the nearest integer, group regions by their rounded degree day values, and compute the average energy premium by taking the mean of the regional betas for each group. We refer to these estimates as *subsample estimates*.

4.2.2 Results

Figure 11 visualises the energy premium (left y-axis) conditional on degree days (x-axis) estimated by Equation (5) using a bold dashed line.[24] The subsample estimates are plotted by light-grey circular markers. While somewhat dispersed, the grey markers are well aligned along the bold dashed line. The energy premia in regions with degree days equal to 10 is 0.10%, and increases by a factor of 2.5 times, to 0.25%, in regions with degree days equal to 40. Therefore, homeowners pay a higher energy premium in colder regions, where marginal improvements in energy efficiency of dwellings are expected to result in larger reductions in energy bills. This provides compelling evidence that homeowners behave rationally and factor utility savings into their investment decisions. Figure 12 illustrates the regional betas through a heatmap. We observe that energy premium is highest in regions exposed to high altitudes and the sea, and lowest in inland temperature geographies.

Therefore, developers or real-estate private equity firms facing a marginal cost of upgrade c may strategically invest in green retrofits for properties in regions corresponding to degree

[24] Degree days have been normalised between 0 and 100; see Section 2.1.4.

Figure 11: Range of Energy Premia Conditional on Degree Days

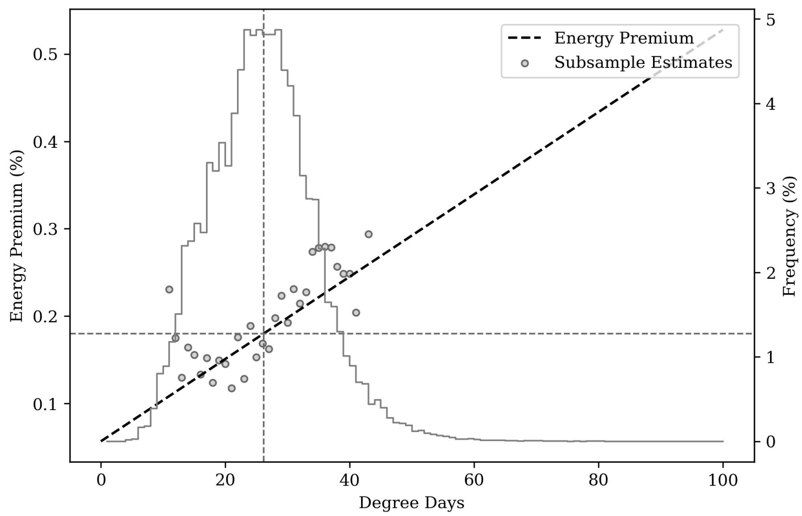

The energy premium (left y-axis) conditional on degree days (x-axis) estimated from Equation (5) is plotted using a bold dashed line. The right y-axis corresponds to the histogram, and measures the proportion of regions with the degree day values marked on the x-axis. The crosshairs show the range of energy premium predicted by the regression line for the region corresponding to the average degree days (i.e., 26) in our sample. The grey dot-markers are obtained by first running separate hedonic regression models for each district and collecting the region-specific estimates for energy premium. Then, we round the values for degree days in each region to the nearest integer, group regions by their rounded degree day values, and compute the average range of energy premium for each group.

Figure 12: Regional Betas

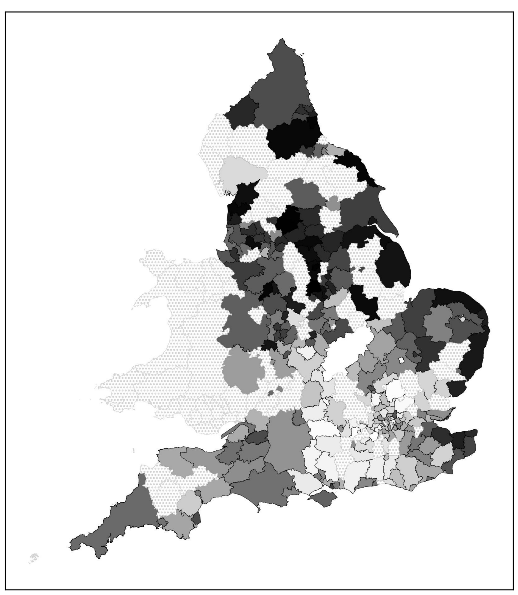

For each district in our sample with at least 10,000 transactions, we run separate hedonic regressions (as in Equation (1a) but without the region-specific fixed effect term as r is now fixed for each subsample) and collect the region-specific estimates for energy premium. Districts marked with a light-grey circular hatch contain less than 10,000 entries.

days d for which $c \leq \beta_{\texttt{Score}} + d\mu$.[25] Similarly, policymakers could offer differential subsidies based on the marginal benefits from investments in energy efficiency in each region.

4.3 Tenure

4.3.1 Methodology

Cajias, Fuerst, and Bienert (2019) advocate that owner-occupied (*buy-to-live*) and private-rental (*buy-to-let*) market segments respond heterogeneously to energy efficiency ratings. The authors attribute this to the different channels through which buy-to-live and buy-to-let landlords recoup their investments in energy efficiency. Homeowners who intend to occupy the dwelling can directly recover the energy premium through savings in utility expenditures; whereas landlords who intend to lease out the property recover the energy premium by passing on the investment costs to tenants through higher rents. Therefore, buy-to-let landlords can be expected to base their willingness to pay for higher energy efficiency on achievable rental values which are net of utility costs, as these are typically covered by tenants.

To study the heterogeneity in energy premium across intended use of the property, we split our sample on residential real estate transactions into two markets by "tenure": *owner occupied* (properties purchased with the intention of occupation) with 4,839,668 observations (88.78%), and *private rental* (properties purchased with the intention of renting out) with 568,565 observations (10.43%). We discard the remaining observations with tenure as *social rental*. We then run period-wise regressions (Equation (3) in Section 3.1) for each subsample to explore how energy premium has evolved for individual market segments over the duration of our sample.[26]

[25] We do not find evidence of such strategic behaviour in our sample.

[26] Fuerst, McAllister, Nanda, and Wyatt (2015b) provide a detailed explanation for why running regressions on subsamples sorted by tenure is preferable over running a single model with an interaction term between energy efficiency score and tenure, as there might be systematic differences in the structural characteristics (Iwata and Yamaga, 2008) and energy profile (Rehdanz, 2007) of the housing stock across the two market segments. In addition, we expect hedonic covariates to be priced differently across the two market segments. For instance, buy-to-let property owners are expected to pay a higher premium for the number of habitable rooms than buy-to-live property owners, as individual rooms can be rented separately to generate more income.

4.3.2 Results

Figure 13 shows the period-wise estimates for energy premium for owner-occupied and private rental market segments. We find that, on average, buy-to-live homeowners pay more than two times the premium (0.195%) for a unit increase in energy efficiency score than buy-to-let landlords (0.09%). The quarterly estimates for the private rental market segments are more volatile, which could be a consequence of smaller period-wise subsamples. We observe that in the second half of the duration of our sample (when policies such as MEES and CGS were introduced; see Section 6.1), energy premium for the owner-occupied market segment remains relatively stable, while that for the private rental market segment *weakly* appreciates. This finding can be explained away by the fact that Minimum Energy Efficiency Standards (MEES) only impacts buy-to-let landlords.

Figure 13: Evolution of Energy Premium by Tenure

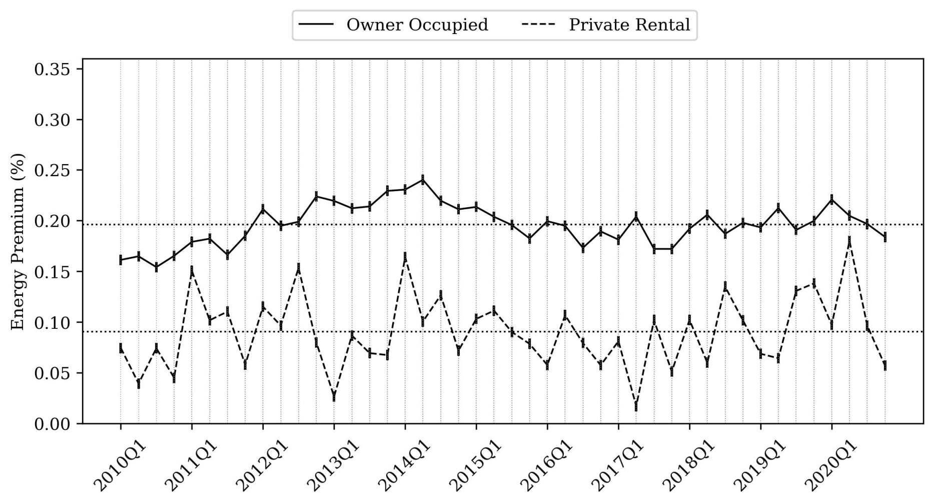

This figure tracks the evolution of energy premium by tenure over the duration of our sample obtained from year-wise subsample regressions in Equation (3). The solid line plots estimates obtained for owner-occupied market segment, and the dashed line plots the estimates obtained for the private rental market segment. Estimates are significant at a 99% confidence level for the owner-occupied market segment, and at a 95% level for the private rental market segment.

Why is energy efficiency priced less in the private rental market segment even though the marginal utility savings from higher energy efficiency are homogenous across both markets over time? In the case of German rental markets, Cajias, Fuerst, and Bienert (2019) argue that *rent caps* prevent landlords from increasing rental prices to recoup their investments in energy efficiency, thereby lowering their willingness to pay. However, the UK rental market has not been subject to such regulations over the duration of our sample.

One potential explanation could be existence of *information asymmetry* in rental markets, whereby landlords are unable to credibly communicate the energy efficiency of various products and features of a property (Gerarden, Newell, and Stavins, 2017) such as appliances (Davis, 2010).[27] However, manufacturers are required to disclose the energy efficiency of their products in the UK; and introduction of EPCs starting October 2008 further reduced this information gap. Therefore, information asymmetries are more likely to originate from the unwillingness of renters to incur *attention costs* required for acquiring home-specific information due to shorter tenures (Clara, Cocco, Naaraayanan, and Sharma, 2022).[28]

Another potential explanation is that renters are irrational (or simply derive lower utility from energy efficiency), which increases the elasticity of the demand curve faced by landlords. Therefore, a lower energy premium paid by landlords may reflect their diminished ability to raise rental prices. If this explanation holds, then we should expect to observe a smaller spread in energy premium paid by the two market segments in urban regions, where demand for rental properties is high and inelastic. Table IA.2 in Section IA.3 of the Internet Appendix shows, however, that this not the case, and the spread in the premium paid by the two market segments is similar across subsamples sorted by level of urbanisation.

Assuming that both landlords and renters are rational, perhaps a convincing argument could be made by considering the existence of *fixed price tariffs*, wherein utility providers charge a fixed rate each month instead of a variable amount that varies according to monthly energy consumption. The cost of fixed price tariffs depends on the conditions of the energy market and on property sizes, but does not factor in the energy profile of individual dwellings. Because utility expenses are typically covered by the tenants, prospective renters considering fixed contracts will be indifferent to dwellings with different energy ratings, as the reductions

[27] Iwata and Yamaga (2008) suggest that a heavier utilisation of dwellings by tenants discourages landlords to make investments in building improvements and energy efficiency, which results in lower market valuations, and a poorer energy profile of the rental housing stock (Rehdanz, 2007).

[28] The 2020-2021 housing survey conducted by the Department for Levelling Up, Housing & Communities found that the typical length of occupation for private-renters is 4.2 years, while that for owner-occupiers is 16 years (UK Government, 2022b).

in energy consumption in a more energy efficient property will not lead to lower utility costs. Therefore, the landlord will be unable to charge higher rents to recoup investments in energy efficiency.

So far, we have benchmarked the energy premium observed in buy-to-let market segment against that observed in buy-to-live market segment. In doing so, we implicitly assumed that the premium paid by buy-to-live homeowners is correct. However, one could consider the reverse. In their paper, which investigates whether exposure to sea-level rise is capitalised in coastal property prices, Bernstein, Gustafson, and Lewis (2019) classify non-owner occupied property owners as *sophisticated investors*, and argue that landlords purchasing a dwelling for investment (or as a second home) make more informed and rational investment decisions (i.e., they exhibit fewer biases in their investment behaviour). This assumption helps them reconcile their findings that when subjected to climate risk, the discount in property prices observed in the non-owner occupied market segment is substantially higher. Given that the average growth rate of energy expenditures over the duration of our sample is 3.33%, and that the marginal upgrade cost (computed in Section 5) is homogenous across owner-occupied and private rental market segments, the rate used by buy-to-live homeowners to discount future utility savings is 4.46% while that used by buy-to-let landlords is 9.66%. If we *assume* that 9.66% is closer to the correct discount rate, then we could conclude that buy-to-live homeowners pay a substantial green premium because they either (i) derive utility from the sustainability of the dwelling they intend to occupy, or (ii) underestimate the risk of future reductions in energy bills. Therefore, homeowners selling the property can generate higher profits by targeting buyers that intend to occupy the dwelling rather than those who intend to lease it out.

5 Evidence of Green Premium

Sections 3 and 4 provide evidence that homeowners factor expected energy savings into their investment decisions. The energy premium remains stable over time, increases when savings from marginal improvements in energy efficiency are higher, and decreases when investments in energy efficiency are difficult to recoup due to market imperfections. Our main motivation in this section is to investigate whether homeowners derive non-pecuniary benefits from the energy efficiency of their dwellings. If this is true, then energy premium should exceed the present value of marginal energy savings. We refer this potential difference as *green premium*.

Section 5.1 describes the methodology employed for the computation of marginal energy savings, and Sections 5.2 outlines the implications on marginal net discount rates and on the existence of green premium.

5.1 Computation of Marginal Savings

Let each observation in the dataset be indexed by transaction $i \in I$ (where $|I|$ is the number of entries in the dataset) and quarter t = 2012-Q1, 2012-Q2,..., 2020-Q4.[29] We denote the current energy score of observation i by $x_{it} \in \{1, 2, ..., 100\}$ and the potential energy score by $y_{it} \geq x_{it}$, $y_{it} \in \{1, 2, ..., 100\}$. Let u^x_{it} be the sum of current annual energy expenditures (heating, lighting, and hot-water) of the property corresponding to transaction i in period t; and let $u^y_{it} \leq u^x_{it}$ denote the potential annual energy expenditure of the property if it is upgraded from x_{it} to y_{it}. The expected annual utility savings are computed as:

$$\Delta u_{it} = u^x_{it} - u^y_{it}. \tag{6}$$

Next, we represent the marginal saving of upgrading from energy efficiency score $s \in \{1, 2, ..., 99\}$ to $s+1$ as $m^u(s)$, and the marginal savings from s to $s+1$ specific to transaction i at time t as $m^u_{it}(s)$. We *assume* that $m^u_{it}(s)$ is uniform between x_{it} and y_{it}, and obtain marginal savings for each observation between its current and potential energy efficiency score as follows:

$$m^u_{it}(s) = \frac{\Delta u_{it}/p_{it}}{y_{it} - x_{it}}, \quad x_{it} \leq s < y_{it}, \tag{7}$$

where p_{it} denotes the transaction price. Note that $m^u_{it}(s)$ is not defined for observations for which $x_{it} = y_{it}$. Therefore, we define I_{ut} as the set of observations in quarter t for which $m^u_{it}(s)$ exists, i.e., $x_{it} < y_{it}$. We compute mean marginal savings from s to $s+1$ for period t

[29]We restrict the merged sample (Section 2.2) to observations for which we have information necessary to compute utility savings; and to properties with energy efficiency labels F through B. Doing so results in a sample of 4,890,511 transactions.

by taking an average over $i \in I_{st}$ as follows:

$$m_t^u(s) = \frac{1}{|I_{st}|} \sum_{i \in I_{st}} m_{it}^u(s). \tag{8}$$

We then compute the marginal savings $m^u(s)$ from s to $s+1$, $s \in \{1, 2, ..., 99\}$, for the full sample by taking weighted averages of marginal costs $m_t^u(s)$ for periods t = 2012-Q1, 2012-Q2,..., 2020-Q4, as follows:

$$m^u(s) = \frac{1}{\sum_t |I_{st}|} \sum_t \{|I_{st}| \times m_t^u(s)\}. \tag{9}$$

Finally, we compute a single measure for marginal savings for each quarter as:

$$m_t^u = \frac{1}{|\Lambda_t|} \sum_{s \in \Lambda_t} m_t^u(s), \tag{10}$$

and that for the complete duration of the sample as:

$$m^u = \frac{1}{|\Lambda|} \sum_{s \in \Lambda} m^u(s), \tag{11}$$

where Λ_t is a subset of values of $s \in \{1, 2, ..., 99\}$ for which $m_t^u(s)$ exists, and Λ is a subset of values of $s \in \{1, 2, ..., 99\}$ for which $m^u(s)$ exists. We compute the median marginal savings metrics by replacing averages with medians in Equations (8) through (11).

The median marginal energy savings measure m^u computed in Equation (11) is used in Section 3.2.4 to reverse engineer the net discount rate. The period-specific mean and median measures of marginal energy savings m_t^u obtained from Equation (10) are plotted in Figure IA.9 in Section IA.4 of the Internet Appendix, and are used to compute the period-wise net discount rates in Section 4.1.

5.2 Implied Marginal Net Discount Rates

For each energy efficiency score $21 \leq s \leq 91$ (x-axis), Figure 14 plots the marginal annual energy savings $m^u(s)$ from a unit increase in energy efficiency score from s to $s+1$ (primary y-axis), expressed a percentage of property prices (Equation (9)). We observe that marginal savings decrease with subsequent increases in energy scores. This suggests declining returns for investments in energy efficiency. The secondary (right) y-axis marks the corresponding

Figure 14: Implied Marginal Discount Rates

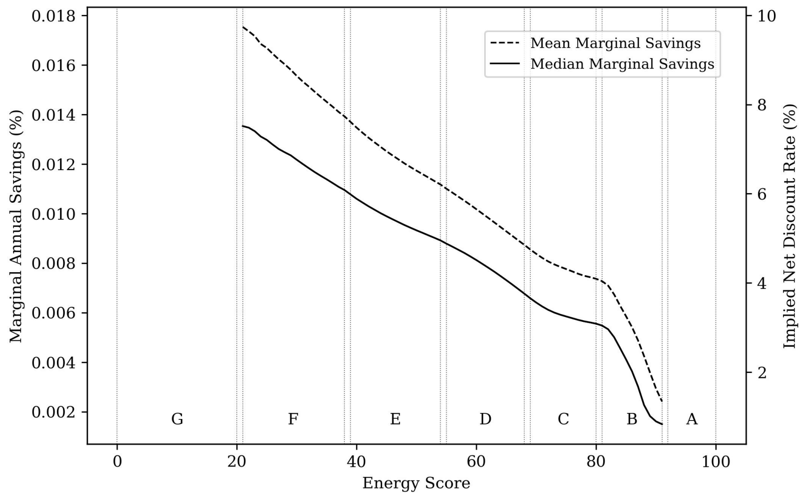

This figure plots marginal energy savings $m^u(s)$ computed from Equation (9) (left y-axis) against energy efficiency scores s (x-axis). Savings are expressed as a proportion of property prices. The secondary y-axis shows the implied net discount rates assuming a constant marginal energy premium of 0.18% estimated in Section 3.

net discount rates implied by a (constant) energy premium of 0.18% estimated in Section 3, computed using the perpetuity formula.

Figure 14 shows that the social discount rate used by homeowners to value reductions in energy expenditures is strongly declining in the initial level of energy efficiency. Therefore, homeowners are willing to accept lower returns – measured in terms of annual energy savings – for greener properties.

One can argue that the decreasing marginal net discount rates are the consequence of using a constant energy premium, and that if we compute and use (presumably declining) marginal energy premia for our analysis instead, we would obtain a straight line in Figure 14. Our piece-wise regression analysis in Section IA.2.2 of the Internet Appendix reveals, however, that a constant energy premium is a sound assumption given that it does not decrease (or increase) systematically with an increase in energy efficiency scores. Furthermore, Section IA.4 of the Internet Appendix shows that estimating non-constant marginal energy premia by including a second-order term for energy efficiency score in Equation (1a) does not alter our conclusions.

Another potential explanation is that more energy efficient properties have a lower time-on-market (Fuerst, Haddad, and Adan, 2020) resulting in an earlier realisation of cash flows. However, time-to-market alone cannot reconcile the difference between the net discount rates. The average net discount rate for marginal energy savings for a property with energy label `E` is 5.4% while that with an energy label `C` is 3.3%. If we *assume* that the "correct" discount rate for properties with an energy label `E` is the same as that for those with an energy label `C`, then for us to empirically observe a discount rate of 5.4% in the data, the cash flows (i.e., energy savings) for properties with an energy label `E` should be realised 9.36 years after those for properties with an energy label `C`, which is implausible.

Therefore, it is reasonable to conclude that homeowners derive non-pecuniary benefits from the energy efficiency of their dwellings; and pay a *green premium.*

6 Energy Upgrades and Policy Impact

Of the 5,769,603 unique properties in our merged sample (Section 2.2), 80.83% of properties were sold once, 16.83% were sold exactly twice, and 2.34% of properties were sold more than two times over the duration of our sample. Table IA.4 and Table IA.5 in Section IA.5 of the Internet Appendix show that properties with multiple transactions are sampled proportionately from the Price Paid Data. Of the 970,760 properties that were sold exactly twice, 226,829 (23.36%) had a new energy certificate issued at the time of their second transaction, of which 156,431 properties had their energy ratings upgraded.[30] In this section, we study how the marginal costs associated with upgrading energy efficiency ratings impact investment decisions of homeowners (Section 6.2), how these decisions evolve in response to regulatory interventions (Section 6.3), and whether properties negatively impacted by the regulations transact at a discount (Section 6.4). Our sample only retains the most recent energy certificate corresponding to each sale recorded in HM Land Registry; hence, subsequent certificate issuances that are not followed by a sale are excluded. However, properties

[30]For 18,258 of the 156,431 properties that had their energy efficiency ratings upgraded, the energy certificate corresponding to the first transaction had expired by the time of the second transaction. However, it is important to note that the requirement for a new energy certificate *does not* carry an obligation to carry out energy upgrades.

that do not transact over the duration of our sample may have their energy ratings upgraded due to environmental concerns, in response to regulatory interventions, or for a variety of other reasons. Therefore, we further replicate our analysis on 2,871,736 properties that had exactly two certificates issued in the Energy Performance of Buildings data, which constitute 17.54% of the total number of properties in the dataset over the duration of our sample (see Table IA.6 in Section IA.5 of the Internet Appendix). We find that of these 2,871,736 properties, 1,613,983 had their energy ratings upgraded.

6.1 Overview of Policies

The Minimum Energy Efficiency Standard (MEES) was announced by the UK Government on 26 March 2015 (UK Government, 2017b), and stated that starting from 01 April 2018, landlords must not grant tenancy to new tenants, or revise tenancy for existing tenants, if their property has an energy efficiency label lower than `E` (score ≤ 38). The regulation further mandated that from 01 April 2020, landlords must not continue letting a property if it has an EPC rating below `E`. Unlike MEES, the Clean Growth Strategy (GGS), published by the UK Government on 12 October 2017 (UK Government, 2017a), was not a regulatory announcement, but outlined the government's agenda to upgrade as many properties as possible to an energy efficiency label of `C` or above (score ≥ 69) by 2035, where practical, cost-effective, and affordable.[31]

Because of the restrictions imposed on private-rental properties with energy efficiency labels `F` and `G`, we expect to observe a greater number of upgrades for these properties post-policy, relative to private-rental properties with greener labels, as well as to the owner-occupied market segment. Furthermore, we expect properties with energy efficiency scores between 1 and 38 to sell at an additional discount, relative to properties with scores greater than or equal to 39, after MEES was announced. Similarly, because properties with energy efficiency scores greater than 69 are not expected to make investments in having their energy profiles upgraded over the next decade, we expect properties with scores between 69 and 100

[31] It is not clear from the policy document what quantitative metrics should be assigned to *practical*, *cost-effective* and *affordable*; see the following *parliamentary report* for a critique.

to command an additional premium, relative to those with scores less than or equal to 68, after CGS was published. We expect the relative impact of CGS to be modest compared to MEES given that CGS is non-binding.

6.2 Marginal Upgrade Costs and Implications on Energy Upgrades

For each energy efficiency score $21 \leq s \leq 91$ (x-axis), Figure 15 plots the median marginal cost $m^c(s)$ for upgrading the energy rating of a property by one unit from s to $s+1$ (y-axis), expressed as a percentage of property prices. We compute marginal costs by following the same methodology used to compute marginal savings in Section 5.1.[32] We observe that upgrading a property becomes increasingly expensive and that marginal costs are substantially higher than the energy premium (0.18%). An increasing marginal cost (together with decreasing marginal savings) suggests that the proportion of properties that have their energy ratings upgraded should decrease with an increase in the initial level of energy efficiency.

Figure 15: Marginal Upgrade Costs

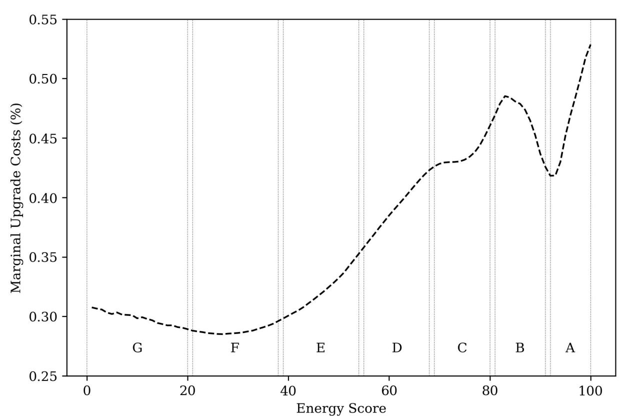

This figure plots marginal upgrade costs savings $m^c(s)$ computed from Equation (9) (y-axis) against energy efficiency scores s (x-axis). Costs are expressed as a proportion of property prices. We restrict the merged sample (Section 2.2) to observations for which we have information necessary to compute upgrade costs. Doing so results in a sample of 4,938,324 transactions.

[32] Similar to the notation used for marginal savings in Section 5.1, we represent marginal cost of upgrade for improving the the energy score of a property by one unit from $s \in \{1, 2, ..., 99\}$ to $s+1$ as $m^u(s)$, and the marginal upgrade cost specific to transaction i at time t as $m^c_{it}(s)$. We *assume* that $m^c_{it}(s)$ is uniform between x_{it} and y_{it}. Finally, we obtain period-wise (m^c_t) and overall (m^c) marginal upgrade cost measures by substituting Δu_{it} with c_{it} (the cost of upgrading the property from x_{it} to y_{it}), and m^u with m^c in Equations (7), (8), (9), (10), and (11).

Figure 16 corroborates our reasoning. For owner-occupied properties that were transacted exactly twice in our merged sample and had a new energy efficiency certificate lodged between the first and the second transaction, the solid black line in Panel (a) plots the proportion of properties that had their energy ratings upgraded (y-axis) against their initial energy efficiency score (x-axis). The dashed (dotted) line plots the proportion of upgrades for properties which had their second energy certificate lodged before (after) the announcement of the Minimum Energy Efficiency Standards (MEES) on 26 March 2015. Panel (b) duplicates the analysis for private-rental properties. Since we retain only the most recent certificates for properties that correspond to a transaction in the HM Land Registry over the duration of our sample, energy upgrades that were not followed by a sale are excluded from our analyses. Therefore, we replicate the analyses in Panels (a) and (b) on the Energy Performance of Buildings Data for properties that had exactly two energy performance certificates issued, and report the results in Panels (c) and (d) respectively.

We observe, across all panels, that properties with higher initial levels of energy efficiency are less likely to have their energy profiles upgraded. Panel (d) reveals that from the pre-policy (dashed line) to post-policy period (dotted line), the increase in proportion of upgrades for private-rental properties with lower energy ratings is substantially greater than those with higher ratings. This is expected because MEES impacts properties with labels `F` and `G`. Surprisingly, although MEES does not impact the owner-occupied market segment, we observe a marked increase in proportion of upgrades post-policy for lower-rated properties in Panel (c). The increase in the proportion of upgrades for properties with labels `F` and `G` post-policy, however, is smaller in the owner-occupied segment (10%) than the private-rental segment (20%). This is consistent with Clara, Cocco, Naaraayanan, and Sharma (2022) who note a very significant increase in the issuance of certificates for lower-rated properties in the private-rental sector (compared to owner-occupied) after MEES approval. We also observe that properties with labels `F` and `G` in owner-occupied market segment were more likely to invest in energy upgrades than private-rental pre-policy, suggesting underinvestment for energy efficiency in the rental markets (see Iwata and Yamaga, 2008; Rehdanz, 2007) in the absence of regulatory intervention. We now focus on properties in our sample which were transacted exactly twice and had a new certificate issued at the time of the second transaction. Panels (a) and (b) show that pre-policy proportion of upgrades for private-rental properties with energy labels `F` and `G` were close to 100%, and over 95% for owner-occupied. Surprisingly, from the pre-policy to post-policy period, the increase in the proportion of upgrades for properties with labels `F` and `G` is lower than those with labels `E`, `D` and `C`.

Figure 16: Proportion of Upgrades by Initial Level of Energy Efficiency

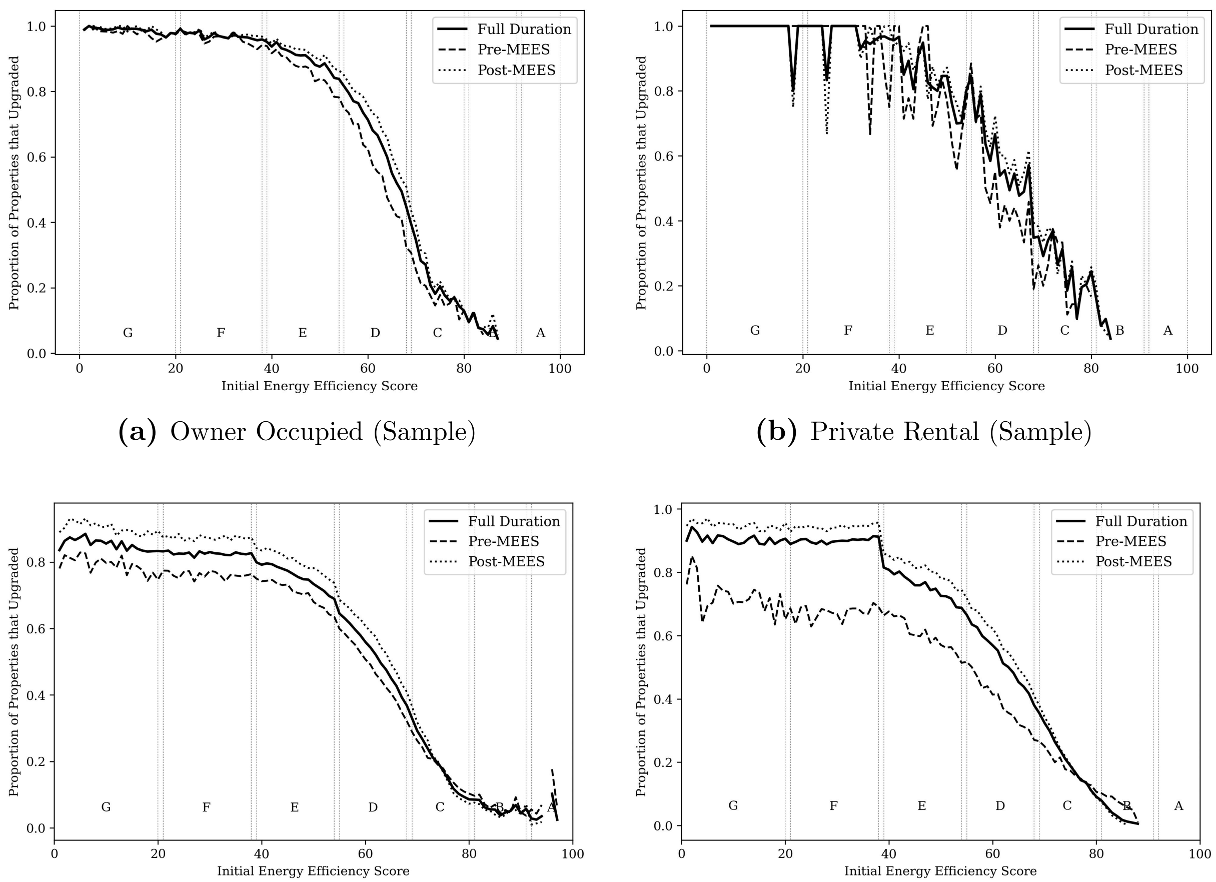

(a) Owner Occupied (Sample)

(b) Private Rental (Sample)

(c) Owner Occupied (EPC Data)

(d) Private Rental (EPC Data)

The solid black line in plots the proportion of properties that upgraded (y-axis) against their initial energy efficiency score (x-axis). The dashed (dotted) line plots the proportion of upgrades for properties which had their second energy certificate lodged before (after) the announcement of the Minimum Energy Efficiency Standards (MEES) on 26 March 2015. For properties that were transacted exactly twice in our merged sample, Panel (a) and Panel (b) plot results for the owner-occupied and private-rental market segments respectively. Panels (c) and (d) replicate the analyses in Panels (a) and (b) respectively on the Energy Performance of Buildings dataset (see Section 2.1.2) for properties that had exactly two energy performance certificates issued.

Therefore, we find three counterintuitive patterns: (i) the proportion of upgrades in the owner-occupied segment increases post-policy even though only the rental segment is subject to regulatory intervention; and for properties that had a new energy certificate issued at the time of their second transaction, the proportion of upgrades (ii) is not only greater for lower-rated properties across both market segments (compared to that observed in the energy certificates data), (iii) but also increases more for higher-rated properties post-policy. These findings suggest that government policies may have had an *indirect* effect in helping improve energy profile of the housing stock in UK by increasing an *already existing* market sentiment towards sustainability. Indeed, Table 8 shows that when energy upgrades are followed by a transaction, properties with labels `F` and `G` are more likely to upgrade to labels `D` and `C`. In contrast, when we look at the population data, there is a higher tendency to upgrade to `E` (just enough to overcome the policy threshold).

Table 8: Upgrade Matrix

	G	F	E	D	C	B	A
A							100 (100)
B						93.72 (96.38)	6.28 (3.62)
C					87.34 (84.32)	12.10 (15.23)	0.56 (0.45)
D				50.05 (51.93)	48.29 (46.31)	1.57 (1.66)	0.10 (0.10)
E			12.37 (19.66)	62.46 (62.94)	24.11 (16.66)	1.01 (0.69)	0.04 (0.05)
F		2.64 (6.28)	22.11 (40.42)	53.69 (40.62)	20.43 (12.29)	1.08 (0.37)	0.05 (0.03)
G	1.34 (5.77)	4.99 (12.68)	16.00 (28.34)	54.53 (38.97)	21.87 (13.85)	1.21 (0.37)	0.07 (0.04)

For each energy label corresponding to a row, this table displays the proportion of properties that upgraded to the energy label corresponding to each column. Each row adds up to 100%. The values outside parentheses report values for properties with upgrades followed by a transaction. The values inside the parentheses provide corresponding values for the energy certificates data.

For each energy label corresponding to a row, Table 8 displays the proportion of upgrades to energy labels corresponding to each column. Therefore, each row adds up to 100%. The values outside parentheses report proportions for properties with upgrades followed by a transaction. The values inside parentheses provide corresponding proportions for the energy certificates data. The magnitude of upgrades for properties with brown labels is significantly greater compared to those with green labels. For instance, while 87.34% of properties with an initial energy label `C` retain their label post-upgrade, 87.63% of properties with an initial label `E` upgrade to a higher label. This is consistent with the fact that marginal upgrades not only become progressively expensive, but also yield lower reductions in energy expenditures.

6.3 Policy Impact on Energy Upgrades over Time

We build upon our discussion in the previous section by tracking energy upgrades over time. For owner-occupied properties that were transacted exactly twice in our merged sample and had a new energy certificate issued, Panel (a) in Figure 17 shows the proportion of properties that had their energy ratings upgraded in each year for each (initial) energy efficiency label. Panel (b) duplicates the analysis for private-rental properties. Panels (c) and (d) visualise results corresponding to the Energy Performance of Buildings data (see Section 2.1.2) for properties that had exactly two energy performance certificates issued. From left to right, the grey vertical lines mark MEES approval on 26 March 2015, CGS publication on 17 October 2017 and MEES implementation on 01 April 2018. In each panel, the average level of trend lines progressively becomes lower for greener energy labels, which is consistent with our observation that properties with higher initial level of energy efficiency are less likely to have their energy ratings upgraded due to increasing marginal costs and decreasing marginal savings.

Panel (d) shows that for private-rental properties with labels `F` and `G` (which are directly impacted by the MEES regulation), the proportion of upgrades increase substantially after MEES is introduced in 2015, and peak in 2018, when the policy came into effect. However, we observe that properties with labels `E` and `D` also feature substantial spikes in proportion of upgrades, together with a modest increase in that for properties with label `C`. Furthermore, Panel (c) shows that proportion of upgrades for owner-occupied properties mirrors the trends observed in Panel (d). We could perhaps attribute the increase in energy upgrades owner-occupied properties with labels `F` and `G` and private-rental properties with labels `E` and `D` to precautionary incentives. Buy-to-live homeowners with properties labelled `F` and `G` may upgrade to retain the option to lease out the property in the future, while buy-to-let landlords with properties labelled `E` and `D` may upgrade in anticipation of further tightening of regulations. However, such incentives would not explain why owner-occupied properties with labels `E` and `D` see substantial upgrades (together with modest upgrades for properties with labels `C`), or why the increase in proportion of upgrades for greener labels is as steep as those for labels `F` and `G`.

Figure 17: Proportion of Upgrades over Time

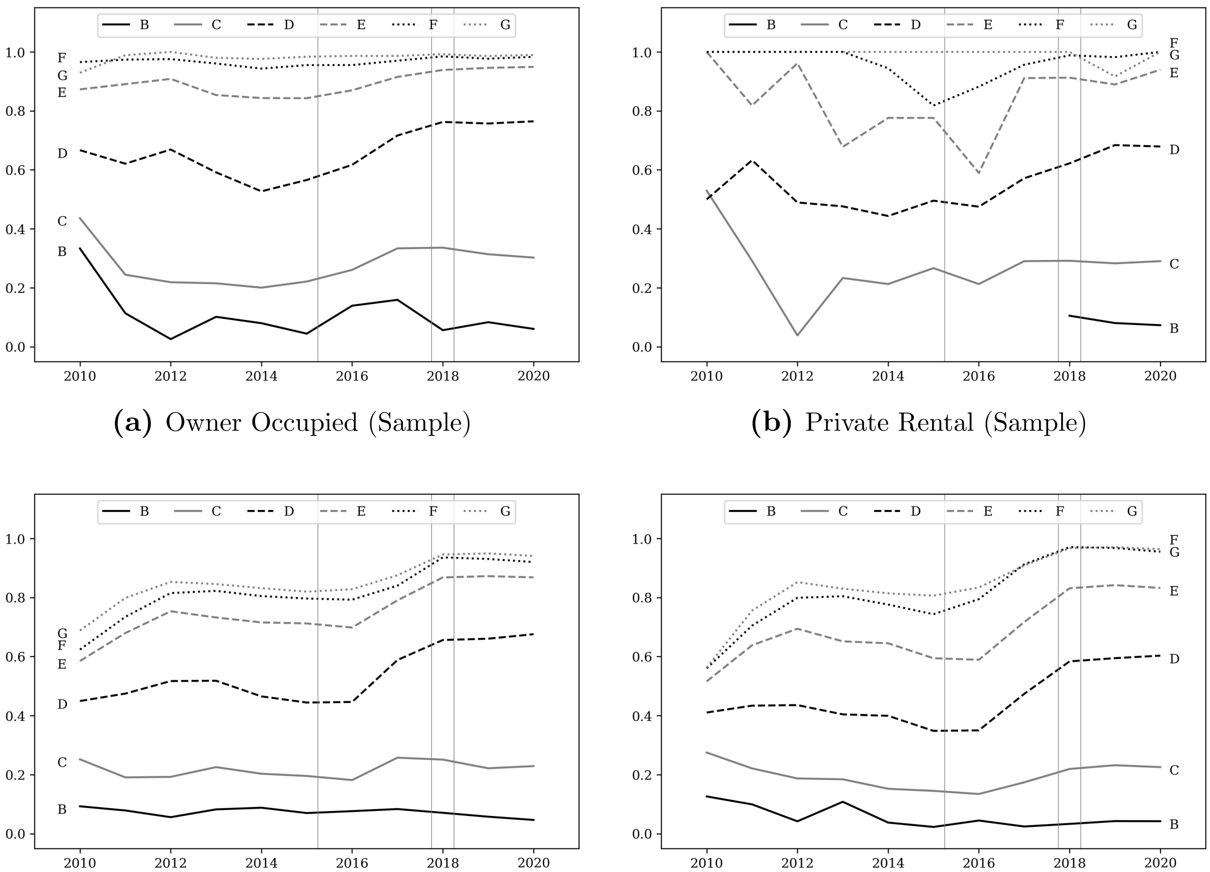

(a) Owner Occupied (Sample)

(b) Private Rental (Sample)

(c) Owner Occupied (EPC Data)

(d) Private Rental (EPC Data)

This figure shows the proportion of properties that had their energy ratings upgraded in each year for each initial energy efficiency label. Panel (a) and (b) show results for owner-occupied and private-rental properties that were transacted exactly twice in our merged sample and had a new energy certificate issued. Panels (c) and (d) visualise results corresponding to the Energy Performance of Buildings dataset (see Section 2.1.2) for properties that had exactly two energy performance certificates issued. From left to right, the grey vertical lines mark MEES approval, CGS publication and MEES implementation.

This reinforces our conclusion in Section 6.2 that government regulations had an indirect effect in helping improve energy profile of the housing stock by increasing an already existing market sentiment towards sustainability. Furthermore, Figure 17 reveals that these indirect effects are primarily attributable to the announcement of MEES in 2015, and not the publication of CGS in 2017, although it is the latter that classifies properties with label `C` or above as sustainable, and is targeted towards the entire housing market. One could argue that binding regulations are perhaps more successful in raising the overall market sentiment towards sustainability than non-binding manifestos, even when the proportion of dwellings directly impacted by the binding regulation is very small. Finally, Panels (a) and (b) show that similar results hold for properties that were transacted exactly twice and had a new energy certificate issued, with the exception that proportion of upgrades for labels before MEES was introduced was much higher across labels `G` through `B`. Therefore, homeowners who intend to sell their properties are more likely to invest in energy upgrades, presumably because buyers pay a green premium for more energy efficient properties (see Section 5).

6.4 Policy Impact on Property Prices

In this section, we investigate whether properties with energy efficiency scores less than or equal to 38 sold at a discount after 26 March 2015, the date at which the Minimum Energy Efficiency Standards (MEES) was announced, using a standard Difference in Difference (DD) approach based on Angrist and Pischke (2008, Chapter 5.2) in Section 6.4.1, and a Sharp Regression Discontinuity (Sharp RD) design based on Angrist and Pischke (2008, Chapter 6) and Lee and Lemieux (2010) in Section 6.4.2. Section 6.4.3 outlines the results.

6.4.1 Methodology 1: Difference in Differences

Difference in Difference (DD) is a quasi-experimental technique that mimics an experimental research design using observational study data. It *assumes* that in the absence of treatment, the differences in potential outcomes between the treatment and the control groups are the same before and after the implementation of the policy.[33] Therefore, DD is applicable in our context if we assume that, in the absence of MEES, the evolution of expected price of a property with an energy efficiency score greater than 38 (the control group) would be *parallel* to that of a property with score less than or equal to 38 (the treatment group), holding all else equal. If we further *assume* that the policy and treatment effects are linear and additive, we can extend the hedonic regression models in Section 3.1 into a Difference in Difference (DD) setup. In particular, consider the following conditional expectation function

[33]This is known as the "counterfactual trends" or "parallel trends" assumption.

for transactions that occurred before MEES was implemented:

$$\begin{aligned}\mathbb{E}\left[\log{(P/A)}_{ihrt}^{\text{untreated, pre-MEES}}\right] &= \alpha_r + \delta_t + \mathrm{S}_{ih}\beta_{\texttt{Score}} + (\mathbb{1}_t.\mathrm{S}_{ih})\beta_{\texttt{Score},t} \\ &\quad + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega,\end{aligned} \tag{12}$$

where the unit of observation in our model is the same as that in Equation (4) in Section 3.1 with $\log{(P/A)}_{ihrt}$ as the target variable; and each symbol has the same meaning as that in Equation (4). Let $\lambda_{\texttt{MEES}} \in \mathbb{R}$ denote the fixed effect of MEES on the target variable (for both treatment and control groups) such that:

$$\begin{aligned}\mathbb{E}\left[\log{(P/A)}_{ihrt}^{\text{untreated, post-MEES}}\right] &= \mathbb{E}\left[\log{(P/A)}_{ihrt}^{\text{untreated, pre-MEES}}\right] + \lambda_{\texttt{MEES}} \\ &= \alpha_r + \delta_t + \mathrm{S}_{ih}\beta_{\texttt{Score}} + (\mathbb{1}_t.\mathrm{S}_{ih})\beta_{\texttt{Score},t} + \lambda_{\texttt{MEES}} \\ &\quad + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega.\end{aligned} \tag{13}$$

Finally, let $\rho_{\texttt{MEES}} \in \mathbb{R}$ denote the casual effect of treatment such that:

$$\begin{aligned}\mathbb{E}\left[\log{(P/A)}_{ihrt}^{\text{treated, post-MEES}}\right] &= \mathbb{E}\left[\log{(P/A)}_{ihrt}^{\text{untreated, post-MEES}}\right] + \rho_{\texttt{MEES}} \\ &= \alpha_r + \delta_t + \mathrm{S}_{ih}\beta_{\texttt{Score}} + (\mathbb{1}_t.\mathrm{S}_{ih})\beta_{\texttt{Score},t} + \lambda_{\texttt{MEES}} + \rho_{\texttt{MEES}} \\ &\quad + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega.\end{aligned} \tag{14}$$

This results in the following regression model:

$$\begin{aligned}\log{(P/A)}_{ihrt} &= \alpha_r + \delta_t + \mathrm{S}_{ih}\beta_{\texttt{Score}} + (\mathbb{1}_t.\mathrm{S}_{ih})\beta_{\texttt{Score},t} + \lambda_{\texttt{MEES}} + \rho_{\texttt{MEES}}X_i \\ &\quad + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt},\end{aligned} \tag{15}$$

where X_i is an indicator variable for treatment; $X_i = 1$ when transaction i takes place on or after 26 March 2015 *and* the energy efficiency score of the underlying property is less than or equal to 38; $X_i = 0$ otherwise.

Because energy premium is heterogenous across tenure (see Section 4.3) and MEES only affects rental properties, we restrict our sample to buy-to-let transactions (i.e., properties purchased with the intention of renting out). In addition, to avoid potential confounding impacts of the 2008 Global Financial Crisis and the 2020 Covid-19 Pandemic on our estimates, we further restrict our sample to transactions that took place between 01 January 2011 and 31 December 2019. Doing so also provides us with a reasonably balanced sample before and after the policy cutoff (i.e., 26 March 2015). However, our reduced sample may still be subject to the confounding impact of CGS, which was published on 12 October 2017 and classified properties with energy efficiency labels C or above (score ≥ 69) as "safe" or "green". In particular, the announcement of CGS in October 2017 would lead to a violation of the parallel trends assumption in Equation (15), as properties with energy efficiency labels C or above can be expected to sell at a premium relative to those with labels D or below, post-CGS. Therefore, price evolution of properties with labels below E (score ≤ 38) and greater to or equal to E (score > 38) can no longer be assumed to be counterfactually parallel in the absence of treatment. To address these problems, we (i) restrict our sample to properties with labels D or below (score < 69), *and* (ii) introduce a fixed effect for CGS in Equation (15) as follows:

$$\begin{aligned}\log{(P/A)_{ihrt}} = \alpha_r + \delta_t + \mathrm{S}_{ih}\beta_{\mathtt{Score}} + (\mathbb{1}_t.\mathrm{S}_{ih})\beta_{\mathtt{Score},t} \\ + \lambda_{\mathtt{MEES}} + \rho_{\mathtt{MEES}} X_i + \lambda_{\mathtt{CGS}} \\ + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt},\end{aligned} \tag{16}$$

where $\lambda_{\mathtt{CGS}} \in \mathbb{R}$ controls for the average change in $\log{(P/A)_{ihrt}}$ for properties with energy efficiency labels D or below (score < 69) that were sold after publication of CGS. Note that in Equation (16), we cannot introduce $\lambda_{\mathtt{CGS}}$ without first restricting our sample to properties with labels D or below. This is because the impact of CGS on properties with labels D or below (score < 69) will be opposite to that on those with labels C or above (score ≥ 69), whereas the fixed effect ($\lambda_{\mathtt{CGS}}$) in Equation (16) is homogenous for all properties post-CGS.[34] Restricting our sample in this manner also helps in alleviating concerns about potential violations of the parallel trends assumption due to variations in market sentiment towards green properties.

In order to conclude that MEES was effective, we require $\rho_{\mathtt{MEES}}$ (the causal effect of interest) to be negative and statistically significant, because properties with energy scores between 1 and 38 are subject to additional legal restrictions post-policy. Additionally, we expect $\lambda_{\mathtt{CGS}}$ (the fixed effect of CGS) to be negative as the properties in our reduced sample would be required to make investments in having their energy profile upgraded over the next decade, resulting in a decrease in transaction values.

[34]Inclusion of properties with labels C or above (score ≥ 69) will require a different model specification; for instance, a triple differences (DDD) approach, which "disentangles" differences in outcomes for properties with labels ABC and FG from those with labels DE and FG.

Finally, Section IA.5 of the Internet Appendix outlines additional DD specifications that (i) eliminate confounding impact of CGS by reducing the duration of the sample around the MEES announcement (i.e., restricting the sample from 01 January 2010 to 11 October 2017 with 26 March 2015 as policy-cutoff), (ii) consider the MEES implementation date instead of the MEES announcement date as the policy-cutoff (i.e., restricting the sample from 12 October 2017 to 31 December 2019 with 01 April 2018 as policy-cutoff), and (iii) study the impact of CGS on "green" properties (i.e., restricting the sample from 26 March 2015 to 31 December 2019 with 12 October 2017 as policy-cutoff). We find that the results obtained for Equation (16), as discussed in Section 6.4.3, remain robust to the additional specifications in Section IA.5 of the Internet Appendix.

6.4.2 Methodology 2: Regression Discontinuity

Regression Discontinuity (RD) can be used to establish causal effects in settings where treatment is a deterministic and discontinuous function of a covariate and the agents have an imprecise control over which side of the treatment cutoff they will land on; i.e., agents in the neighbourhood of the cutoff have approximately the same probability of being just above (receiving the treatment) or just below (being denied the treatment). In such a situation, we can think of the assignment as a randomised experiment and draw causal inferences on the treatment effect. Lee and Lemieux (2010) remark that RD designs require milder assumptions compared to those needed for other non-experimental approaches, and that causal inferences from RD designs are potentially more credible than those from typical *natural experiment* strategies (e.g., difference-in-differences or instrumental variables). For example, the instrumental variable approach *assumes* that instrument is extraneously generated; an assumption which is often hard to justify. Furthermore, comparisons of average outcomes in a small enough neighbourhood to the left and right of the cutoff should provide an estimate of the treatment effect that does not depend on the correct specification of the model (Angrist and Pischke, 2008). This is important, as it sets RD apart from other potential policy impact estimators such as DD, which requires that all trends and interactions are properly included.

We now explain how RD is applicable in our context. Once the property owner lodges a request for an energy efficiency inspection, a surveyor examines the property physically, takes pictures of the items relevant to energy certification, and inputs his observations into a digital tablet. A back-end software then automatically calculates the numerical energy efficiency rating between 1 and 100. The numerical score is then converted into an alphabetical label between G and A; see Table 4. To use RD to analyse the impact of MEES, we are most interested in the numerical cutoff (i.e., 39) at which a property is labelled E if the numerical score is 39 or more, and F, if it is lower. Given the human errors in examining the property and feeding information into the software, and the inability to accurately observe information such as the age of the dwelling, a property close to the numeric threshold (i.e., 39) could have easily landed on either side. If we only consider properties with energy efficiency scores in the proximity of the threshold, then there is no reason to suspect that homeowners whose properties have an energy label E are more concerned about energy efficiency than those with properties with label F. Further, if we restrict our sample to dwellings for which energy efficiency certificates were issued *before* the MEES regulation was announced (that is, before the landlords had an incentive to distinguish between ratings E and F) but sold *after* MEES (so that the rating E around the threshold acts as treatment of being marked safe), we can assume random assignment and draw a causal inference on whether the policy led to a discount for dwellings affected by the policy.[35]

This motivates the use of RD, which comes in two flavours - *sharp* and *fuzzy*. Sharp RD is relevant to settings where assignment of treatment is perfectly known, whereas Fuzzy RD is a two-step IV-like approach used in settings where assignment of treatment around the cut-off is not perfectly known (e.g., when we are trying to predict assignment instead of knowing it). Since we have perfect information about energy ratings, the corresponding labels, and their treatment, we deploy a Sharp RD model. We start by selecting private rental properties that were transacted after 01 April 2018 (implementation of MEES) but were issued an energy certificate before 26 March 2015 (announcement of MEES). We further restrict our sample to properties which have a numerical EPC between 33 and 44. Lastly, we remove transactions

[35]The random assignment assumption will not hold for green labels (i.e., A, B and C). This is because homeowners who own properties with higher energy efficiency ratings may care about the energy label of their property, and may therefore opt for an energy rating in a non-random way.

that took place before 2011 or after 2019 to avoid the confounding impact of the 2008 Global Financial Crisis and the 2020 Covid-19 Pandemic on our estimates.

We denote the numerical cutoff at which a property is labelled E as $c = 39$. The unit of observation in our model is the same as that in Equation (4) in Section 3.1 with $\log{(P/A)_{ihrt}}$ as the target variable. We distinguish between properties that receive the treatment (properties that were impacted by the policy) as $\log{(P/A)^1_{ihrt}}$ and those that do not as $\log{(P/A)^0_{ihrt}}$. Consider the following conditional expectation formulation:

$$\begin{aligned}\mathbb{E}\left[\log{(P/A)^0_{ihrt}}\right] = \alpha_r + \delta_t + (\mathrm{S}_{ih} - c)\beta_{\texttt{Score}} + \mathbb{1}_t.(\mathrm{S}_{ih} - c)\beta_{\texttt{Score},t} \\ + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega,\end{aligned} \tag{17}$$

where $\mathrm{S}_{ih} - c$ is the centered numerical energy efficiency score, and all other symbols have exactly the same meaning as that in Equation (4). Let $\rho \in \mathbb{R}$ denote the casual effect of treatment such that:

$$\log{(P/A)^1_{ihrt}} = \log{(P/A)^0_{ihrt}} + \rho. \tag{18}$$

This yields the following regression model:

$$\begin{aligned}\log{(P/A)_{ihrt}} = \alpha_r + \delta_t + (\mathrm{S}_{ih} - c)\beta_{\texttt{Score}} + \mathbb{1}_t.(\mathrm{S}_{ih} - c)\beta_{\texttt{Score},t} + \rho X_i \\ + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt},\end{aligned} \tag{19}$$

where X_i is an indicator variable defined as:

$$\mathrm{X}_i = \begin{cases} 1 & \text{if } S_{ih} \geq c \\ 0 & \text{if } S_{ih} < c \end{cases} \tag{20}$$

and ρ is the causal effect of interest. Note that because Equation (19) models differences in outcomes between treatment and control groups during the same time period, it does not require the parallel trends assumption like Equation (16) in Section 6.4.1. To conclude that MEES was effective, we require ρ (the causal effect of interest) to be negative and statistically significant, as properties with energy scores below the cutoff are subject to additional restrictions.

6.4.3 Results

We find that the $\rho_{\texttt{MEES}}$ (the treatment effect of MEES on properties with energy efficiency score less than 39) and $\lambda_{\texttt{CGS}}$ (the fixed effect of CGS on properties with energy efficiency score less than 69) in Equation (16) are positive and statistically insignificant. This is *opposite* to what we expected at the outset. These results remain robust against the alternative DD specifications discussed in Section 6.4.1 and elaborated in Appendix IA.5. We may attribute these surprising results to model misspecification and violations of the parallel trends assumption. In Section 6.4.2, we discuss how an RD specification overcomes these potential challenges. Indeed, estimating Equation (19) yields a negative causal effect of interest; i.e., $\rho = -0.0035$ which corresponds to a 0.35% decline in property prices per unit area for rental properties subject to leasing restrictions post-MEES. However, this estimate is statistically insignificant, with a p-value of 0.828. We may suspect that expanding or contracting the neighbourhood (i.e., 33 to 44) around the cutoff (i.e., 39) in the RD design may provide more precise estimates. We find the magnitude of ρ changes based on how wide the neighbourhood is, but the estimates consistently remain close to zero, and insignificant with p-values greater than 0.5. Therefore, we conclude that MEES did not lead to an additional discount for rental properties with energy labels below E, and that CGS did not lead to an additional premium for rental properties with energy labels above C (Section IA.5 of the Internet Appendix).

One can argue that because CGS is not enforceable, the market simply did not react to its publication. This reasoning would be consistent with our observation in Section 6.3 that the increase in proportion of energy efficiency upgrades are primary attributable to MEES, which is enforceable. Notwithstanding, because MEES led to substantial increase in the proportion of energy upgrades across the entire housing sector, the absence of treatment effect on properties with labels `F` and `G` that did not upgrade post-MEES is confounding.

Figure 17 shows that the proportion of properties with labels `F` and `G` that did not upgrade post-MEES is less than 5%; thus, one potential explanation for the absence of a treatment effect is that most of these properties are eligible for one of the MEES-exemptions enumerated in Table 9. For example, in our sample, an average household with energy label

Table 9: Exemptions to Minimum Energy Efficiency Standards (MEES)

Exemption	Description
High Cost	The prohibition on letting property below an EPC rating of E does not apply if the cost of making even the cheapest recommended improvement would exceed £3,500.
All Improvements Made	Where all the relevant energy efficiency improvements for the property have been made (or there are none that can be made) and the property remains sub-standard.
Wall Insulation	The landlord has obtained written expert advice indicating that the measure is not appropriate for the property due to its potential negative impact on the fabric or structure of the property.
Consent	Certain energy efficiency improvements may legally require third party consent (e.g., local authority planning consent, consent from mortgage lenders, etc.) before they can be installed.
Devaluation	An exemption from meeting the minimum standard will apply where the landlord has obtained a report from an independent surveyor who is on the Royal Institution of Chartered Surveyors (RICS) register of valuers advising that the installation of specific energy efficiency measures would reduce the market value of the property, or the building it forms part of, by more than 5%.
New Landlord	If a person becomes a landlord in circumstances where it would be unreasonable for them to be required to comply with the regulations immediately, a temporary 6 month exemption will apply from the date they become the landlord.

This table provides a partial list and description of exceptions from the MEES policy that landlords can file for. The list of exemptions has been adapted from `Guidance on PRS Exemptions` published by UK Government.

`F` is priced at £255,526 and requires its energy efficiency score to be upgraded by at least 7 units to meet the regulation. This translates to an expected upgrade cost of £8,764 that is significantly higher than the £3,500 threshold over which landlords can claim a "High Cost" exemption from MEES. Similarly, properties that cannot be upgraded to an energy label greater than or equal to `E`, and those for which homeowners can argue that energy efficiency improvements would be detrimental to the dwelling's structural integrity, are also exempt from MEES.

It is therefore reasonable to suspect that MEES has a negative and significant impact on prices of only those properties that do not qualify for these exemptions. Based on our dataset, it is not possible to identify the subset of properties that do not qualify for all potential MEES exemptions that homeowners can file for. Nonetheless, in Appendix IA.5, we perform robustness checks by running each DD model on a sample of properties with energy label `F`, and assigning treatment to properties that (i) can be upgraded to an energy efficiency score greater than or equal to E (score $\geq$ 39), and for which (ii) the expected cost of upgrade to a score of 39 is less than £3,500. Therefore, in the absence of policy exemptions, price evolution of the treatment and control groups can be assumed to be identical. Furthermore, we use a similar method to test an alternative RD specification. We find that our conclusions remain robust to these additional specifications.

7 Conclusion

In this paper, we take an important step towards understanding to what extent homeowners care about the environment. We show that homeowners price energy efficiency in a rational manner. In colder climates, where each unit improvement in energy efficiency is expected to result in larger utility savings, the energy premium is higher. In contrast, premia is lower in private rental markets, attributable to market imperfections that make it harder to recoup investments in energy efficiency. We also contribute to the ongoing discourse in the economics of climate finance on the appropriate social discount rate to be used for climate abatement investments, by empirically computing the rate at which homeowners discount future energy savings. More importantly, we show that this social discount rate is declining in the initial level of energy efficiency of a dwelling. Therefore, homeowners derive non-pecuniary benefits from the energy efficiency of their dwellings. Surprisingly, we find that government regulation did not lead to a price impact and was followed by a commensurate increase in the proportion of energy upgrades across market segments targeted and not targeted by the policies.

Our findings could be of interest to homeowners, developers and real estate private equity firms, who wish to understand to what extent investments in energy efficiency will be priced by the market. For example, homeowners may choose to invest in energy efficiency upgrades in regions where marginal upgrade costs are less than the (conditional) premia. Surprisingly, we do not find evidence of such strategic behaviour in our sample. Because we do not observe ownership of properties, it is possible that such behaviour is only displayed by institutional investors and not retail homeowners. This is a potential future direction of research.

Our paper also is also relevant to policymakers who wish to obtain insights into potential barriers and incentives to improving the energy profile of the housing stock. An example of this is the misalignment between the social discount rate used by the regulator and that empirically observed in our sample. Our findings also highlight various channels along which the regulator can offer differential subsidies to incentivise investments in energy upgrades. For example, policymakers can offer higher subsidies in temperate regions where the marginal benefit of investing in energy upgrades is lower.

References

H. Amecke. The impact of energy performance certificates: A survey of german home owners. *Energy Policy*, 46:4–14, 2012.

J. D. Angrist and J.-S. Pischke. *Mostly harmless econometrics*. Princeton university press, 2008.

A. Ayala, I. Galarraga, and J. V. Spadaro. The price of energy efficiency in the spanish housing market. *Energy Policy*, 94:16–24, 2016. ISSN 0301-4215. doi: https://doi.org/10.1016/j.enpol.2016.03.032. URL https://www.sciencedirect.com/science/article/pii/S0301421516301355.

M. Baldauf, L. Garlappi, and C. Yannelis. Does climate change affect real estate prices? Only if you believe in it. *The Review of Financial Studies*, 33(3):1256–1295, 2020.

A. Baranzini, J. Ramirez, C. Schaerer, and P. Thalmann. *Hedonic methods in housing markets: Pricing environmental amenities and segregation*. Springer Science & Business Media, 2008.

A. Bernstein, M. T. Gustafson, and R. Lewis. Disaster on the horizon: The price effect of sea level rise. *Journal of financial economics*, 134(2):253–272, 2019.

D. Brounen and N. Kok. On the economics of energy labels in the housing market. *Journal of environmental economics and management*, 62(2):166–179, 2011.

M. Cajias and D. Piazolo. Green performs better: energy efficiency and financial return on buildings. *Journal of Corporate Real Estate*, 2013.

M. Cajias, F. Fuerst, and S. Bienert. Tearing down the information barrier: the price impacts of energy efficiency ratings for buildings in the german rental market. *Energy Research & Social Science*, 47:177–191, 2019.

P. Cerin, L. G. Hassel, and N. Semenova. Energy performance and housing prices. *Sustainable Development*, 22(6):404–419, 2014.

N. Clara, J. F. Cocco, S. L. Naaraayanan, and V. Sharma. Investments that make our homes greener: The role of regulation. *Available at SSRN 4009054*, 2022.

L. W. Davis. Evaluating the slow adoption of energy efficient investments: Are renters less likely to have energy efficient appliances? Working Paper 16114, National Bureau of Economic Research, June 2010. URL http://www.nber.org/papers/w16114.

P. T. Davis, J. A. McCord, M. McCord, and M. Haran. Modelling the effect of energy performance certificate rating on property value in the belfast housing market. *International Journal of Housing Markets and Analysis*, 2015.

P. Eichholtz, M. Korevaar, T. Lindenthal, and R. Tallec. The total return and risk to residential real estate. *The Review of Financial Studies*, 34(8):3608–3646, 2021.

European Commission. The Energy Performance of Buildings Directive (EPBD). 2002. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32002L0091&qid=1649018545268.

Eurostat. Energy Consumption in Households. 2021. URL https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Energy_consumption_in_households.

F. Fuerst, P. McAllister, A. Nanda, and P. Wyatt. Does energy efficiency matter to home-buyers? an investigation of epc ratings and transaction prices in england. *Energy Economics*, 48:145–156, 2015a.

F. Fuerst, P. McAllister, A. Nanda, and P. Wyatt. Variations in the implicit pricing of energy performance by dwelling type and tenure: A study of wales, 07 2015b.

F. Fuerst, P. McAllister, A. Nanda, and P. Wyatt. Energy performance ratings and house prices in wales: An empirical study. *Energy Policy*, 92:20–33, 2016a.

F. Fuerst, E. Oikarinen, and O. Harjunen. Green signalling effects in the market for energy-efficient residential buildings. *Applied Energy*, 180:560–571, 2016b. ISSN 0306-2619. doi: https://doi.org/10.1016/j.apenergy.2016.07.076. URL https://www.sciencedirect.com/science/article/pii/S0306261916310170.

F. Fuerst, M. F. C. Haddad, and H. Adan. Is there an economic case for energy-efficient dwellings in the uk private rental market? *Journal of Cleaner Production*, 245:118642, 2020.

T. D. Gerarden, R. G. Newell, and R. N. Stavins. Assessing the energy-efficiency gap. *Journal of Economic Literature*, 55(4):1486–1525, December 2017. doi: 10.1257/jel.20161360. URL https://www.aeaweb.org/articles?id=10.1257/jel.20161360.

S. Giglio, M. Maggiori, K. Rao, J. Stroebel, and A. Weber. Climate change and long-run discount rates: Evidence from real estate. *The Review of Financial Studies*, 34(8):3527–3571, 2021.

K. Gillingham, R. G. Newell, and K. Palmer. Energy efficiency economics and policy. *Annu. Rev. Resour. Econ.*, 1(1):597–620, 2009.

B. Groom and D. Maddison. New estimates of the elasticity of marginal utility for the uk. *Environmental and Resource Economics*, 72(4):1155–1182, 2019.

L. Högberg. The impact of energy performance on single-family home selling prices in sweden. *Journal of European Real Estate Research*, 2013.

M. Hyland, R. C. Lyons, and S. Lyons. The value of domestic building energy efficiency evidence from ireland. *Energy Economics*, 40:943–952, 2013. ISSN 0140-9883. doi: https://doi.org/10.1016/j.eneco.2013.07.020. URL https://www.sciencedirect.com/science/article/pii/S0140988313001655.

S. Iwata and H. Yamaga. Rental externality, tenure security, and housing quality. *Journal of Housing Economics*, 17(3):201–211, 2008.

A. Jaffe, R. Stavins, and C. Cleveland. *Economics of Energy Efficiency*, page 7990. Elsevier, 2004.

O. M. Jensen, A. R. Hansen, and J. Kragh. Market response to the public display of energy performance rating at property sales. *Energy Policy*, 93:229–235, 2016.

L. Kaplow, E. Moyer, and D. Weisbach. The Social Evaluation of Intergenerational Policies and Its Application to Integrated Assessment Models of Climate Change. *The B.E. Journal of Economic Analysis & Policy*, 10(2):1–34, November 2010. doi: 10.2202/1935-1682.2519. URL https://ideas.repec.org/a/bpj/bejeap/v10y2010i2n7.html.

D. S. Lee and T. Lemieux. Regression discontinuity designs in economics. *Journal of Economic Literature*, 48(2):281–355, June 2010. doi: 10.1257/jel.48.2.281. URL https://www.aeaweb.org/articles?id=10.1257/jel.48.2.281.

J. Murfin and M. Spiegel. Is the risk of sea level rise capitalized in residential real estate? *The Review of Financial Studies*, 33(3):1217–1255, 2020.

W. D. Nordhaus. A review of the stern review on the economics of climate change. *Journal of economic literature*, 45(3):686–702, 2007.

A. Ramos, A. Pérez-Alonso, and S. Silva. Valuing energy performance certificates in the portuguese residential sector. *Economics for Energy. WP*, pages 02–2015, 2015.

F. P. Ramsey. A mathematical theory of saving. *The economic journal*, 38(152):543–559, 1928.

K. Rehdanz. Determinants of residential space heating expenditures in germany. *Energy Economics*, 29(2):167–182, 2007.

M. Schneider, C. Traeger, and R. Winkler. Trading off generations: Equity, discounting, and climate change. *European Economic Review*, 56(8):1621–1644, 2012. URL https://EconPapers.repec.org/RePEc:eee:eecrev:v:56:y:2012:i:8:p:1621-1644.

N. H. Stern. *The economics of climate change: the Stern review*. cambridge University press, 2007.

UK Government. Clean growth strategy. 2017a. URL https://www.gov.uk/government/publications/clean-growth-strategy.

UK Government. Domestic private rented property: minimum energy efficiency standard - landlord guidance. 2017b. URL https://www.gov.uk/guidance/domestic-private-rented-property-minimum-energy-efficiency-standard-landlord-guidance.

UK Government. The Green Book. 2022a. URL https://www.gov.uk/government/publications/the-green-book-appraisal-and-evaluation-in-central-governent/the-green-book-2020#valuation-of-costs-and-benefits.

UK Government. English Housing Survey. 2022b. URL https://www.gov.uk/government/statistics/english-housing-survey-2020-to-2021-headline-report.

J. Zhuang, Z. Liang, T. Lin, and F. De Guzman. Theory and practice in the choice of social discount rate for cost-benefit analysis: a survey. Technical report, ERD working paper series, 2007.

Internet Appendix to

Do Homeowners Care about Sustainability?

IA.1 Data

This appendix supplements Section 2.

IA.1.1 Problems with Levenshtein distance

As discussed in Section 2.1.2, in order to investigate the relationship between property values and energy profiles, we must link each transaction recorded in the HM Land Registry with a valid EPC through address matching. Unfortunately, addresses are not entered consistently within and between datasets. For example, the address `FLAT 42, 16A BROADWAY STREET, 413` may also be recorded as `42 BROADWAY STREET, 16A 413`. One method to link addresses is to use *fuzzy* matching techniques such as the Levenshtein distance, which computes the minimum number of single-character edits (insertions, deletions, or substitutions) required to change one word into the other.

Although programming packages are readily available for implementing such inexact techniques, it is unclear what the correct threshold to set is in order to maximise the ratio of correct to incorrect matches in our use case. For example, consider housing units within the same building: `FLAT 42A, BROADWAY STREET`, `FLAT 42B, BROADWAY STREET`, `FLAT 42B, BROADWAY STREET`, and so on. The addresses of the housing units only differ by a single letter. The minimum threshold that can be set for an algorithm implementing Levenshtein distance to allow for inexact matches is also one. Therefore, all housing units in this building will be identified as the same across and within datasets, as it takes a single substitution to convert one of these addresses to the other.

In addition to being computationally intensive, these techniques are also sensitive to the manner in which addresses are formatted. For instance, `42 BROADWAY STREET, 16B 413` would be considered closer to `42 BROADWAY STREET, 16A 413` than `FLAT 42, 16A BROADWAY STREET, 413`, as deleting `FLAT` requires more operations than replacing `B` with `A`. Furthermore, we find that in several instances, parts of addresses are repeated across fields. For example, `42 BROADWAY STREET, 413` may also be stored as `42 BROADWAY, BROADWAY STREET 413`. Deleting the word `BROADWAY` will take 8 operations, and therefore, the two addresses will be treated as different by the an algorithm with a threshold less than 8. On the other hand, any algorithm with such high threshold will yield a large number of inexact matches. Given the heterogenous nature of real estate data, inexact matches may distort results substantially.

IA.1.2 Exact Address Matching Algorithm

We develop a custom algorithm that produces *exact* matches, but results in a smaller dataset post-compilation. The procedure is relatively straightforward, but requires careful investigation of how the data is stored and the potential errors that can arise when attempting to match entries. We provide a step-by-step outline here.

Addresses in both Price Paid Data and The Energy Performance of Buildings Data are split across multiple subfields. For instance, the Price Paid Data records `PAON` (building number), `SOAN` (apartment number if a property contains multiple housing units), and `STREET`. The Energy Performance of Buildings Data splits addresses into `ADDRESS1`, `ADDRESS2`, and `ADDRESS3`. Upon manual inspection, we find that the manner in which addresses are recorded in the Energy Performance of Buildings Data presents two challenges. First, for certain local authorities, locational identifiers (such as building names) present in `ADDRESS2` are repeated in `ADDRESS1`. We correct for these duplications. Second, `ADDRESS2` often contains the name of the post-town of the property, which is supposed to be in a separate subfield and is not required for matching addresses, as we have information on postcodes, which are exact and more granular than post-towns. We further find that the post-towns mentioned in `ADDRESS2` are often incorrect. This we purge `ADDRESS2` of all post-town names available in the dataset.

In addition, we discover that the addresses in the Energy Performance of Buildings Data may not always uniquely identify a property. This typically occurs when two housing units within the same building omit the Secondary Addressable Object Name (SAON) from their respective addresses. For instance, both `FLAT 12, 20 BROADWAY STREET` and `FLAT 14, 20 BROADWAY STREET` may be recorded as `20 BROADWAY STREET`. To ensure that each property in the mapped (or linked) dataset is uniquely identified, we remove entries with *address keys* that map to more than one *Building Reference Number* (BRN) in the Energy Performance of Buildings Data, as BRNs are uniquely assigned to each property even if the address recorded in the database is not. Out of 17,827,487 EPCs issued between 01 January 2010 and 31 December 2020, there are 14,807,313 unique address keys but 14,960,081 BRNs.

Then, we start by concatenating the address subfields in each dataset in descending order of granularity. Thus, apartment numbers come before property numbers, which in turn come before the street address. We convert the concatenated address to uppercase letters and remove keywords that are commonly omitted between one address and another. These are `FLAT`, `APARTMENT`, and `BUILDING`.[36] Therefore `FLAT 42, 16A BROADWAY STREET, 413` becomes `42, 16A BROADWAY STREET, 413`. We then filter out non-alphanumeric characters (i.e., spaces, punctuations, and special characters are removed) and reorganise the address so that numbers (both with and without an alphabetical qualifier such as `16` or `16A`) are moved in front of words. These operations convert `42, 16A BROADWAY STREET, 413` to `4216A413BROADWAYSTREET`. Finally, we add the formatted text to the postcode of the building producing a unique *address key*, e.g., `NW14SA4216A413BROADWAYSTREET` where `NW1 4SA` is the building's postcode. The main limitation of our algorithm is that we are unable to account for spelling mistakes in addresses.

The HM Land Registry records 9,808,400 transactions between 01 January 2010 and 31 December 2020. Of these, 9,692,971 transactions take place in postcodes for which we have entries in the Energy Performance of Buildings Data. Of these, we are able to uniquely map 7,239,549 transaction entries (73.8% of 9,808,400) using our exact-matching technique.

We could potentially consider using Levenshtein distance on the alphanumeric characters after the non-alphanumeric characters are moved to the front. This would make the algorithm computationally (and memory) intensive, as within each postcode, we will have to compare all addresses with one another. However the main deterrent is that once we process the data and drop records with missing values for the variables used in our analysis, we are only left with 5,451,054 out of 7,239,549 entries. Thus, attempting to increase the number of records matched may only result in a marginal increase in the final regression sample, and therefore, may not justify the increased computational complexity and a potential for inexact matches.

[36] The algorithm can be potentially improved even further by identifying more such keywords.

IA.1.3 Detailed Notes of Feature Selection and Formatting

We are able to uniquely map 7,239,549 transaction entries using our exact matching algorithm. We can classify the set of features in the *mapped sample* into two types: (i) those required to construct dependent variables (i.e., price and total floor area) and independent variables whose coefficients we are primarily interested in (i.e., energy efficiency scores), and (ii) those that act as hedonic controls (e.g., built form, transaction type, age) or facilitate investigative analysis (e.g., utility costs and environmental impact). Since, entries in the mapped dataset are incomplete and contain `null` values, we must trade-off the number of type (ii) features with the total number of entries in the dataset. Note that entries for which a type (i) has a `null` value *must* be removed for analysis, and therefore do not present a trade-off. There is no fixed rule on how to accomplish this. Nonetheless we are able to retain all features that are of first order importance; and are enumerated in Table 2.

Thereafter we format (or clean) the mapped dataset feature-by-feature. This section walks the reader through feature-by-feature implementation details. The compiled dataset post-processing contains 7,022,645 entries, i.e., we loose roughly 1% of 7,239,549 entries from wrangling, formatting, and cleaning. Table 3 provides a quick summary of the key operations in the order in which they are carried out

Energy Ratings. We filter out entries for which the Potential Energy Score is less than that of the Current Energy Score. This results in a loss of 890 entries. We cap the Potential Energy Scores to 100. Figure IA.1 illustrates the distribution of Current and Potential Energy Labels.[37]

Building Characteristics. We restrict ourselves to properties with a Total Floor Area between the 0.01 (34 m^2) and 0.99 (243 m^2) percentiles, which eliminates 23,637 observations. We remove 22 properties with Property Type as "Park Home". We remove 3,023 entries with 0 or more than 12 habitable rooms. In instances where two or more Construction Age Bands are not mutually exclusive, we club them together. For example, we combine "2007 onwards", "2007-2011", and "2012 onwards" into a single category, "2007 onwards".

[37]The figures shown throughout Section 2 are based on the final dataset produced post-processing, and therefore, provide an accurate description of the final dataset used for analysis.

Figure IA.1: Energy Label Breakdown

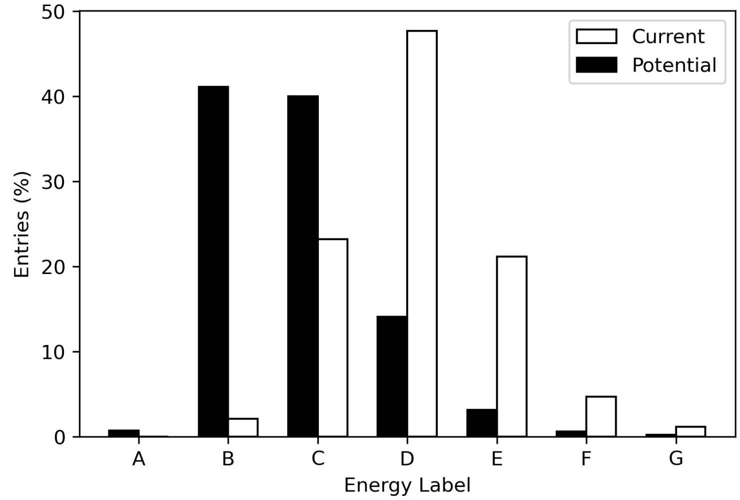

We also rename categories for various categorical variables to make them more readable. For example, "Y"("N") values in feature New were changed to "Yes"("No") and the prefix from Construction Age Bands was removed, rendering "England and Wales: 1900-1929" to "1900-1929".

For Glazed Area, we subsume the 88 "Much Less Than Typical" and 1731 "Much More Than Typical" values into "Less Than Typical" and "More Than Typical" respectively. Given that the Multi-Glaze Proportion for more than 75% of properties is 100%, we convert the feature into a categorical variable, with "High" (≥66.5%), "Low" (≤33.3%), and "Medium" categories. Finally, we cap Low-Energy Lighting Proportion to 100.

Transaction Characteristics. The same category in Tenure is stored in different formats; we clean these category names to obtain three classifications, "Owner Occupied", "Rental (Social)", and "Rental (Private)". We drop 10,669 entries for which Tenure can not be determined. Several categories in Transaction Type have very few entries, or are closely related to one another. We combine "Eco" and "FiT" assessments into a single category; all categories related to "Rental" properties are subsumed into one; "Stock Condition Survey" and other miscellaneous categories are classified as "Other".

Note: When formatting the dataset based on price, cost and energy measures, as described next, we divide them by the Total Floor Area to enable comparisons between properties. We also filter the dataset based on such features at the end as these filters involve elimination of extreme values based on percentiles, which would lead to a higher loss of entries if carried out before the formatting steps carried out in the previous sections.

Price. We remove properties with price per unit area less than 0.0001 or more than 0.9999 percentiles, resulting in a loss of 35,975 entries.

Environmental Metrics. We remove entries outside the 0.0001 and 0.9999 percentile range for Current Environmental Impact, Carbon Emissions per unit area, and Energy Consumption per unit area, resulting in a loss of 2,249 entries. We also filter out properties for which Potential Carbon Emissions and Energy Consumption values are lower and Current ones, and for which the Current Environmental Impact score is higher than the Potential score, resulting in a further loss of 41,179 entries.

Utility Costs. We start by removing properties for which Current Lighting, Heating, and Hot Water Cost is outside the 0.0001 and 0.9999 percentile range, loosing 3,480 entries. Ideally, we would like to filter out all entries where Potential costs are higher than the Current ones; however, doing so results in a loss of roughly 25% of the dataset. Therefore, we introduce a small *threshold* set to 1/20th of the median of cost per unit area. Entries for which Current costs are less than Potential costs minus the threshold are removed. For entries that are within this threshold and have Potential costs greater than the Current ones, we set the Potential costs equal to the Current costs.

For example, we calculate the median of Current Heating Cost per unit area and divide it by 20 to obtain the threshold τ. If Current Heating Cost per unit area is less than Potential Heating Cost per unit area minus τ, we remove the entry. For the remaining entries, if the Potential Heating Cost is greater than the Current Heating Cost, we set it to the Current value. We repeat this process for Lighting and Hot Water costs, resulting in a total loss of 95,780 entries.

IA.1.4 Compiling Multiple Deprivation Indices

The Multiple Deprivation Indices (MDI) are available for the years 2007, 2010, 2015, and 2019. There are four considerations in compiling the indices for analysis. First, the format in which these indices are recorded is inconsistent across reports. Therefore, we manually reorganise the composite Index of Multiple Deprivation (IMD) and the component indices (Income, Employment, Health Deprivation, Education, Crime, Housing Barrier, and Living Environment) into tabulated files with a consistent format so that they can be processed using a script.

The second consideration is that the 2007 and 2010 MDI are reported for the 32,482 LSOA regions constructed in 2001, whereas the 2015 and 2019 MDI are reported for the 32,844 LSOA regions constructed in 2011. We use the LSOA 2001 to LSOA 2011 Lookup table published by the Office of National Statistics to link the two. However, LSOA 2001 to LSOA 2011 conversions are not one-to-one. There are splits (S), merges (M), exact matches (U), and best fits (X). Therefore, we group by LSOA 2011 and take an average. If an LSOA 2001 was split into two zones in 2011, then both zones will have same 2001 entries, and taking average does not impact 2001 scores. If 2001 areas were merged into a 2011 area, then this operation takes an average.[38]

The third consideration is to select one of two formats in which the indices are reported: *scores* or *ranks*. We opt to use ranks in our analysis as they involve fewer mathematical transformations in their construction, are less polarised, and in general, the recommended measure for analysis in government documentation and reports. Figure IA.2 illustrates the distribution of ranks through a heat-map. We normalise the ranks from 0 to 100 by dividing 2007 and 2010 ranks by 32,482, and 2015 and 2019 ranks by 32,844.[39]

Finally, we must interpolate indices for those years between 2010 and 2020 for which we do not have a MDI report. Two natural candidates are *linear interpolation (and extrapolation)* and *stepwise assignment*. The former would be a good approach if the direction of change

[38] Note that taking naive averages is not completely accurate, as ideally, we should weight the average by population, number of houses, or area of the region.

[39] We find that using scores instead of ranks, using non-normalised ranks, or sorting ranks into 100 quantiles does not effect our regression estimates for energy premium in Section 3.

Figure IA.2: IMD Ranks

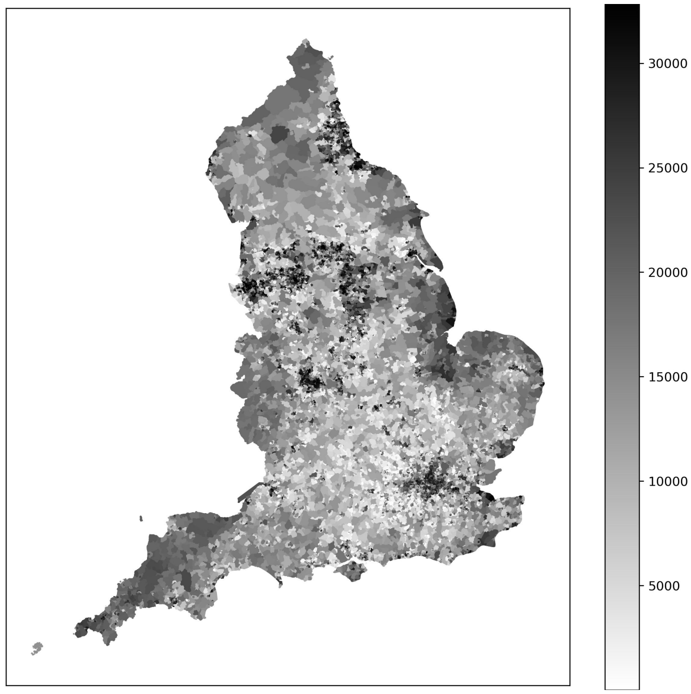

in ranks from one report to another was somewhat predictable. However, we find that this is not the case: ranks for 22.9% of LSOAs continued to increase from 2010 to 2015 and from 2015 to 2019; 21.36% of LSOAs continued to decrease; whereas ranks for 27.04% of LSOAs increased from 2010 to 2015, but decreased from 2015 to 2019; and those for 23.16% of LSOA decreased from 2010 to 2015 but increased from 2015 to 2019. We therefore opt for a stepwise approach, and for each year, assign the rank corresponding to the most recent MDI report. For example, the ranks for 2018 are taken from the 2015 MDI report, and those for 2020 are taken from the 2019 MDI report.

IA.1.5 Constructing Degree Days

For each of one the 10,432 5×5km grids represented by coordinates, we work with average monthly temperature values recorded by the Meteorological Office from January 2007 to December 2020. We calculate degree days (DD^o) for month m in year t for grid g as:

$$DD^o_{gmt} = \max(0, B - T_{gmt}) \times N_m,$$

where B is a pre-specified baseline temperature value, typically set to 15.5 °C, T_{gmt} is the average temperature recorded for month m in year t for grid g, and N_m is the number of days in the month. For each year, we then sum over the monthly degree days to obtain the annual degree day measure $DD^o_{gt} = \sum_m DD^o_{gmt}$.[40] The higher the degree days, the colder the climate, and the more the heating requirements for a given building at a specific location.

It is also useful to think about how Degree Days would factor into a property transaction. The buyer and seller cannot know the aggregate Degree Days for the year in which the transaction occurs. Additionally, an unusually hot or cold year is unlikely to factor into property valuation. Therefore, for each year from 2010 to 2020, we use the average of Degree Days value taken over the preceding three years, denoted by $\overline{DD^o}_{gt} = (1/3)\sum_{k=t-3}^{t-1} DD^o{}_{gk}$. For example, we use the average of degree days from years 2017 to 2019 for 2020, and from 2007 to 2009 for 2010.

Because we have degree days for grids represented by a unique set of coordinates, we use the LSOA 2011 Boundaries dataset published by the Office of National Statistics to extract the representative coordinate for each of the 32,844 LSOA 2011 and assign to them the Degree Days values for years 2008 through 2021 for the grid that is closest in (Euclidean) distance to each LSOA.[41] Figure IA.3 illustrates the average of Degree Degree days over the analysis period (2008 to 2021) for each LSOA 2011 in the UK.

Lastly, because degree days computed in this manner depend on the frequency at which temperature observations are recorded, the unit of measurement is not a "day" and the values should be interpreted *relative* to each other. Therefore, we use `max-min` normalisation to rescale the values between 0 and 100 as follows:

$$DD_{rt} = 100 \times \frac{\overline{DD^o}_{rt} - \min_{rt} \overline{DD^o}_{rt}}{\max_{rt} \overline{DD^o}_{rt} - \min_{rt} \overline{DD^o}_{rt}},$$

where DD_{rt} is the final degree days measure for region r in year t that we use in our analysis.

[40] Typically, these calculations are done on a daily basis, or even an intraday basis, and then aggregated to monthly or annual measures. By using monthly average temperature values instead, we will underestimate degree days, since if the mean temperature of the month is greater than 15.5 °C, the HDD for the month will be 0, but if we used daily data, this might not be the case. Because downloading and processing daily data is significantly more computationally intensive, we opt for the less granular approach.

[41] Typically, Euclidean distance must be avoided in geospatial distance measurements as it does not take into account the curvature of Earth, but since we are interested in the closest match (which is less than 5km here), using Euclidean distance will produce reasonably accurate matches.

Figure IA.3: Average Degree Days

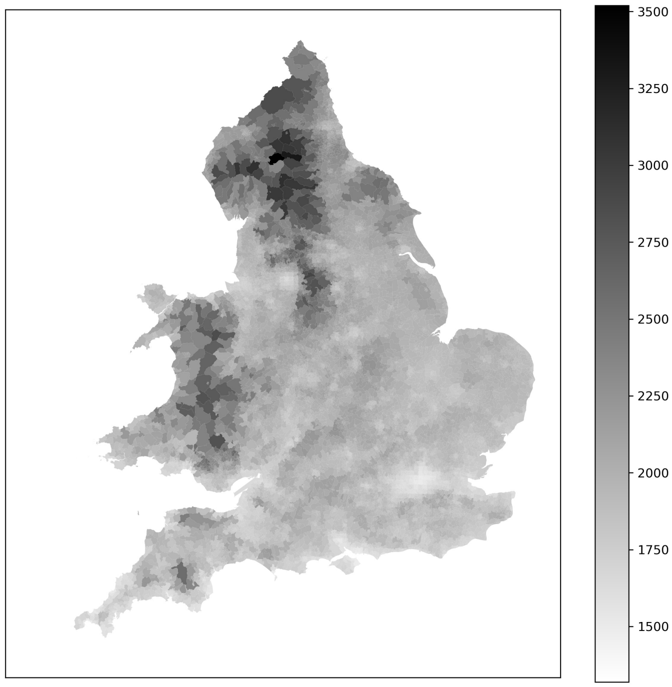

IA.1.6 Supplement to Sample Properties

This section contains figures that supplement Section 2.2.

Figure IA.4: Proportion of New Entries Sample per Quarter

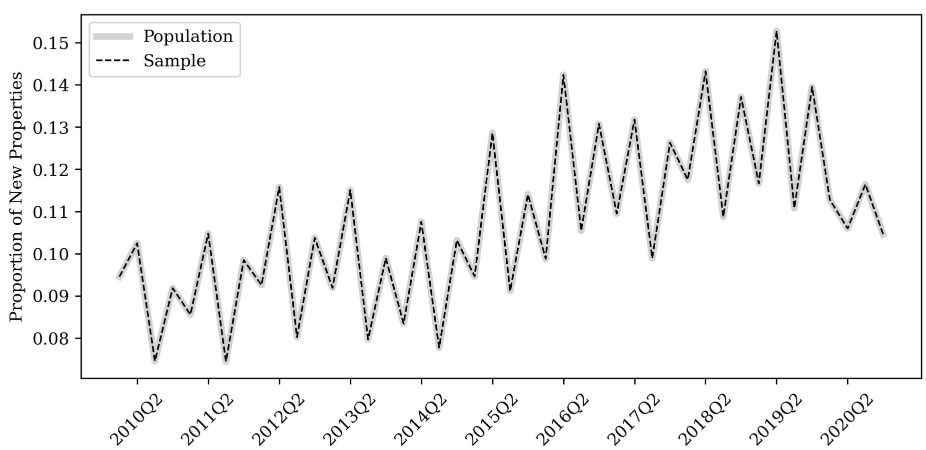

For each quarter, the tick grey line represents the proportion of properties that are marked new in the Price Paid Data. The dashed black line represents the proportion of properties that are marked new in the regression sample.

Figure IA.5: Entries Sampled per Construction Age Band

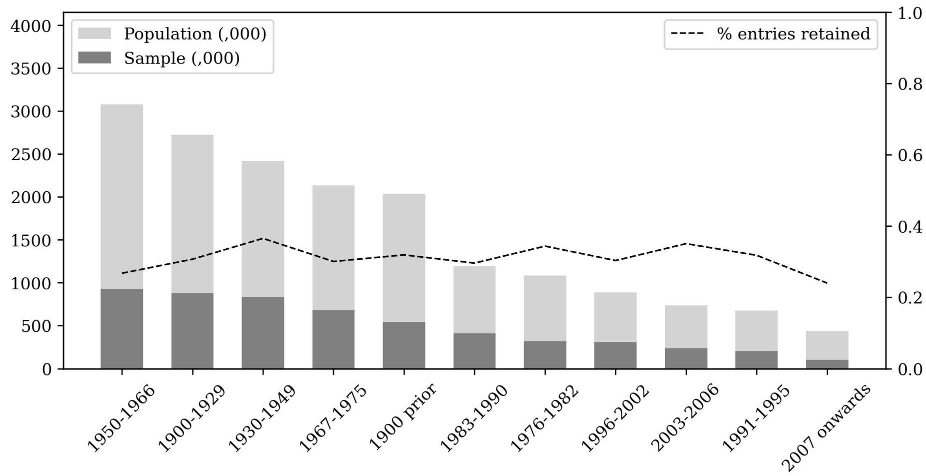

The x-axis marks the various Construction Age Band categories in the Energy Performance of Buildings Data. For each age band, the height of the light grey columns correspond to the primary (left) y-axis, and represent the number of entries present in the dataset, in thousands. The dark grey columns represent how many entries, in thousands, were retained in the regression sample. The black dashed line corresponds to the secondary (left) y-axis and represents the proportion of entries sampled from the Energy Performance of Buildings Dataset for each age band.

Figure IA.6: Population vs. Transactions Sampled by Borough

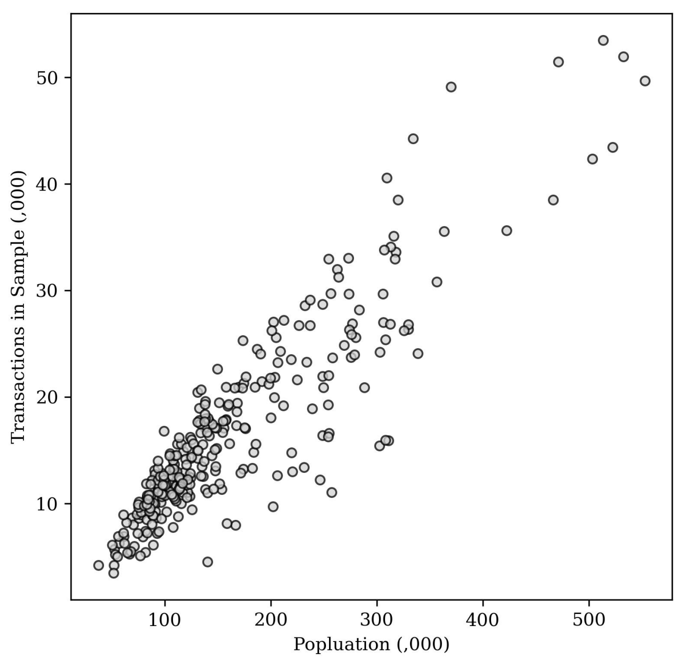

Each point corresponds to one of the 341 local authorities (administrative regions) in the Energy Performance of Buildings Data. The x-axis corresponds the total population in each local authority, obtained from the Rural Urban Classification (RUC) data, published by Department for Environment, Food & Rural Affairs. The y-axis corresponds to the number of transactions that belong to each local authority in the regression sample.

IA.2 Estimation of Energy Premium

This appendix supplements Section 3.

IA.2.1 Consistency, Endogeneity, and Bias

A. Consistency of Premium Across Different Levels of Aggregation

In this appendix, we discuss how the estimates for energy premia reported in Table 5 of Section 3.2 are consistent across different levels of aggregations of energy ratings, and with those observed in the existing body of literature. Columns (a_1) and (a_2) provide estimates for $\beta_{\texttt{Score}}$ in Equation (1a). Column (a_1) excludes properties with energy efficiency labels A (score $\geq$ 92) and G (score $\leq$ 20) while column (a_2) includes them. Columns (b), (c), and (d) provide estimates for $\beta_{\texttt{Label}(i,h)}$, $\beta_{\texttt{Group}(i,h)}$, and $\beta_{\texttt{Class}(i,h)}$ corresponding to Equations (1b), (1c), and (1d) respectively. We observe in column (a_2) that a unit increase in the numerical energy efficiency score is associated with a 0.21% increase in transaction value, holding all else equal. This premium declines to 0.18% in column (a_1) when properties with labels A (score $\geq$ 92) and G (score $\leq$ 20) are excluded. Column (b) tells us that properties with labels A (+6.08%), B (+1.84%) and C (+1.62%) command a premium relative to D, while those with labels E (−2.58%), F (−5.54%) and G (−15.23%) transact at a discount; and we learn from column (d) that when the aggregated, *green* (labels C and higher) dwellings command a 2.25% premium over *brown* (labels D and lower).

We report the estimates for select building properties and transaction controls in Table 6, and those for degree days and the seven multiple deprivation indices in Table 7. As noted in Section 3.2.2, the results for hedonic covariates act as a robustness check across specifications, and we see that the coefficients of numerical features (e.g., Total Floor Area), and the differences in levels of categorical features (e.g., New, Tenure), are consistent across specifications (a) through (d). Similar to estimates for hedonic controls, estimates for region and time fixed effects are consistent across model specifications, and realistic from an economic standpoint.

Notwithstanding, at a first pass, it is quite striking to observe the dramatic increase in the range of energy premium reported in Table 5 as we move from aggregated to more granular energy ratings; from 2.25% in column (d) to 20.79% in column (a_2).[42] However, if we aggregate the estimates the energy premium based on how the energy ratings were aggregated in Table 4, we observe that the estimates are indeed consistent across specifications. We start by noting that our estimates for numerical energy ratings (20.79%) in column (a_2) and alphabetical labels (21.31%) in column (b) are very close. If we take an average of coefficients of labels D and E (0.4665) in column (b), and subtract it from the average of labels B and C (0.4967), we obtain 3.02%, which is in the same ballpark as that of the difference in coefficients of groups BC and DE (2.09%) in column (c).

B. Residuals and Endogeneity

Figure IA.7: Residuals

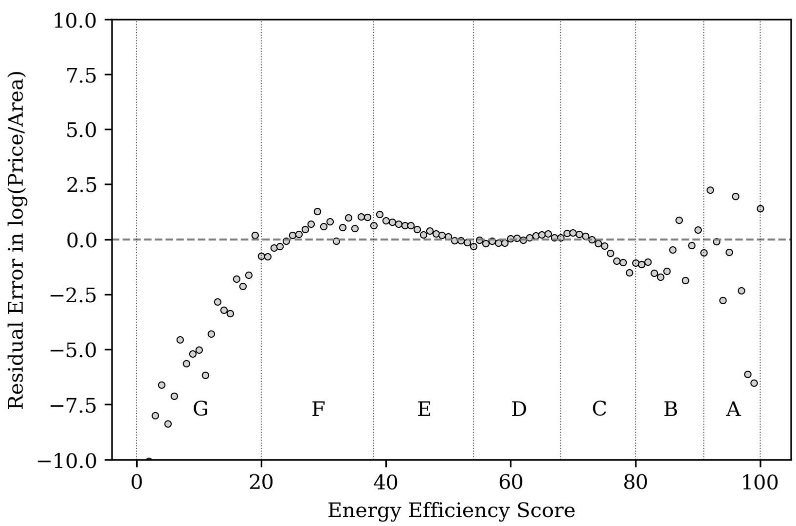

This figures plots the conditional mean of the residual $\mathbb{E}[\varepsilon_{ihrt}|\mathrm{S}_{ih}]$ (y-axis) obtained from Equation (1a) for each numerical energy efficiency score S_{ih} (x-axis). We obtain the conditional expecations as follows. First, we obtain the residual for each transaction in the regression sample by subtracting the actual values of the target variable ($\log(Price/Area)$) from the fitted (or predicted) values. Then we group transactions by their current energy efficiency scores $\mathrm{S}_{ih} \in \{1, 2, ...100\}$. For each of the 100 groups thus obtained, we take the mean of the residuals and then plot them in this figure.

However, taking an average of groups DE and FG and subtracting it from coefficient of group BC yields 5.2% which is high compared to the premium for Green properties (2.25%)

[42] Recall that we define the *range of energy premium* as the difference between the price of a property with highest energy efficiency rating minus the that of the lowest rating, ceteris paribus.

in column (d). The main reason for this are the stark increases in energy premium from energy label G to F (9.69%), and then from B to A (4.24%). A closer analysis reveals that these jumps in energy premia are a consequence of endogeneity. Figure IA.7 shows that $\mathbb{E}[\varepsilon_{ihrt}|\mathrm{S}_{ih}] \neq 0$ for properties with energy labels A (score $\geq$ 92) and G (score $\leq$ 20) when estimating Equation (1a) in Section 3. Recall that exogeneity implies that the conditional expectation of the residual should be zero everywhere. Therefore, restricting our sample to properties with labels F (score $\geq$ 21) through B (score $\leq$ 91) yields an unbiased estimate for $\beta_{\texttt{Score}}$, which is reported in column (a_1) in Table 5. Doing so also reconciles the differences observed in energy premia when comparing differences in estimates across columns (b), (c), and (d).

IA.2.2 Matched Estimate for Energy Premium

This section performs a robustness check by running piecewise linear hedonic regressions (i.e., Equation (1a)) for subsamples sorted by energy efficiency labels. Table IA.1 enumerates the *piecewise estimates*. Then, we take a weighted average of the subsample estimates to arrive at the *matched estimate* for energy premium (each piece-wise estimate is weighted by the proportion of observations present in the subsample), which equals 0.14%.

Table IA.1: Piecewise Regression Estimates

	Estimate	p-value	Obs.	Adj. R^2
A	0.002742	0.52	899	0.78
B	0.003996	0.00	99,990	0.68
C	-0.001126	0.00	1,301,408	0.77
D	0.002267	0.00	2,666,960	0.80
E	0.001626	0.00	1,095,163	0.79
F	0.002570	0.00	234,464	0.76
G	0.005284	0.00	52,170	0.76

The target variable is logarithm of price per unit area. Thus, as estimate of 0.002742 corresponds to a 0.27% change in the target variable per unit change in the explanatory variable.

We draw four insights from this analysis. First, we see that energy premium survives this procedure, and the matched estimate (0.14%) is in the proximity for the value we report in Section 3.2.1 (0.18%). Second, our concerns around endogeneity are further alleviated. We find that not only does energy premia persist across subsamples sorted by energy efficiency labels, the piece-wise estimates are often higher than 0.18%. Third, we see that the piece-wise estimate for subsample corresponding to energy label C is negative. Because government regulations (informally) classify properties with energy labels C or above as brown, a negative estimate suggests that homeowners attach value to the "classification" of the energy label of their properties. Therefore they do not pay a premium for subsequent improvements in energy efficiency once the energy label of a property meets the regulator's threshold. But we do see a significantly high energy premium for subsample corresponding to energy label B (0.40%). Therefore, this reasoning cannot be generalised to all homeowners. Fourth, we observe that energy premium does not systematically goes up or down as we move across subsamples corresponding to labels F through B. This means that a linear specification with the first order term for energy efficiency score is a reasonably good modelling choice. A consistently decreasing or increasing premium would indicate that the model can be potentially enriched by adding a second order term (or a similar monotonic transformation) for energy efficiency score.

IA.2.3 Pricing Potential Upgradeability

In this appendix, we investigate whether *potential upgradeability* is priced by the market. We define potential upgradeability as:

$$\mathrm{U}_{ih} = \frac{y_{ih} - x_{ih}}{c_{ih}},$$

where $x_{it} \in \{1, 2, ..., 100\}$ represents the current energy score of property h associated with transaction $i \in I$ ($|I|$ is the number of entries in the dataset), $y_{it} \geq x_{it}$, $y_{it} \in \{1, 2, ..., 100\}$ represents the potential energy score, and c_{ih} denotes the cost of upgrading the property from x_{ih} to y_{ih}. We then augment Equation (1a) in Section 3.1 as follows:

$$\log{(P/A)_{ihrt}} = \alpha_r + \delta_t + \mathrm{S}_{ih}\beta_{\texttt{Score}} + \mathrm{U}_{ih}\pi + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt},$$

where $\pi \in \mathbb{R}$ is parameter associated with U_{ih} and all other symbols have exactly the same meaning as those in Equation (1a). We find that potential upgradeability is not priced (i.e., $\pi \simeq 0$).

IA.2.4 Alternative Dependent Variable

In this appendix, we estimate Equation (1a) in Section 3.1 by replacing energy efficiency score with the *environmental impact score* of the property. Each property is provides a score between 1 and 100 based on its annual carbon emissions per unit area. Dwellings with higher environmental impact scores generate lower emissions. Therefore, we estimate:

$$\log{(P/A)_{ihrt}} = \alpha_r + \delta_t + \mathrm{E}_{ih}\beta_{\texttt{Impact}} + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt},$$

where E_{ih} is the environmental impact score of property h associated with transaction $i \in I$ ($|I|$ is the number of entries in the dataset), $\beta_{\texttt{Impact}} \in \mathbb{R}$ is the parameter associated with E_{ih}, all other symbols have exactly the same meaning as those in Equation (1a).

Next, we run year-wise subsample regressions to track the evolution of $\beta_{\texttt{Impact}}$ over the duration of our sample; that is, for each t = 2009, 2010,..., 2021, we estimate:

$$\log{(P/A)_{ihrt}} = \alpha_r + \mathrm{E}_{ih}\beta_{\texttt{Impact},t} + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt},$$

where the unit of observation is transaction $i \in I_t$ (where $|I_t|$ is the number of entries in the dataset for year t) of property h, in region r, at time t. We include the subscript t in parameter associated with environmental impact score, $\beta_{\texttt{Impact},t}$, as it is now specific to the year for which the regression is run. All other symbols have exactly the same meaning as those in Equation (3) in Section 3.1.

Figure IA.8 plots the period-wise subsample estimates associated with environmental impact score ($\beta_{\texttt{Impact},t}$) together with those for energy efficiency scores ($\beta_{\texttt{Score},t}$ in Equation 3). Because environmental impact scores are used to compute energy efficiency scores, both measures are highly correlated (0.95). Therefore, the regression estimates for both dependent variables are similar.

Figure IA.8: Alternative Dependent Variable

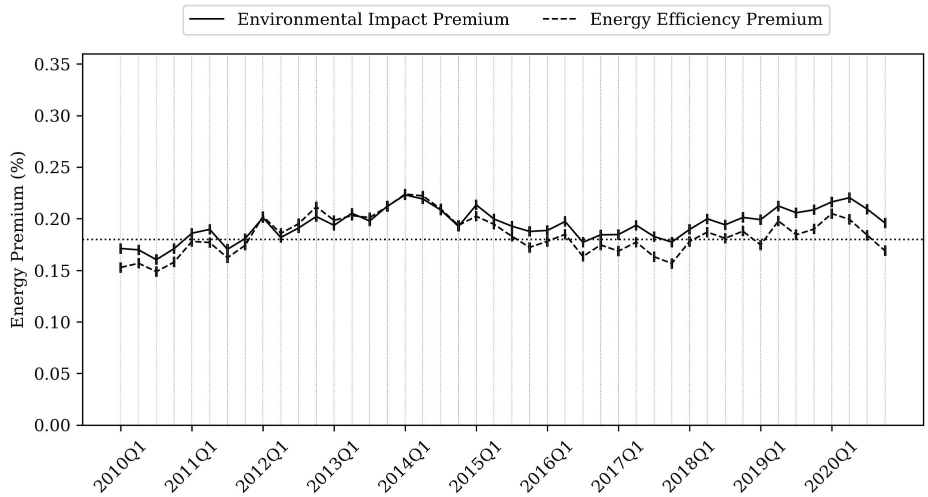

This figure tracks the evolution of energy premium over the duration of our sample. The solid line plots estimates obtained from year-wise subsample regressions in Equation (3). The dashed line plots the estimates obtained from the time-interacted effects model in Equation (4).

IA.3 Heterogeneity in Energy Premium

This appendix supplements Section 4. Table IA.2 reports the energy premium estimates for subsamples sorted by tenure and rural-urban classifications. Table IA.2 reports energy premium estimates reported for subsamples sorted by tenure and property type.

Table IA.2: Estimates for Subsamples across Tenure and RUC

Tenure	RUC	Estimate	Error	p-Value	N.Obs.	Adj. R^2
Owner Occupied	1	0.20	0.00	0.00	421540	0.56
	2	0.22	0.00	0.00	592,109	0.70
	3	0.22	0.00	0.00	642,481	0.74
	4	0.23	0.00	0.00	1,299,944	0.77
	5	0.33	0.01	0.00	158,108	0.66
	6	0.24	0.00	0.00	1,476,547	0.84
Rental (Private)	1	0.09	0.02	0.00	34556	0.53
	2	0.07	0.01	0.00	49,199	0.70
	3	0.12	0.01	0.00	58,190	0.72
	4	0.09	0.01	0.00	152,445	0.78
	5	0.17	0.02	0.00	18,554	0.58
	6	0.13	0.01	0.00	230,313	0.84

Table IA.3: Estimates for Subsamples across Tenure and Property Type

Tenure	Property Type	Estimate	Error	p-Value	N.Obs.	Adj. R^2	MS	MC
Owner Occupied	House	0.23	0.00	0.00	2,836,170	0.81	0.01	0.49
	Flat	0.14	0.01	0.00	312,204	0.79	0.01	0.23
	Bungalow	0.28	0.00	0.00	433,937	0.72	0.01	0.40
	Maisonette	0.27	0.01	0.00	47,859	0.80	0.01	0.21
Rental (Private)	House	0.16	0.01	0.00	197,350	0.84	0.01	0.65
	Flat	0.11	0.01	0.00	82,625	0.77	0.01	0.23
	Bungalow	0.19	0.02	0.00	14,034	0.71	0.01	0.45
	Maisonette	0.05	0.04	0.18	9,300	0.74	0.01	0.22

MS (MC) refer to marginal savings (costs).

IA.4 Computation of Green Premium

This appendix supplements Section 5. Figure IA.9 plots the period-specific mean and median measures of marginal energy savings m_t^u obtained from Equation (10).

Figure IA.9: Evolution of Marginal Energy Savings

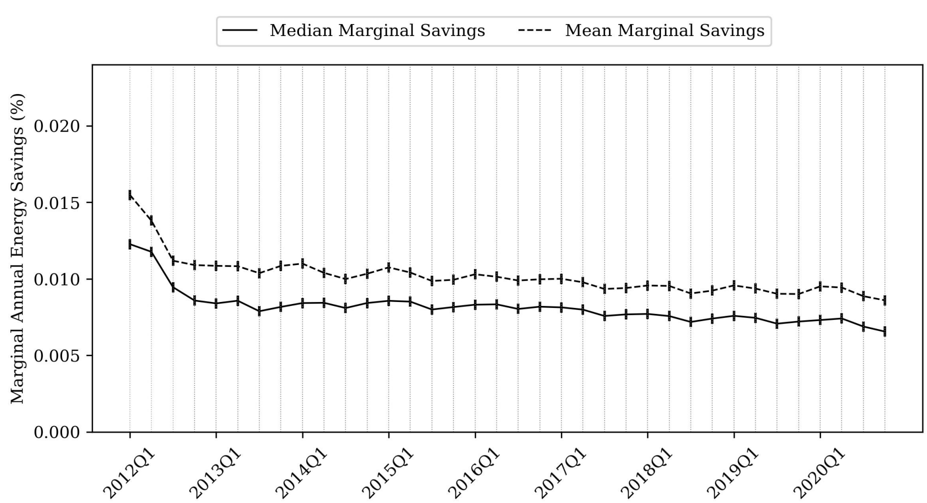

For each quarter, the solid (dashed) line plots the median (mean) marginal energy savings m_t^u obtained from Equation (10) in Section 5.1.

Figure IA.10 plots the implied marginal net discount rates based on non-constant marginal energy premia estimated by including a second-order term for energy efficiency score in Equation (1a) as follows:

$$\log{(P/A)}_{ihrt} = \alpha_r + \delta_t + \mathrm{S}_{ih}\beta_{\texttt{Score}} + \mathrm{S}_{ih}^2\eta_{\texttt{Score}} + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt},$$

where $\eta_{\texttt{Score}}$ is the coefficient of the second order term for energy efficiency score S_{ih}^2; all other symbols have exactly the same meaning as those in Equation (1a). The energy premium is no longer independent of the initial level of energy efficiency; this can be seen by taking the derivative of the target variable with respect to S_{ih}. Let $\beta(s)$ denote the marginal energy premium associated with property with energy efficiency score s. Then, we have:

$$\beta(s) = \beta_{\texttt{Score}} + 2 \times s \times \eta_{\texttt{Score}}$$

We obtain the marginal implied net discount rate associated with property with energy efficiency score s as $m^u(s)/\beta(s)$, where $m^u(s)$ is obtained from Equation (9). The dashed line in Figure IA.10 plots the marginal discount rates obtained from the conditional energy premium $\beta(s)$ while the solid black line plots those obtained from using a constant premium of 0.18%. We find that our conclusions in Section 5.2 remain unchanged.

Figure IA.10: Evolution of Marginal Energy Savings

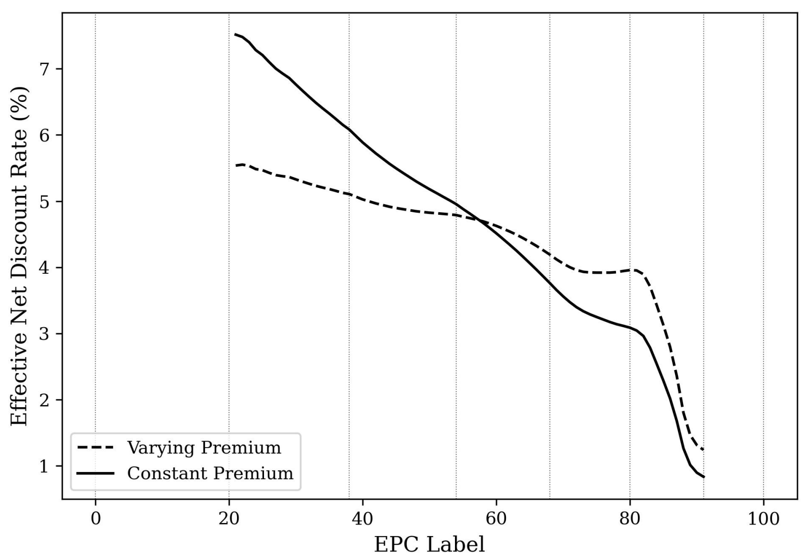

IA.5 Energy Upgrades and Policy Impact

This appendix supplements Section 6.

IA.5.1 Alternative Specifications for Policy Impact

In this section, we discuss additional Difference in Difference (DD) and Regression Discontinuity (RD) specifications to corroborate our findings in Section 6.4.3.

We begin my describing alternative model specifications for the Difference in Difference (DD) model proposed in Section 6.4.1. In order to control for confounding impact of CGS in Equation (15) and to address concerns around potential violations of the parallel trends assumption, we (i) restricted our sample to properties with labels D or below (score < 69) and then (ii) introduced a fixed effect for CGS, resulting in Equation (16). An alternative method would be to eliminate the confounding impact of CGS by following the same approach as we did to account for the confounding impacts of the financial crisis and the coronavirus pandemic; i.e., restricting the duration of the sample between 01 January 2011 (post-financial crisis) and 11 October 2017, the day before CGS was published. Doing so eliminates the need to include a fixed effect for CGS and results in a specification identical to Equation (15):

$$\begin{aligned}\log{(P/A)_{ihrt}} = \alpha_r + \delta_t + \mathrm{S}_{ih}\beta_{\texttt{Score}} + (\mathbb{1}_t.\mathrm{S}_{ih})\beta_{\texttt{Score},t} + \lambda_{\texttt{MEES}} + \rho_{\texttt{MEES}}X_i \\ + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt},\end{aligned} \tag{IA1}$$

where symbols have exactly the same meaning as those in Equation (15). However, one might suspect that the negative impact of MEES on property prices with energy labels below E comes into play *after* the policy was implemented on 01 April 2018. Therefore, we rerun the specification outlined in Equation (IA1) by restricting the sample from 12 October 2017 to 31 December 2019 with 01 April 2018 as policy-cutoff. Hence, $X_i = 1$ when transaction i takes place on or after 01 April 2018 *and* the energy efficiency score of the underlying property is less than or equal to 38; $X_i = 0$ otherwise. Similarly, $\lambda_{\texttt{MEES}}$ is the coefficient of an indicator variable which is equal to one when a transaction occurs on or after 01 April 2018, and zero otherwise. As before, we eliminate the need to include the fixed effect for CGS; this time, by forcing the sample to begin at 12 October 2017, the date when CGS was published. In both instances, we conclude that MEES was unsuccessful, which is consistent with our results in Section 6.4.3.

Further, we had observed that the fixed effect for CGS in Equation (16), $\lambda_{\texttt{CGS}}$, was positive, which is contrary to our expectations. By (i) restricting the sample from 26 March 2015 (when MEES was announced) to 31 December 2019 (pre-pandemic) with 12 October 2017 (CGS publication) as policy-cutoff date, and (ii) including properties with energy efficiency scores both above and below 69, we specify a DD model to study the impact of CGS on "green" properties (score $\geq$ 69), as follows:

$$\begin{aligned} \log{(P/A)_{ihrt}} = \alpha_r + \delta_t + \mathrm{S}_{ih}\beta_{\texttt{Score}} + (\mathbb{1}_t.\mathrm{S}_{ih})\beta_{\texttt{Score},t} + \lambda_{\texttt{CGS}} + \rho_{\texttt{CGS}}X_i \\ + \mathrm{B}_h^T\theta + \mathrm{T}_i^T\gamma + \mathrm{MDI}_{rt}^T\nu + \mathrm{DD}_{rt}\omega + \varepsilon_{ihrt}, \end{aligned} \tag{IA2}$$

where $\lambda_{\texttt{CGS}}$ is the fixed effect of CGS and, X_i is the indicator variable for treatment; $X_i = 1$ when transaction i takes place on or after 12 October 2017 *and* the energy efficiency score of the underlying property is greater than or equal to 69; $X_i = 0$ otherwise. All other symbols have exactly the same meaning as those in Equation (4). In order to conclude that CGS was successful, we expect $\rho_{\texttt{CGS}} \in \mathbb{R}$, the causal effect of interest, to be positive and statistically significant. We find that CGS did not lead to a premium for properties with labels C or above (score $\geq$ 69), or conversely, brown properties did not incur an additional discount post-CGS. These findings our consistent with the non-negative fixed effect for CGS obtained for Equation (16).

Section 6.4.3 also provides potential explanations for why the policies were unsuccessful. In particular, one potential reason could be that MEES had a negative and significant impact on transaction values of only those properties that do not qualify for policy exemptions, such as those listed in Table 9. Based on our dataset, it is not possible to identify the subset of properties that do not qualify for all potential MEES exemptions that homeowners can file for. Nonetheless, to check the robustness of our results, we make a partial attempt by estimating Equations (16), (IA1), and (IA2) by restricting our sample properties with energy efficiency labels F, and then assigning treatment to properties that (i) can be upgraded to an energy efficiency score greater than or equal to E (score $\geq$ 39), and for which (ii) the expected cost of upgrade to a score of 39 is less than £3,500. Restricting our sample to properties with energy labels F strengthens the parallel trends assumption as the counterfactual price evolution of the treatment group should be exactly the same as that of the control group in the absence of MEES exemptions. However, our conclusions remain robust to the results obtained from these alternative specifications.

Lastly, we test an alternative Regression Discontinuity (RD) specification where we restrict ourselves to a sample or properties with energy efficiency scores less than 39, and assign treatment to properties that (i) can be upgraded to an energy efficiency score greater than or equal to E (score ≥ 39), and for which (ii) the expected cost of upgrade to a score of 39 is less than £3,500. Because it is not possible to determine whether a dwelling will be exempt from MEES prior to its publication (and therefore, the criteria for exemptions), we can assume random assignment of treatment. However, we find that our conclusions remain unchanged from those discussed Section 6.4.3.

IA.5.2 Supplementary Figures and Tables

Figure IA.11: Energy Labels Composition Over Time

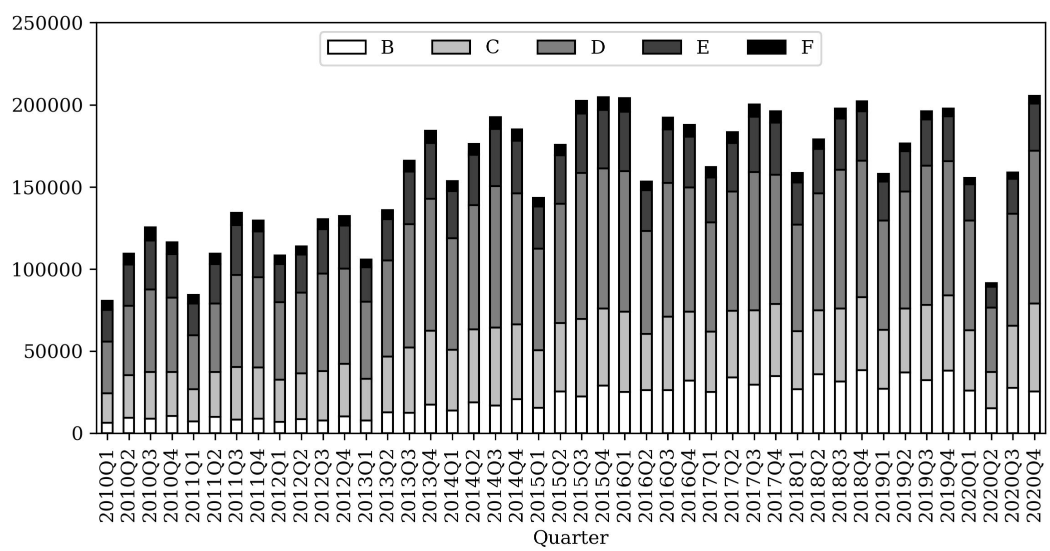

Table IA.4: Multiple Transactions in Merged Sample

Transactions	Address Key	Proportion (%)
1	4,663,761	80.83
2	970,760	16.83
3	124,103	2.15
4	10,107	0.18
5	734	0.01

Table IA.5: Multiple Transactions in Price Paid Data

Transactions	Address Key	Proportion (%)
1	6,355,649	79.96
2	1,381,325	17.38
3	190,622	2.40
4	18,398	0.23
5	1,991	0.03

Table IA.6: Multiple Certificates in the Energy Performance of Buildings Data

Certificates	Building Reference Number	Proportion (%)
1	12,919,914	78.90
2	2,871,736	17.54
3	461,353	2.82
4	91,394	0.56
5	21,713	0.13
>5	8,388	0.05