Stranded Houses?
The Price Effect of a Minimum Energy Efficiency Standard

Konstantinos Ferentinos
Lancaster University

Alex Gibberd
Lancaster University

Benjamin Guin
Bank of England *

February 28, 2023

Abstract

Climate policies aimed at reducing greenhouse gas emissions can lead to decreases in the values of assets, risking them to become "stranded". We study the price effects of the introduction of a specific climate policy, a minimum energy efficiency standard, in the housing market. Leveraging a unique data set of the population of all residential house transactions in England and Wales, we show that prices of energy-inefficient properties affected by this policy decreased on average by about £5k to £9k relative to efficient ones. We interpret this evidence as being consistent with semi-strong market efficiency in the housing market.

Keywords: Climate Policy, House Prices, Energy Efficiency

JEL Classification: C54, Q54, Q58

*We are grateful for comments from Thiemo Fetzer, Bill Francis, Lucas W. Davis, Simon Hall, Daire McCoy, Jennifer Owladi, Martin Sheppard, Richard S.J. Tol, Veronica Veggente and a reviewer of the Staff Working Paper Series as well as two journal reviewers. Any views expressed are solely those of the authors and should not be taken to represent those of the Bank of England or as a statement of Bank of England policy. This paper should therefore not be reported as representing the views of members of the Monetary Policy Committee, Financial Policy Committee or the Prudential Regulation Committee.

1 Introduction

While residential housing is a major vehicle of household wealth accumulation (Bhatia, 1987; Bach et al., 2020), it accounts for a significant share of greenhouse gas emissions. For example, in the United Kingdom the house stock contributes 15% of overall emissions[1] which is a significant contribution to climate change. Public policies have been proposed to mitigate these emissions. However, they can potentially decrease prices of affected properties, risking them to become "stranded" (Caldecott et al., 2016).[2] To the best of our knowledge, evidence on the adjustment of house prices to climate policies is non-existing. This lack of evidence may be due to a lack of available data (NGFS, 2021) and as a large number of climate policies will only be implemented in the future (UNEP, 2021).

This paper addresses this research gap in the context of a specific policy intervention in the UK housing market targeting houses that were the least energy-efficient. This Minimum Energy Efficiency Standard (MEES) which came into force in England and Wales on 1 April 2018 aimed at encouraging landlords and property owners to improve the energy efficiency of their properties in order to reduce overall greenhouse gas emission. It restricted the granting and continuation of existing tenancies for energy-inefficient properties, those with Energy Performance Certificate (EPC) ratings of F and G. Otherwise fines of up to £5,000 could be applied to non-compliant landlords.[3]

In this paper, we examine how sales prices of properties adjust to the implementation of this policy. We expect prices of the least energy-efficient properties which are affected by the policy to decrease relative to properties that are unaffected. To test this hypothesis, we construct a novel data set of the population of all property transactions in England and Wales which we match with the Energy Performance Certificate (EPC) ratings of these buildings. The panel dimension of our data allows us to track the prices of the same properties over time. This mitigates concerns around confounding factors such as economic conditions or the composition of the housing stock changing over time.

We conduct a *Difference-in-Difference* (DiD) analysis which we combine with *Propensity Score Matching* (PSM). As a first step, we conduct the PSM to ensure that we compare only properties that are like-for-like. We match each property affected by the MEES 2018 with a "twin" property, a different property that is otherwise very similar in terms of observable characteristics, such as the property size or the region it is located in. We then gauge the effect of the policy on property prices, by conducting a DiD analysis in which we compare how transaction prices of properties affected by this policy change relative to those unaffected.

Our results suggest that prices of properties affected by the MEES 2018 decreased by about £5,000 to

[1] Source: https://www.gov.uk/government/statistics/final-uk-greenhouse-gas-emissions-national-statistics-1990-to-2018, accessed on 09 November 2020.

[2] We acknowledge that there are different sources of stranding of assets. These include changing resource landscapes, i.e. price and availability of different natural resources, as well as new government regulations (e.g. carbon pricing or minimum standards); a change in demand (e.g. a shift towards renewable energy because of lower energy costs); or even legal action against high emitters (Caldecott et al., 2015; Matikainen, 2022). In comparison to the traditional use of the term "stranded assets", from a policy perspective is neither desirable nor feasible to render a large proportion of the housing stock unusable or stranded. In case of of a minimum energy efficiency standard, there is a risk of stranding if properties' energy efficiency cannot improve in response to the intervention. We acknowledge that other government interventions, e.g. subsidies, can counteract the risk of stranded assets.

[3] I.e. we study one specific policy tool with modest calibration and limited jurisdiction.

£9,000 relative to unaffected ones. The magnitude of this effect compares well to our priors. As energy performance certificates are publicly available, we expect potential buyers to consider them when making their purchase decisions. The degree of the price decrease should reflect the expected cost necessary to improve the energy efficiency of the property.[4] and the possible fine if a landlord does not comply. Costs are initially capped at £3,500 but homeowners will need to incur the remaining costs and improve the energy efficiency of the property within the following 5 years. If landlords, i.e. the owners of properties that let them out, are unwilling to improve the energy efficiency, they will face a fine up to £5k and will need to incur the cost of improvement anyway. Hence, we expect the magnitude of the price discount to be in the range £5k to £9k. Overall, we take this result as evidence in favour of semi-strong market efficiency (Fama, 1970), as real estate markets seem to price in the publicly available information about the energy efficiency of the underlying property. We also discuss several further consequences from the introduction of the MEES 2018 policy intervention on wealth inequality and financial instability which go beyond the average price effect.

The present paper contributes to three strands of the literature. First, we document how a climate policy can lead to adjustments of asset prices, in this case transaction prices of houses. On the one side, there is growing evidence that risks from climate policy is priced into corporate loans (Delis et al., 2021; Fard et al., 2020; Javadi and Masum, 2021) in particular if lenders have an informational advantage screening such loans (Degryse et al., 2020). This can lead to reallocation of exposures away from risky borrowers (Reghezza et al., 2021). On the other side, there is a growing body of evidence that a carbon premium exists in stock markets. Investors already demand compensation for their exposure to carbon emission risk in stocks (Bolton and Kacperczyk, 2021a) as well as protection against downside risks (Ilhan et al., 2020). Evidence suggests that a significant group of investors take a long-term perspective assessing stock prices (Ramelli et al., 2021; Bolton and Kacperczyk, 2021b) which may lead to a green paradox of increases in short-term emissions (Diluiso et al., 2020; Lemoine, 2017; Di Maria et al., 2014). Whether risks are fully priced in depends on whether the announcement of the policy constitutes a surprise or if it was anticipated (Van der Ploeg and Rezai, 2020). Large, unanticipated changes in climate policy that may result in the stranding of carbon-intensive assets (Rozenberg et al., 2020; Van der Ploeg and Rezai, 2020). Indeed, there is growing evidence that investors react to short-term changes in the probability of climate policy proposals in Germany (Sen and von Schickfus, 2020) and in the US (Carattini and Sen, 2019) which can increase the likelihood of transition risk from sudden adjustments of asset prices (Carattini et al., 2021; Battiston et al., 2017). However, this strand of the existing literature has largely ignored the potential of stranded assets in the housing market (Edenhofer et al., 2020).[5] The present paper addresses this gap in the context of one specific policy, a minimum energy efficiency standard. Examining the MEES 2018, we find a negative effect on transaction prices in an order of magnitude to be expected from the calibration of fines and exemptions of the policy.

Second, by providing such evidence, the present paper also contributes to a fundamental question in economics and finance: whether and to what extent markets are efficient. Going back to Mandelbrot (1963)

[4]EPC certificates also contain recommended measures and their anticipated costs.

[5]Edenhofer et al. (2020) argues that this is surprising as real estate constitutes a large and idiosyncratic household wealth component and the largest component of an economy's capital stock.

and Samuelson (1965), there has been an ongoing and unresolved debate whether publicly available information are reflected in asset prices, something Fama (1970) called semi-strong information efficiency. Empirical evidence on the efficiency of the real estate market is ambiguous casting doubt on its efficiency (Case and Shiller, 1989, 1990), see Herath and Maier (2015) for a meta-analysis of the literature.[6] The MEES 2018 offers an ideal quasi-natural experimental set-up to revisit the question of efficiency of real estate markets: First, energy performance certificates are publicly available and, hence, accessible to all market participants. Second, by setting fines and the value of exemptions from the policy, there are clear lower bounds on the expected effect on transaction prices. Hence, we have a clear theoretically-founded empirical prior of the order of magnitude of the effect of the policy on property prices which is testable. Thereby, it relates to the literature providing evidence that the availability of energy efficiency ratings is reflected in property prices (Eichholtz et al., 2010; Walls et al., 2017; Kahn and Kok, 2014) documenting that higher ratings are associated with higher prices (Fuerst et al., 2015) possibly because home buyers being attentive to fuel costs (Myers, 2019). Our paper complements this literature by showing that a salient policy change, the introduction of a minimum energy efficiency standard, can affect house prices.[7]

For our empirical analysis, we employ a Difference-in-Difference design. Its validity relies on one key identifying assumption, the parallel trends assumption. It requires that in the absence of the policy intervention, the difference in the average of the outcome variable, in our case the transaction price of the property, between properties affected by the policy and the ones unaffected would have evolved in parallel (e.g. Lechner, 2010).[8] Researchers often assess the plausibility of this parallel trends assumption by examining the pre-trends, i.e. the differences in trends between the treatment and control group prior to the policy intervention. But standard approaches to verifying it do not typically go beyond visual inspections or testing for differences in linear pre-intervention trends (Ryan et al., 2019; Roth, 2020; Rambachan and Roth, 2019). In fact, out of the 16 papers in the 2016 American Economic Review that use a linear panel data model, eleven are concerned with the existence of pre-trends as a sign of endogeneity. Of these papers, 9 include a plot of pre-trends, of which provide a formal test of whether pre-trends are zero (Freyaldenhoven et al., 2019). Given the importance of this assumption for the validity of the DiD inference, any visual conclusion should ideally be verified under a statistical framework. Moreover, statistical tests for differences in pre-intervention outcome trends between treated and non-treated units, should not make overly restrictive assumptions about the functional form of the trend. Instead, a more realistic approach is to avoid the assumption of a linear relationship and allow the observed data to determine the shape of the trend in the outcome variable. To this end, we here present a novel application of *Generalized Additive Models (GAMs)* (Hastie and Tibshirani,

[6] The existing literature on housing market efficiency can be grouped into several dimensions: by tests of semi-strong vs. weak market efficiency, the type of property (residential vs. commercial) and by evidence in favour or against efficiency. The majority of studies does not find evidence of semi-strong efficiency neither in the residential segment (Case and Shiller, 1990; Clayton, 1996; Barkham and Geltner, 1996; Fu and Ng, 2001) nor in the commercial market segment (Fisher et al., 2009; Barkham and Geltner, 1995) with only little competing evidence (Anas and Eum, 1984). Similarly, there is little evidence on weak efficiency for both residential properties (Case and Shiller, 1989; Englund et al., 1999; Clapp et al., 1995; Dolde and Tirtiroglu, 1997) and commercial ones (MacKinnon and Al Zaman, 2009) with only little competing evidence on the first (Rosenthal, 2006) and the latter (Bardhan et al., 2008).

[7] Thereby it addresses the research gap on resolving known agency problems, e.g. between homeowners and renters, which are pointed out by Fowlie et al. (2018), Gillingham et al. (2012) or Myers (2020) among others.

[8] Hereby, it indirectly tests whether there are anticipation effects prior to the implementation of the policy.

1986), allowing for smooth non-linear trends over time.

The remainder of the paper is as follows: Section 2 summarises the policy background; section 3 offers a conceptual framework. Section 4 details the data used in our analysis; section 5 presents the key methodology used, alongside results. Last, section 6 provides a discussion of the results and conclusion.

2 Policy Background

The 2018 Minimum Energy Efficiency Standard initiative (MEES 2018) came into force in England and Wales on 1 April 2018 under the Energy Efficiency (Private Rented Property) Regulations 2015 (UK Statutory Instruments, 2015). It is best understood as a part of a large-scale market transformation strategy, that uses the EPCs as a springboard to transform the overall UK housing market and reduce greenhouse gases emissions due to the building stock's heavy energy use.

The policy supported two statutory objectives of UK government. First, it aims to reduce energy demand and greenhouse gas emissions. By improving the energy efficiency of privately rented homes, this policy would cut energy use and the greenhouse gas emissions, contributing to the government's climate change commitments. Second, it is also intended to tackle fuel poverty. Raising energy efficiency standards to an EPC rating of E by 2020, mirrors the government's interim target to raise as many fuel poor homes in England to a rating of E by the same date (Department for Business, Energy & Industrial Strategy, 2018).

The policy aimed to address a range of market failures and barriers to energy efficiency improvements which provided a rationale for government intervention in the private rental market. First, there was a concern of misaligned incentives as the costs of upgrading a property fall to landlords but the benefits of lower energy costs and/or a warmer home accrue to the tenant. The landlord was not necessarily able to capture the benefits through increases in rent. Second, there were possible externalities, such as energy prices not fully reflecting the climate change costs of burning fossil fuels, or the public health benefits of warmer homes not fully accruing to those who pay for energy efficiency upgrades. Third, there was a concern of incomplete information as landlords or tenants might not have a good understanding of the benefits of energy efficiency (Department for Business, Energy & Industrial Strategy, 2018).

The MEES 2018 initiative means that, since 1 April 2018, landlords of domestic private rented properties in England and Wales are prohibited to let a tenancy to new or existing tenants if the property that is to be let has an EPC rating of F or G (Department for Business, Energy & Industrial Strategy, 2017, p. 11). This minimum standard takes effect from the point at which a new tenancy is issued, or where an existing tenancy is renewed. For properties that have an EPC rating below E, the landlord needs to improve the energy performance of the property in order to ensure that the standard of EPC E is met or exceeded (Department for Business, Energy & Industrial Strategy, 2017, p. 10).

The MEES 2018 includes an exemption. The requirement to meet the minimum level of energy efficiency (EPC E) does not apply where a landlord has made all the relevant energy efficiency improvements to the property that can be made within a £3,500 cap and the property remains sub-standard (Department for

Business, Energy & Industrial Strategy, 2017, p. 57).[9] The exemption remains valid for five years. After that, the landlord must implement the required improvements of the property, so as to comply with the policy. In cases where the property does not meet the criteria for an exemption and the landlord does not comply with the policy, they may face a penalty of up to £5,000 (Department for Business, Energy & Industrial Strategy, 2017, p. 89). A further exemption may be registered after 5 years. It means that a buyer still has the option of holding off potentially very large capital investments for a longer period of time.[10]

The original regulation also included a "no cost to the landlord" exemption. It permitted landlords to register an exemption where they were unable to make improvements to their sub-standard property at no cost to themselves (Department for Business, Energy & Industrial Strategy, 2017, p. 57). Instead, they could use external funding, i.e. by the central government, a local authority or any other person, either separately or in combination. Following an amendment to MEES regulations, the "no cost to the landlord" provision became unavailable from 1 April 2019. After this date, landlords had to use their own funding to cover the cost of improving their property to EPC E, still subject to a spending cap of £3,500 on each property. Those who had already registered for the "no cost to the landlord" exemption prior to regulation changes, would no longer be exempt for five years. Any "no cost to the landlord" exemption registered between 1 October 2017 and 31 March 2019 would end on 31 March 2020. Therefore, landlords who had registered such "no cost to the landlord" exemptions must make the necessary improvements to ensure their property meets EPC E by 1 April 2020. All landlords who had registered "no cost to the landlord" exemptions prior to 1 April 2019 have been contacted personally via the PRS Exemptions Register to alert them in good time to the adjusted exemption length. The PRS Exemptions Register is updated by the government so that all "no cost to the landlord" exemptions are automatically cancelled on 31 March 2020 (Department for Business, Energy & Industrial Strategy, 2017, p. 58).

From 1 April 2020, the standard was expanded to all relevant properties, even where there has not been any change in tenancy. Landlords must not continue letting a relevant domestic property which is already let if that property has an EPC rating F or G (as shown on a valid EPC for the property). This change of policy is not part of our sample period.

In the same time period, there were changes to the taxation of buy-to-let (BTL) markets. First, there were changes to the taxation of income from BTL properties. Before 6 April 2017, landlords were able to

[9] In other words, even if the measure or package of measures purchased and installed by the landlord does not improve the property to EPC E, the landlord does not need to take any further action if there are no additional measures which can be selected without pushing the overall costs above the £3,500 cap.

[10] As of early March 2020, 9,269 exemptions had been registered with the government for domestic properties. Hence, their magnitude is small relative to the total number of properties with EPC ratings of below E. We discuss how this amendment affects our analyses in the robustness section.

deduct any of their mortgage interest payment from their rental income before paying tax.[11][12] Both tax changes in the BTL market apply uniform across all EPC ratings. Hence, we expect neither of these changes to affect our results.

3 Conceptual Framework

In this section, we provide a conceptual framework which discusses the implications of this policy for equilibrium prices in the housing market. Our framework is in the spirit of a discounted cash flow model which has been shown to be suitable for real estate pricing (Ghysels et al., 2007). We take the view of a prospective property buyer of an energy-inefficient property and ask how their willingness to pay decreases in light of the policy intervention given the magnitude of exemptions and penalties. We then discuss the implications for transaction prices.

First, we consider a risk-neutral buyer who wants to ultimately let the property. This buyer considers purchasing a property with an EPC rating of below E. Initially, the costs of home improvements are capped at £3,500 by the policy. The buyer must still conduct the remaining home improvement after T years. By the design of the policy, homeowners have up to five years to make the relevant improvements, i.e. $T \leq 5$ years. However, they may register a further exemption after 5 years which implies that T may be larger than 5 years in some cases. Hence, their overall willingness to pay for the property should decrease by £3,500 plus the expected remaining costs of improvements in T years discounted at a rate of $\rho \geq 0$. Alternatively, this buyer can choose not to invest in the property to increase its energy efficiency. In this case, they might be caught with a probability of $0 \leq \pi \leq 1$ and will need to pay the fine of up to £5,000.[13] In this case, the buyer will need to incur the cost for the home improvement anyways. Hence, their willingness to pay should decrease by the probability of getting fined times the fine plus the costs of improvement. The change in the willingness to pay δ^* of a risk-neutral buyer will be determined by the minimum of the expected costs of improvement or the expected costs of not improving with the risk of getting fined, summarized according to

$$\delta^* = -\min \begin{cases} 3500 \ + \ (\frac{1}{1+\rho})^T (\ E[Costs] \ - \ 3500 \ ) \\ \pi(\ E[Fine] \ + \ E[Costs] \ ) \end{cases} \tag{1}$$

[11] UK government provides further details including case studies: `https://www.gov.uk/guidance/changes-to-tax-relief-for-residential-landlords-how-its-worked-out-including-case-studies`. Thereafter, a new BTL tax system was phased in with the target to be fully in place from 6 April 2020. The percentage of mortgage interest payment deductible from rental income was decreased by 25 percentage points per year reaching zero percent in the tax year 2020/21. From April 2020, the entire sum of the interest payment would then be subject to a relatively low 20% tax relief. As a result, high-income landlords in higher tax brackets could then end up paying much more tax than before. Second, there were changes to stamp duty surcharges for BTL. Since 1 April 2016, buy-to-let or second properties, i.e. any residential property not classed as a main residence, carry an additional 3% stamp duty surcharge if they are priced higher at £40,000 or more.

[12] `https://www.gov.uk/stamp-duty-land-tax/residential-property-rates`.

[13] A public release of information by the UK government suggests that, as of March 2020, 449 compliance notices and 17 financial penalties have been issued for breaches of MEES 2018 with fines totalling £65,600. This implies that the average fine was about £3,859.

How does this affect this change in a buyer's willingness to pay affect transaction prices, $\Delta Price$? For simplicity, we assume the supply properties to be perfectly inelastic, i.e. it does not change with price. This may be a reasonable assumption for a relatively short time period which we study as houses or flats cannot be easily built. In this case, the decrease in a buyer's willingness to pay should be directly reflected in transaction prices. I.e. the decrease in the equilibrium price, or transaction price, should be equal to the decrease of a buyer's willingness to pay, i.e. $\Delta Price = \delta^*$.

Examining Eq. 1, describing a buyer's willingness to pay δ^*, we can see that the conceptual framework consists of several elements, some of which depend on the size of the property and others that do not. On the one hand, there is a set of elements which consists of the magnitude of the fine, $E[Fine]$, probability of getting caught, π, and the cost cap of £3,500. All of these elements do not vary by property size or property value. On the other hand, there is also the expected cost of home improvement, $E[Costs]$, which may be a function of the size of the property. While the first set of elements justifies measuring price changes in absolute terms (in £), the latter can justify a measurement in relative terms, for example, the percent (%) change relative to the price in the pre-intervention period.[14]

This simple conceptual framework is clearly not exhaustive. First, potential home buyers take their decision under uncertainty. Their willingness to pay δ^* would decrease further if they were risk-averse rather than risk-neutral. Second, potential buyers might perceive this policy change as a signal that goes beyond its current scope. For example, they might believe that there is a risk of additional policy tightening in the future, for example via higher fines or lower exemptions. We would expect their willingness to pay δ^* to decrease further for this reason.

While this framework discusses the implications for the private-rented property market, it may also be relevant to the owner-occupied market for several reasons. First, buyers of residential properties value the option of being able to ultimately let the property. Second, the implementation of the MEES 2018 may be seen as a signal for policy interventions in the residential market which might be implemented in the future. Thus, even if the intervention does not directly apply to buyers of residential properties, their willingness to buy might be affected in a similar way.

4 Data

4.1 Data sources

In order to conduct an empirical assessment of the MEES 2018 policy, we combine data from three different sources. These data cover information on transaction of properties, the energy performance of these properties as well as demographic characteristics of the regions in which they are located. We summarise these data sets below. We provide additional details in Appendix A.

Property transactions: First, we derive a data set which includes information on property transactions,

[14] As the property price can be seen as a function of its size.

sourced from HM Land Registry (HM Land Registry, 2014). This data set includes information on all residential property transactions in England and Wales since 1995. In particular, it includes information on the date of the transaction and the price paid as well as the exact address of each property. Our empirical strategy requires data on transactions from both before and after the MEES 2018 policy intervention. For that reason, we focus only on those properties which have multiple transactions, both before and after the intervention date.[15] Specifically, our final sample consists of such properties which were repeatedly transacted in the time period between 2015 and 2019. I.e. our data covers only a short period after the implementation of the MEES 2018 in April 2018.[16]

Energy Performance Certificates (EPC): To identify the properties which are affected by the MEES 2018 policy intervention, we use data from the public register on Energy Performance Certificates (EPC) which is the official source for all EPCs issued for all domestic buildings and building units in England and Wales. This register covers information on all properties that have been constructed, sold or let since 2008, and were sourced from the Ministry of Housing, Communities & Local Government (MHCLG) (Ministry of Housing, Communities & Local Government, 2020a). EPCs classify properties into seven categories which range from A (being the most energy efficient) to G (being the least energy efficient).[17] In addition to these broad categories, EPCs provide a continuous measure of energy efficiency underlying the discrete EPC ratings ranging from 0 to 100. EPCs ratings can be obtained by an accredited energy assessor who visits the property to collect the necessary information. This process of obtaining an EPC costs between GBP 60 and GBP 120 and EPCs are valid on a property for 10 years. Whenever properties are built, sold, or rented, the owners need to ensure that their properties have EPCs.[18] EPCs are a good proxy of energy efficiency but can come with some measurement error (Hardy and Glew, 2019; Crawley et al., 2019).[19]

Geodemographic characteristics We complement these data with information characterising the area surrounding each property, making use of geodemographic classifications in England and Wales. Specifically, we use the classification produced from the 2011 census which clusters communities into eight different types (Gale et al., 2016): 1) Rural Residents; 2) Cosmopolitans; 3) Ethnicity Central; 4) Multicultural Metropolitans; 5) Urbanites; 6) Suburbanites; 7) Constrained City Dwellers; 8) Hard-Pressed Living.[20] These data are at the output area (OA) level, with each OA consisting of one type of postcode units, either entirely urban or entirely rural postcodes. OAs are designed to share similar population sizes, and attain a high level of social homogeneity, with respect to the type of tenure and

[15]Intervention date is taken to be the 1 April 2018.

[16]We acknowledge that it may lead to an underestimation of the effects of the MEES 2018 if they unfold over a longer period of time.

[17]In our empirical analyses, we make use of these discrete ECP ratings. The continuous measure could be used in a RDD design. However, there is bunching of observations at the cutoffs for each of the categories (e.g. from F to E).

[18]`https://www.gov.uk/buy-sell-your-home/energy-performance-certificates`.

[19]In our analyses, we implicitly assume this measurement error to be random and non-systematic.

[20]For simplicity we only use the broadest classification into eight clusters, the so-called upper *Supergroup* tier. In addition to these 8 *Supergroups*, the classifications provides a set of more granular hierarchical clusters consisting of 26 *Groups*, and 76 *Subgroups*. We provide a description of each of these 8 *Supergroups*, or clusters, in Appendix A.

accommodation, with a recommended size of 125 households.[21] We sourced these data from the Open Geography Portal (Office for National Statistics, 2011).

While the MEES 2018 policy intervention applies specifically to domestic private rented properties, the analysis in this paper examines all repeated property sales in England and Wales, that are sold for value and are lodged with HM Land Registry for registration. The rationale behind this decision is that a part of the impact of the MEES 2018 policy on the house prices should spill over to, from the rental market to the owner-occupied market, since all properties might be ultimately let out (i.e. a homeowner of an owner-occupied house might want to keep the option of renting the property out in the future). Hence, it is assumed that all properties with EPC rating F and G, regardless of their type of tenure, are eligible to be affected in a similar way by the MEES 2018 policy.

4.2 Creating the repeated sales data set

The original data set of property transactions lists 5,110,623 transactions between 2015 to 2019. In a first step, we construct the sub-sample of repeated sales, i.e. properties that were transacted both before and after the MEES 2018 policy intervention in April 2018. We then create a panel version of these data by identifying which transactions belong to the same property by using address information.[22] Creating this panel dimension allows us to examine how the price of the same property changes over time. This sub-sample includes 454,085 transactions. In a second step, we then match the energy performance certificates to each of these transactions using the common address. Specifically, we match each transaction with the closest prior registered EPC rating. If there was no prior EPC rating, the transaction we match it with is the EPC rating registered on the closest date afterwards.[23] This choice enables us to study the full population of all relevant properties with EPC information, whilst also ensuring that the energy and property information is accurate and reflective of a dwelling at time of sale. In the final step, we link each property in our matched data set with the corresponding geographic information using postcode information.

We also derive several additional variables from our matched data. First, we construct an indicator variable of whether a property is affected by the MEES 2018. We classify a property as being affected if it has an EPC rating of below E in its latest assessment. It takes on the value of one if that is the case and zero otherwise. Second, we create a variable which represents the high-level geographic region associated with the property. This variable is useful when examining regional effects. In this case, we adopt the *Nomenclature of Territorial Units for Statistics (NUTS)* Level 1 regions for the UK, and match the local-authority registered for each transaction with the corresponding region. The NUTS system is used for referencing the current administrative and electoral areas of the UK for statistical purposes. The NUTS Levels 1, 2, and 3 all stay fixed for a minimum of three years, with Level 1 referencing the regions, Level 2 the counties and grouped

[21] The total number of 2011 OAs is 171,372 for England and 10,036 for Wales. `https://webarchive.nationalarchives.gov.uk/20160107193025/http://www.ons.gov.uk/ons/guide-method/geography/beginner-s-guide/census/output-area--oas-/index.html`.

[22] Details on data processing can be found in Appendix A.

[23] This reflects the possibility that EPC might be registered only afterwards but could be available to prospective buyers at the time of the purchase.

Figure 1: Distribution of property prices

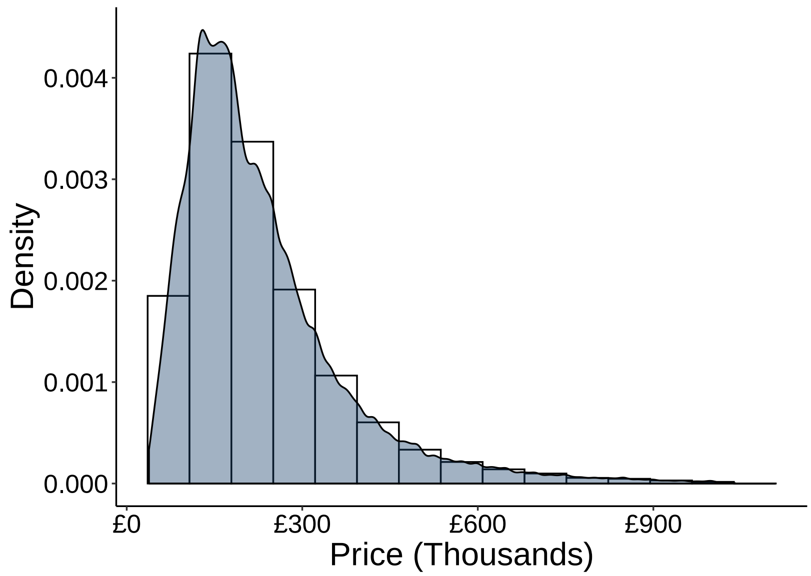

Note: This figure shows the empirical distribution of the price variable in our final data set, trimmed at the 1st and 99th percentiles. It shows a histogram of the price as well as a density plot of the price variable. Source: House price data are obtained from HM Land Registry.

London boroughs, and Level 3 the unitary authorities and districts in the UK.

A simple analysis of the transaction data set shows some clear outliers. For instance, the minimum price is £1 and the maximum price is £247 million. To account for the concern that our results are driven by such outliers, we delete those observations that are below the 1st and above the 99th percentiles of this price distribution.[24] In the trimmed data set, there are 392,495 transaction. Their minimum price is £38,075 and their maximum is £1.04 million. We then remove all of those transactions with missing values in relevant variables, assuming the missing values are random. After ensuring that each property has at least one sale after the MEES 2018 policy intervention and one sale prior to this date, we are left with 305,337 transactions across 147,842 properties. In Figure 1, we show the empirical distribution of the property prices in our final data set. The distribution is right-skewed, i.e. a large share of properties are sold at prices between £150k and £250k with very few properties being sold at prices above £300k. The mean house price is £231k, with a standard deviation of £141k.

4.3 Descriptive analyses

We also provide other descriptive statistics of our panel data set. Figure 2 presents the geographical distribution of properties in our sample. Panel a) shows the average transaction prices per region. As expected, the highest average transaction prices are in London and in regions in the South-East of England. Panel b) illustrates the number of properties per region. There is a high concentration of properties in the South-East, with relatively few in Wales and the North-West.

[24] In robustness tests, we also trim the data set at other percentiles. Section 5.1.4 discusses the sensitivity of our results depending on the trimming.

Figure 2: Prices and transactions by regions

(a) Average price by region (b) Number of properties by region

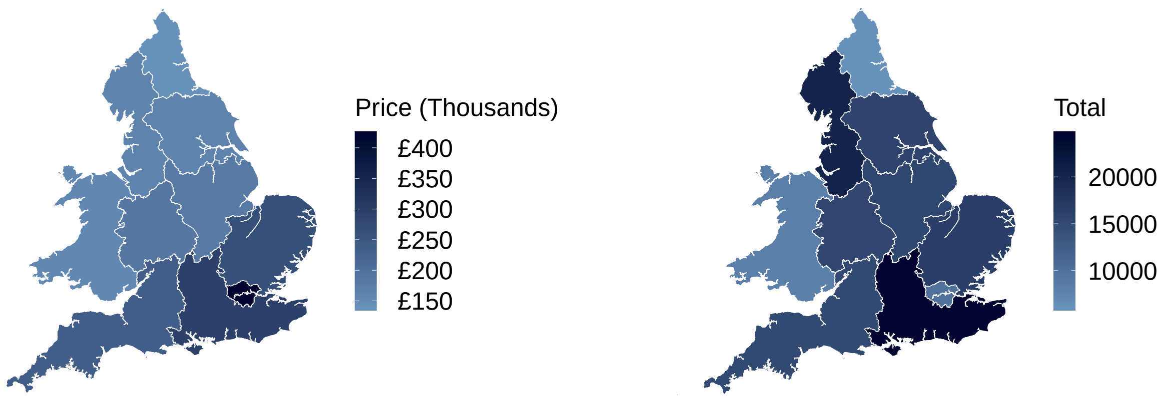

Note: This figure shows maps of geographic regions in England and Wales: Panel a) shows the average price in each region in our final sample of property transactions. Panel b) shows the number of properties in our final sample of property transactions. Source: House price data are obtained from HM Land Registry.

Panel a) of Figure 3 shows the relative frequency of properties by EPC ratings before and after the MEES 2018 policy intervention. It illustrates several interesting stylised facts: first, the majority of properties are of medium energy efficiency with EPC ratings of C, D or E. Second, the share of properties with energy efficiency ratings of F and G is much lower than the share of properties with higher ratings. In fact, they make up only approximately 7% of the total sample size in the time period before the MEES 2018 policy intervention. In total, we observe 10,032 properties with such EPC ratings. This number is much lower than the 137,810 energy-efficient properties with EPC ratings of E or higher, which are not affected by the MEES 2018 policy intervention. Third, it appears that the share of properties with high EPC ratings increases over time, suggesting that properties become more energy-efficient after the MEES 2018 policy. We take this as evidence that a considerable number of properties appear to have upgraded their energy efficiency rating. From these descriptive analyses, it is unclear whether this is due to the MEES 2018 policy intervention, or other reasons for improving efficiency.

In Panel b) of Figure 3, we illustrate average transaction prices by EPC ratings over time, i.e. before and after the MEES 2018 policy intervention. We observe that the prices appear to rise across all EPC ratings. However, they seem to increase more for properties which are initially less energy-efficient. The descriptive evidence of these two panels of Figure 3 seems rather counter-intuitive. Over time, properties have become more energy-efficient, possibly suggesting higher demand for higher efficiency homes. By contrast, the prices in the least energy efficient class appear to rise the most. However, these plots present a highly aggregated view. Indeed, there can be many confounding variables which impact price, besides the EPC rating.

Figure 3: Number of properties and price by EPC over time

(a) Share of properties by EPC rating

(b) Average price by EPC rating

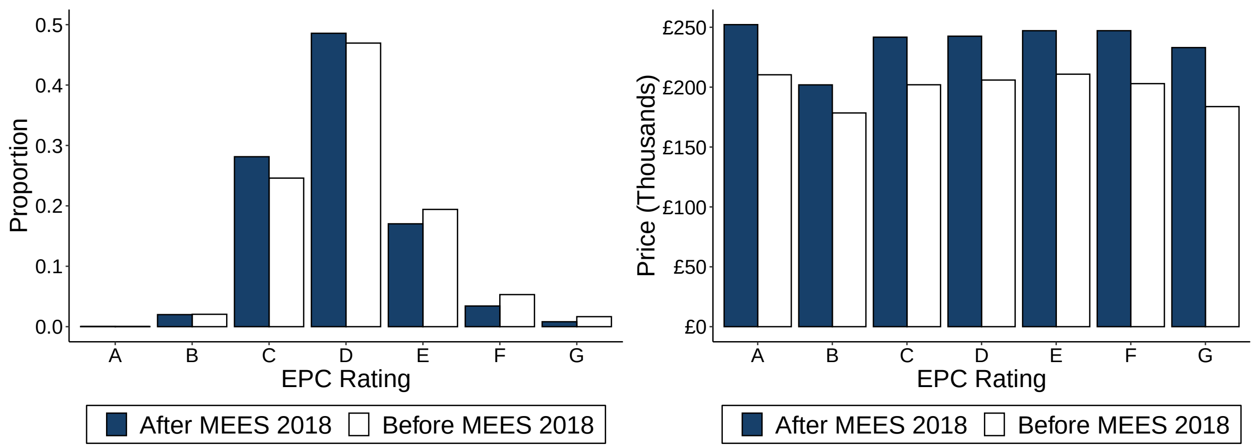

Note: Panel a) shows the relative frequency of EPC ratings before and after the MEES 2018 policy intervention. Panel b) shows barplots of the average transaction price by EPC ratings before and after the MEES 2018 policy intervention. Source: House price data are from HM Land Registry. EPC ratings are from the MHCLG.

4.4 Comparability of the treatment group and the control group

In this section, we examine the characteristics of properties in our sample. Table 1 compares characteristics depending on the MEES 2018 policy intervention. On the one hand, there are those properties with EPC ratings below E. They are directly targeted by the MEES 2018 policy intervention which applies to them. Hence, these properties form our *Treatment group*. On the other hand, there are properties with EPC ratings of at least E. They are not affected by the MEES 2018 policy intervention, our *Control group*. Just from examining these simple descriptive statistics, we can observe some sizeable differences. For instance, there are differences in the property type. As compared with more energy-efficient properties, those with EPC ratings below E are more often houses (77.4% vs 73.2%). Moreover, there is a higher share of older properties among them. Low EPC properties are more often built before 1900 than higher EPC ones (30.7% vs 11.3%), and in 1900-1966 (57.7% vs 43.2%). Similarly, the rural-residents demographic is more predominant in the low-energy efficiency properties (23.3% vs 9.3%).

To quantify the differences between these two groups, we calculate the *standardized mean difference (SMD)* for each covariate (see Appendix B for details). The SMD value 0.1 is commonly used as a threshold for indicating imbalance (Stuart et al., 2013). Values above this threshold suggest that the treated and non-treated units are not comparable with respect to the corresponding variable. We report the calculated SMD values in column 2 of Table 1. As we can see, the two groups appear to be imbalanced in almost all covariates, with SMD values being greater than the critical threshold of 0.1. The most striking differences are in terms of the CONSTRUCTION AGE, with a SMD value of 0.924, and the DEMOGRAPHIC characteristic reporting a SMD value of 0.397. The only exception is property TENURE whose SMD is relatively small (0.071).

This imbalance between our treatment and control groups raises the concern that comparing transaction prices across these two groups might suffer from a selection bias. For this reason, we seek to ensure compara-

Table 1: Characteristics of properties by EPC rating

Group by EPC rating	Treatment group	Control group	Control group
Sample	Full sample	Full sample	Matched sample
	(1)	(2)	(3)
PROPERTY TYPE		(SMD=0.203)	(SMD=0.017)
Bungalow	11.9%	9.4%	11.5%
Flat	9.3%	15.5%	9.1%
House	77.4%	73.2%	78.1%
Maisonette	1.4%	1.8%	1.3%
Park home	0.0%	0.0%	0.0%
CONSTRUCTION AGE		(SMD=0.924)	(SMD=0.008)
before 1900	30.7%	11.3%	31.0%
1900-1929	27.2%	15.6%	27.1%
1930-1966	30.5%	27.6%	30.3%
1967-1995	11.2%	30.1%	11.1%
1996-2006	0.4%	12.6%	0.4%
2007 onwards	0.0%	2.9%	0.0%
TENURE		(SMD=0.071)	(SMD=0.026)
owner-occupied	88.9%	86.6%	89.2%
rental (private)	10.4%	12.5%	9.9%
rental (social)	0.7%	0.9%	0.9%
DEMOGRAPHIC		(SMD=0.397)	(SMD=0.020)
Constrained city dwellers	5.0%	5.5%	5.0%
Cosmopolitans	4.1%	4.8%	3.9%
Ethnicity central	2.0%	2.7%	1.9%
Hard-pressed living	15.5%	17.9%	15.2%
Multicultural metropolitans	9.8%	9.9%	9.8%
Rural residents	23.3%	9.3%	23.9%
Suburbanites	16.1%	22.2%	15.9%
Urbanites	24.2%	27.7%	24.3%
REGION		(SMD=0.181)	(SMD=0.044)
East Midlands (England)	12.2%	10.4%	11.0%
East of England	10.6%	11.2%	11.0%
London	5.9%	6.8%	5.7%
North East (England)	2.8%	4.3%	2.7%
North West (England)	13.5%	13.8%	13.4%
South East (England)	13.2%	16.7%	13.3%
South West (England)	12.7%	10.0%	13.4%
Wales	7.6%	5.5%	8.0%
West Midlands (England)	11.0%	10.4%	10.8%
Yorkshire and The Humber	10.4%	10.7%	10.8%
TOTAL FLOOR AREA		(SMD=0.160)	(SMD=0.016)
Floor area	$90.69m^2$	$84.47m^2$	$91.38m^2$
HABITABLE ROOMS		(SMD=0.155)	(SMD=0.007)
less than 4	22.3%	28.8%	22.4%
4 to 5	58.2%	54.5%	57.8%
more than 5	19.5%	16.7%	19.7%
Number of observations	(N=10,032)	(N=137,810)	(N=10,032)

Note: This table examines the balance of each covariate: Column 1 shows covariates of properties with an EPC rating below E, our *Treatment group*. Column 2 shows covariates of properties with an EPC rating of at least E, our *Control group*, in the full sample. Column 3 shows covariates of properties with an EPC rating of at least E, our *Control group*, in our matched sample after PSM matching. For each categorical covariate, percentages of each level are reported. For the continuous covariate TOTAL FLOOR AREA, the mean is displayed. Source: House price data are from HM Land Registry. EPC ratings are from the MHCLG. Demographic data are from the ONS.

bility between properties in these two groups by employing *Propensity Score Matching* (PSM) (Rosenbaum and Rubin, 1983). In our set-up, PSM aims to match each property that is assigned to the treatment group with one property assigned to the control group. This ensures that the properties we examine in our empirical analyses are similar on observable characteristics, such as property size, region or property age. PSM can be used to reduce selection bias that could influence the estimation of the effects of the MEES 2018. Specifically, we implement PSM on the properties before the MEES 2018 policy intervention. If a property was sold multiple times before the MEES 2018 policy, we match them based on the transaction whose date was closest to 1 April 2018, the day of the MEES 2018 policy intervention.

We conduct the *Propensity Score Matching* (PSM) in two steps (see Appendix C for details). In a first step, we estimate the propensity of being affected by the MEES 2018 policy intervention, conditional on a set of observed covariates (Rosenbaum and Rubin, 1983; Rubin, 2007). The "propensity score" is the predicted probability of the treatments assignment. In a second step, we then use these propensity scores to find for each "treated" property the most similar property unaffected by the MEES 2018. For properties with similar propensity scores, the idea of PSM is that the treatment assignment for these two properties is independent of all confounding variables (Westreich et al., 2010), i.e. the selection bias decreases.

To estimate propensity scores, the existing literature typically uses one of two parametric models, either Logistic regression or Probit regression models (Caliendo and Kopeinig, 2005). More recently, there has been an increasing use of machine-learning methods to construct scores (Westreich et al., 2010; Cannas and Arpino, 2019; Luellen et al., 2005; Setoguchi et al., 2008; McCaffrey et al., 2004). For example, Lee et al. (2010) show that machine-learning methods can improve the covariate balance between treated and non-treated units in the resulting matched data set using a simulation study. For that reason, we complement a standard Logistic regression with two popular tree based machine-learning classifiers. Since single tree classifiers have the unfortunate reputation of being rather weak prediction models (James et al., 2013, p. 316; Efron and Hastie, 2016, p. 324), we employ the following two ensemble methods that usually enjoy good predictive performance:

Random Forest is a method that averages across many trees based on a bootstrap resampled data set. Further to simply bagging the trees, the method attempts to decorrelate the trees by further resampling of the covariate set (over which each split is estimated) (Breiman, 2001). For each tree a random sample of $m < p$ covariates are chosen for inclusion. A typical value of this hyperparameter is $m = \sqrt{p}$, however, this can also be assessed via cross-validation. Using random forest with $m = p$ is equivalent to a bagged tree model (Efron and Hastie, 2016, p. 327).

Boosted Trees work by repeatedly extending a basic seed model (Friedman et al., 2000). More specifically, for each iteration of the algorithm, a shallow classification tree is added to the previous model (estimated based on the residuals), and hence we build up an additive model consisting in a sum of trees. Overall, boosting has three hyperparameters that can alter the predictive performance of a boosted model: the number of trees B, the number of splits in each tree d, and the shrinkage parameter (which scales each tree to be added) λ (Efron and Hastie, 2016, p. 333-334).

When fitting the above models, we split the data into training and test sets, respectively 70 and 30 percent of the data. We also ensured the class balance between treatment and control groups was consistent across these sets. Logistic regression models require no tuning parameters and we use a full model with all covariates. However, for random forest and boosting, we implement 5-fold cross-validation (CV) (Hastie et al., 2009, p. 241-249) on the training data.[25]

For each model, the hyperparameters that maximized the cross-validated *area under the curve (AUC)* score were selected. One way to interpret the AUC score, is to think that if we randomly choose a positive instance and a negative instance, then AUC represents the probability that the classifier ranks the positive instance higher than the negative instance. For the random forest, the hyperparameter set of values that maximized the CV AUC score was equal to $m = 2$, while for XGBoost was $B = 500$ trees, $d = 6$ and $\lambda = 0.01$. The final performance of each method in terms of AUC is given in Appendix C. Interestingly, in this application all methods perform similarly by this metric, with a slight edge given to the boosting method. In addition to the slightly higher AUC, it becomes apparent that the boosted method provides superior alignment of covariate balance as measured through SMD. For this reason, we choose the boosted model to continue our analysis confident that the matched sample is far more comparable than the unmatched one.

Indeed, PSM matching improves the covariate balancing across treatment and control groups. In column 3 of Table 1 we present covariates and SMD values of each coviariates of our matched data set, using the noosted model. We can see that SMD values are now lower than 0.1 for all variables. Some key examples of this can be observed by comparing column 1 and column 3. For instance, after matching the mean total floor area for treated properties was only 0.96m^2 smaller than for non-treated properties (as opposed to 6.2m^2 in column 2), while the percent difference between treated and non-treated properties built before 1900, and in the periods 1967-1995 and 1996-2006, was reduced to -0.3% from 19.4%, 0.1% from -18.9%, and 0% from -12.2%, respectively. We conclude that PSM has reduced any difference that had existed between the treatment group and the control group prior to matching.

We also examine how the price distributions of properties in our treatment and control groups change following the implementation of PSM. In Figure 4, we plot density estimates of transaction prices both before ("pre") and after ("post") the MEES 2018 policy intervention. Panel a) shows the distributions of prices in our unmatched data set before the MEES 2018 policy intervention. First, it shows prices of properties with EPC ratings of at least E, our control group. It also illustrates the distribution of prices of those properties with EPC ratings of below E, our treatment group, indicated by the shaded area. These prices are on average lower than prices of energy-efficient properties, indicated by the white area under the curve. Panel b) then illustrates the distribution of prices in our matched sample where we keep only those energy-efficient properties that are similar to our energy-inefficient properties. In this matched sample, the distribution of prices of energy-inefficient properties shifts to the left. I.e. these properties in our control group not only

[25] The particular implementation of boosting used was via the *eXtreme Gradient Boosting (XGBoost)* package (Chen and Guestrin, 2016). The primary benefit of this package is that it enables the usage of parallel computing to boost the speed over the classical boosting algorithm, whilst also including some options for additional model regularisation.

Figure 4: Distribution of transaction prices of properties

(a) Pre-intervention density for unmatched data set

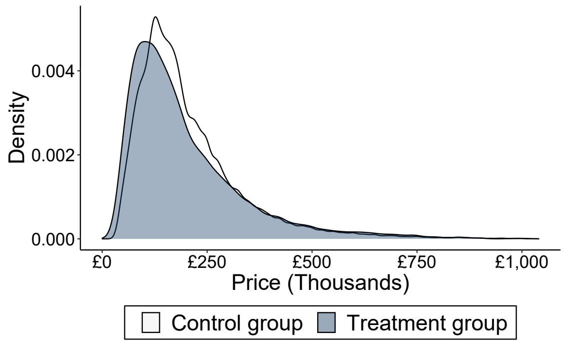

(b) Pre-intervention density for PSM data set

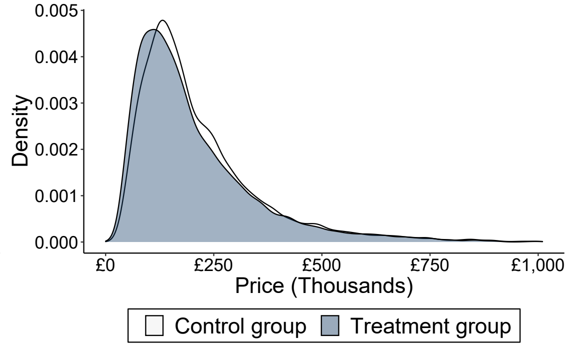

(c) Post-intervention density for unmatched data set

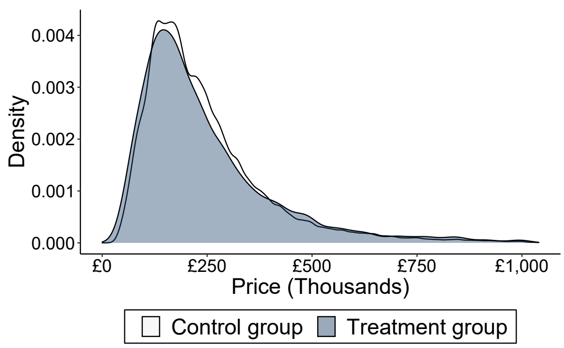

(d) Post-intervention density for PSM data set

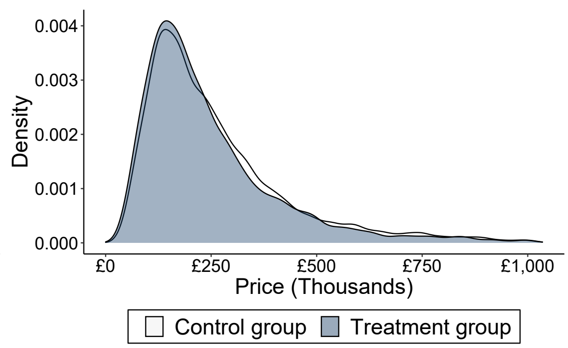

Note: This figure shows density estimates of transaction prices by energy efficiency. Our control group consists of properties with EPC ratings of at least E. Our treatment group consists of properties with EPC ratings below E. Plots in the top row show density estimates before ("pre") the MEES 2018 policy intervention and those in the bottom row show density estimates after ("post") the MEES 2018 policy intervention. Plots on the left show densities of the unmatched data set before implementing PSM and those on the right show the densities of the matched data set after PSM matching. Source: House price data are from HM Land Registry. EPC ratings are from the MHCLG.

become more similar in terms of observable characteristics, but also their price distribution becomes more similar to the distribution of prices of energy-inefficient properties. In the bottom part of this figure, we then illustrate the distributions of prices after MEES 2018 policy interventions. Comparing Panel c) with Panel d), we observe, also for this time period, that properties become more similar in terms of their price distribution following the implementation of PSM.

In summary, PSM improves the comparability of our treatment and control groups along several dimensions. While this adds credibility to the internal validity of our analyses and the causal interpretation of our DiD estimates, it remains open to what extent they extrapolate to the entire population, i.e. are externally valid. Indeed, in section 5.3, we show that our PSM-matched sample differs from the population in a number of dimensions.

5 Main Analyses

In this section, we present our main analyses of the effect of the MEES 2018 policy intervention on transaction prices. We employ a Difference-in-Difference design which has been widely used in the applied economics literature to examine treatment effects of policy interventions (e.g. Card and Krueger, 1994; Brown et al., 2016).[26] We also examine the relevance of our results by discussing the characteristics of our estimation sample and by examining the validity of the parallel trend assumption underlying this research design.

5.1 Difference-in-Difference analyses

In our set-up, the Difference-in-Difference (DiD) design quantifies the effect of the MEES 2018 policy intervention by calculating the average change in transaction prices among properties in the treatment group which consists of those properties affected by the MEES 2018, those with EPC ratings of at least E. It then compares it to the average change in transaction prices among properties in the control group, i.e. properties unaffected by the MEES 2018 policy intervention, those with EPC ratings of F or G. In our estimation, we use our matched sample consisting of properties that are very similar on observable characteristics. Specifically, we examine the subset of $n = 20,064$ properties with at least two transactions, one before and one after the policy intervention on 1 April 2018.[27]

5.1.1 The average effect of the MEES 2018 on house prices

In this sample, we observe each property over time whenever it is transacted. This panel dimension enables us to calculate the change in price of each individual property, ΔPrice_i. It allows us to estimate the regression model in first differences where ΔPrice_i is our dependent variable. Estimating the regression in first differences enables us to estimate the effect of the MEES 2018 *within* a property over time, thus removing the effects of unobserved time-invariant confounding variables *between* the properties. Such confounders may include the size and the architectural style of the building and features of the location of the property such as its proximity to the sea. The first differences also absorb the effects of micro location which is not time-varying.[28] We also control for the broader region in which each property is located acknowledging that house price trends differ across the country. Specifically, we estimate the following regression model

$$\Delta\text{Price}_i = \delta \ \mathbb{1}(EPC_i < E) + \sum_{r=1}^{R} \tau_r \mathbb{1}(\text{region}_i = r) + \Delta u_i \ , \tag{2}$$

where the subscript $i = 1, \ldots, n$ indexes the properties. The model components are as follows:

- ΔPrice_i is the change in price of property i between two consecutive transactions.

[26] For example, Gropp et al. (2019); Arnould et al. (2020) studying the effects of prudential policy and Rodnyansky and Darmouni (2017); Fatouh et al. (2021) examining the effects of monetary policy.

[27] To create a balanced panel, we only keep the post-intervention observation with the most recent transaction date.

[28] We decide to estimate our DiD model in first difference as opposed to levels as it lowers the number of parameters an increases the degrees of freedom. We report the estimated R^2 which measures the proportion of the variation in the change in price variable within properties, that is explained by the respective model.

- $\mathbb{1}(EPC_i < E) = 1$, if property i has an EPC rating below E in post-intervention period, and zero otherwise;
- $\mathbb{1}(\text{region}_i = r)$ if property i is located in region r. There are $R = 10$ regions in our sample;
- Δu_i are idiosyncratic errors;

We use ordinary least-squares to estimate the model given by Equation 2. Of primary interest is the estimated coefficient of $\mathbb{1}(EPC_i < E)$ which measures the change of the transaction price of properties affected by the MEES 2018 policy relative to unaffected properties, those with EPC ratings of at least E. Our Difference-in-Difference design tests the combined effect of properties with exemptions from the policy and those without as we do not observe exemptions at the property level.

Table 2 summarises these estimates. It shows the estimated coefficients using the full sample of all properties in our matched data set. The estimated coefficient of $\mathbb{1}(EPC_i < E)$ points to a decrease in the prices of affected properties with EPC ratings below E by £8,910 relative to unaffected properties with an EPC ratings of at least E. This estimate is statistically significant at the 1% level of significance. Overall, the MEES 2018 policy intervention appears to have significantly impacted the house prices of affected properties.

Table 2: The average effect of the MEES 2018 on house prices

Dependent variable	$\Delta Price$
	(1)
$\mathbb{1}(EPC < E)$	$-8,910$
	(1,053)
	p < 0.001
Region FE	YES
Observations	20,064
R^2	0.349

Note: This table shows the results on the matched sample. Column 1 shows the regression model in eq. (2). Standard errors are reported in brackets. All specifications are estimated in first differences. Source: House price data are from HM Land Registry. EPC ratings are from the MHCLG.

5.1.2 Heterogeneity: Post-2019 Policy Amendment

On 1 April 2019, UK government introduced an amendment to the original MEES 2018 policy. Following this date, the original "no cost to the landlord" provision became unavailable. As a result, landlords had to use their own funding to cover the cost of improving their property to an EPC rating of at least E, still subject to a spending cap of £3,500 on each property. As this amendment should increase the expected costs of improving the EPC rating, we expect the MEES 2018 to have a more negative impact on property prices. We account for this policy change in a revised empirical specification in which we introduce a dummy variable corresponding to the policy amendment date to our baseline model given in Equation 2. We also introduce an interaction term with our treatment indicator which models the additional treatment effect of this policy amendment. The revised model is given by

$$\Delta\text{Price}_i = \delta\ \mathbb{1}(EPC_i < E) + \beta\ \mathbb{1}(EPC_i < E) \times \mathbb{1}(Post\ 2019_t = 1) + \gamma\ \mathbb{1}(Post\ 2019_t = 1) \qquad (3)$$
$$+ \sum_{r=1}^{R} \tau_r \mathbb{1}(\text{region}_i = r) + \Delta u_i\ ,$$

where $\mathbb{1}(Post\ 2019_t = 1)$ is a dummy variable taking on the value of one if the second transaction of property i happens in April 2019 or later (zero otherwise). It estimates the marginal price change since April 2019. Its interaction effect with $\mathbb{1}(EPC_i < E)$ provides insight onto the heterogeneous effect of the amendment on our treatment group of properties under the scope of the policy.

Table 3 shows these regression results. In column 1, we replicate the baseline regression results given by Equation 2 for comparison. In column 2, we then show the regression results for the revised model in Equation 3. The main effect increases in magnitude slightly from $£-8,910$ to about $£-10,918$, still being statistically significant at all conventional levels. The estimated coefficient of $\mathbb{1}(EPC < E) \times \mathbb{1}(Post\ 2019)$ is about $£4k$ being statistically significant at 5% significance level. Hence, there is evidence that the post-2019 policy amendment reduced the magnitude of the drop in prices of energy-inefficient properties under the scope of the policy compared to energy-efficient ones.

Table 3: The DiD effect with post-2019 policy amendment

Dependent variable	$\Delta Price$	$\Delta Price$
	(1)	(2)
$\mathbb{1}(EPC < E)$	$-8,910$	$-10,918$
	(1,053)	(1,391)
	$p < 0.001$	$p < 0.001$
$\mathbb{1}(EPC < E) \times \mathbb{1}(Post\ 2019)$		4,383
		(2,133)
		$p = 0.04$
$\mathbb{1}(Post\ 2019)$		$-3,311$
		(1,157)
		$p = 0.004$
Region FE	YES	YES
Observations	20,064	20,064
R^2	0.349	0.349

Note: This table shows the results on the matched sample. Column 1 shows the regression model in eq. (2). Column 2 shows the regression model in eq. (3). Standard errors are reported in brackets. All specifications are estimated in first differences. Source: House price data are from HM Land Registry. EPC ratings are from the MHCLG.

5.1.3 Heterogeneity: Testing for Changes in Tenure

Properties can be owner-occupied or private-rented. Moreover, the tenure of the dwelling can change over time. As a result, a property classified as owner-occupied can become a private rental property (and vice versa). Our data includes information on the tenure of each property.[29] This allows us to model such changes in tenure status. To that end, we introduce dummy variables reflecting the tenure status of a property, alongside their interactions with the treatment variable. The revised model is given by

$$\begin{aligned}\Delta\text{Price}_i = {} & \delta\mathbb{1}(EPC_i < E) \\ & + \beta_{12}\mathbb{1}(EPC_i < E) \times \mathbb{1}(i = \text{Owner To Rented}) \\ & + \beta_{22}\mathbb{1}(EPC_i < E) \times \mathbb{1}(i = \text{Rented To Owner}) \\ & + \beta_{32}\mathbb{1}(EPC_i < E) \times \mathbb{1}(i = \text{Remained Rented}) \\ & + \beta_{11}\mathbb{1}(i = \text{Owner To Rented}) + \beta_{21}\mathbb{1}(i = \text{Rented To Owner}) + \beta_{31}\mathbb{1}(i = \text{Remained Rented}) \\ & + \textstyle\sum_{r=1}^{R} \tau_r \mathbb{1}(\text{region}_i = r) + \Delta u_i \ .\end{aligned} \tag{4}$$

where $\mathbb{1}(i = \text{Owner To Rented})$ is a dummy variable taking on the value of one if the property was owner-occupied in the first transaction and becomes rented in the second transaction. $\mathbb{1}(i = \text{Rented To Owner})$ is a dummy variable taking on the value of one if the property was rented in the first transaction and becomes owner-occupied in the second transaction. $\mathbb{1}(i = \text{Remained Rented})$ is a dummy variable taking on the value of one if the property was rented in the first transaction and remains rented in the second transaction. Their interaction effects with $\mathbb{1}(EPC_i < E)$ provide insight onto the impact of the MEES 2018.

Table 4 shows the estimated coefficients. As before, in Column 1, we replicate the main regression shown in Equation 2 for comparison. In Column 2, we then show the regression results of the augmented model by Equation 4. Our DiD effect, the coefficient of $\mathbb{1}(EPC < E)$, is very similar, becoming slightly more negative in size ($-£9,285$). It remains statistically significant at all conventional levels. We can interpret this as the effect on properties that remain owner-occupied, which is the omitted category in Equation 4.

The effect of the MEES 2018 on properties that change their tenure status from owner-occupied to private-rented is of similar magnitude.[30] Interestingly, the coefficient of $\mathbb{1}(\text{Remained Rented})$ suggests that prices of properties that remained private-rented decreased on average by about $£21k$ more than the average price of properties that remained owner-occupied.[31] By contrast, the marginal effect of the MEES 2018 on the energy-inefficient subset of private-rented properties is close to zero.[32] This suggests that the MEES 2018 did not decrease this price drop by even more. Instead, the treatment effect of about $£9k$ seems to be driven by the remaining subsets of properties, those stay owner-occupied or have changed their tenure status.

There are possible explanations for this heterogeneous effect across tenure status. On the one hand, buyers of private-rented properties might believe that there is a risk of additional policy tightening in the future beyond EPC rating E. As a result, we see a price drop across all EPC ratings. As a result, we see their willingness for low EPC properties to decrease. On the other hand, the willingness to pay for private-rented

[29] In our sample, there are 539 properties that transform from private-rented to owner-occupied. There are 468 properties that transition from owner-occupied to private-rented and there are 1,651 properties that remain private-rented (the remaining ones remain owner-occupied).

[30] The net effect of this subset is $-£9,285 + £2,281 - £2,051 = -£9,055$. The effect on the subset of properties that were rented and become owner-occupied is $-£9,285 + £8,137 + £6,653 = £5,505$.

[31] Given by the estimated coefficient of $\mathbb{1}(i = \text{Remained Rented})$.

[32] Given by sum of the estimated coefficients $-£9,285 + £10,376 = £1,091$.

properties should have decreased due to additional BTL taxes, which we outlined above. But buyers of private-rented properties do not believe that price drops beyond a certain level vis a vis owner-occupied ones are justified. Hence, we do not observe an additional effect of the MEES 2018 policy on low-EPC properties that remain rented.

Table 4: The DiD effect accounting for changes in tenure

Dependent variable	$\Delta Price$	$\Delta Price$
	(1)	(2)
$\mathbb{1}(EPC < E)$	$-8,910$ (1,053) $p < 0.001$	$-9,285$ (1,123) $p < 0.001$
$\mathbb{1}(EPC < E) \times \mathbb{1}$(Owner To Rented)		2,281 (11,469) $p = 0.842$
$\mathbb{1}(EPC < E) \times \mathbb{1}$(Rented To Owner)		8,137 (9,776) $p = 0.405$
$\mathbb{1}(EPC < E) \times \mathbb{1}$(Remained Rented)		10,376 (3,634) $p = 0.004$
$\mathbb{1}$(Owner To Rented)		-2,051 (3,355) $p = 0.541$
$\mathbb{1}$(Rented To Owner)		6,653 (3,166) $p = 0.036$
$\mathbb{1}$(Remained Rented)		-21,057 (2,226) $p < 0.001$
Region FE	YES	YES
Observations	20,064	20,064
R^2	0.349	0.352

Note: This table shows the results on the matched sample. Column 1 shows the regression model in eq. (2). Column 2 shows the regression model in eq. (4). Standard errors are reported in brackets. All specifications are estimated in first differences. Source: House price data are from HM Land Registry. EPC ratings are from the MHCLG.

5.1.4 Sensitivity Analysis

We conduct several sensitivity tests. First, we assess the sensitivities of our results with respect to outliers. Specifically, we repeat our analyses on a sample which trims our data at the 5th and 95th percentile of the

Price variable.[33] Interestingly, the size of the DiD effect decreases slightly from £8,910 to £5,072. It remains statistically significant at all conventional levels (p-value: < 0.0001). We also see that the model fit, measured by the R^2 value, increases. For completeness, we also performed the DID regression adjusting for spatial autocorrelation via Conley Regression (Conley, 1999). The conclusions are consistent (p-value < 0.001) with those reported in Table 2 and the results are robust for a range of distance cut-offs.

Second, we conduct a logarithmic transformation of our Price variable in the PSM-matched data set. In this case, the GAM model still does not provide significant evidence against the parallel trend assumption in the pre-intervention period. The subsequent DiD analysis demonstrates a negative effect of -0.0221 which is again statistically significant (p-value < 0.0001). With the logarithmic transformation of the price variable, the interpretation of the DiD effect changes. We can now interpret it as a change in price relative to the initial property price. I.e. the estimated coefficient corresponds to a roughly 2.2% reduction in pre-intervention price.[34] The R^2 increases to about $R^2 = 0.478$, again accounting for regional fixed effects.

Overall, we conclude that our main conclusion is robust to different modelling choices of the dependent variable. It is also robust to outliers as indicated by the trimming exercise. Whilst the logarithmic transformation of the price variable may provide a more appropriate statistical model, we prefer to present our main analysis in terms of absolute price changes as its interpretation more closely aligns with the conceptual framework outlined above.[35] In summary, there is evidence that prices of the least energy-efficient properties affected by the MEES 2018 policy intervention decreased by about £5k to £9k relative to unaffected ones.

5.2 Validity of the parallel trends assumption

The inference for the estimated intervention effect using our Difference-in-Difference (DiD) design is valid only under the *parallel trend assumption*.[36] It means that the outcomes in treatment and control group follow the same time trend in the absence of the policy intervention (Lechner, 2010; Ryan et al., 2019). This assumption is not testable as one cannot verify the counterfactual, i.e. the price trend of treated properties had they not been affected by the a policy intervention. Existing research typically looks for evidence against the assumption in the pre-intervention period, i.e. the time leading up to the MEES 2018 policy intervention. But standard approaches to verifying (or not rejecting) it do not typically go beyond visual inspections of trends in the outcome variable, or testing for differences in linear pre-intervention trends (Ryan et al., 2019; Roth, 2020; Rambachan and Roth, 2019). Both methods come with drawbacks. While the first approach does not allow for statistical inference, the second method assumes a linear relationship between the outcome variable and the time variable.

Given the importance of the *parallel trend assumption* for the validity of our Difference-in-Difference

[33] Under this trimming, the performance of the different PSM algorithms is similar to that of Section 4.4, with the *XGboost* model again performing best. Using this model for PSM and then estimate the GAM, which leads to a deviance explained of 48.7% with a difference-in-smooth p-value of 0.781, again demonstrating little evidence against the common trend assumption.

[34] Consider the proportional decrease in value for the mean of our sample $2.2\% \times £224,000 \approx £4,900$. The scale of this effect aligns with that reported in the trimmed sample.

[35] Last, we also cluster standard errors at different levels. Results do not change and they are available upon request.

[36] It is also known as the *common trend assumption.*

estimates, we propose a novel, more realistic approach. It avoids the assumption of a linear relationship between outcome and time. Instead, we present a novel use of a *Generalized Additive Model (GAM)* (Hastie and Tibshirani, 1986). The advantage of a *GAM* is that it allows for smooth non-linear trends over time. Similar to our set-up, Rose et al. (2012) estimate trends in the pollutant profiles at individual sites in three regions of Scotland. However, our application to testing the parallel trend assumption of the DiD model appears novel.

In Appendix D, we illustrate how a *GAM* can be used to model property prices as smooth non-linear functions of a time trend. In the context of our DiD analyses, we aim to understand if time trends differ across properties affected by the MEES 2018 from those unaffected prior to its implementation in April 2018. To that end, we estimate the difference in smooth time trends across these two groups of properties. We then test if the difference in time trends is a statistically significant different from zero. To that end, we fit a *GAM* with factor-smooth interactions which we implement with the model

$$\text{Price}_i = \beta + f(\text{Time}_i) + (\beta_\Delta + f_\Delta(\text{Time}_i))\,(1 - \mathbb{1}(EPC_i < E)) + \gamma^\top \mathbf{X_i} + \epsilon_i \;, \tag{5}$$

where the subscript $i = 1, \ldots, n$ indexes the properties. The model components are as follows:

- Price_i is the price variable of property i;
- β is the intercept that represents the mean price of a treated property, with EPC ratings of below E. β_Δ is the difference between the mean price of a non-treated property and a treated property;
- $f(\text{Time})$ and $f_\Delta(\text{Time})$ are the smooth functions of Time which is a continuous time covariate; $f(\text{Time})$ represents the time trend of prices for the treatment group, those low EPC properties affected by the policy.[37] By contrast, $f_\Delta(\text{Time})$ represents the difference in time trends between treated and non-treated groups.
- $\mathbb{1}(EPC_i < E) = 1$, if property i has an EPC rating below E in post-intervention period, and zero otherwise;
- $\mathbf{X_i}$ is a vector of property characteristics;[38]
- ϵ_i are idiosyncratic errors that are distributed as follows $\epsilon_i \sim N(0, \sigma^2)$.

We implement this *GAM* on our PSM-matched sample using cubic regression spline smooths, with the knots distributed evenly throughout the time covariate values.[39] The estimated functions $\hat{f}(\text{Time})$ and $\hat{f}_\Delta(\text{Time})$ are given in Figure 5. Figure 5a suggests that the estimated price trend for the treatment group is non-linear, which justifies modelling it via *GAM*. More importantly for the parallel trend assumption is the difference smooth $\hat{f}_\Delta(\text{Time})$ which we plot in Figure 5b). It possesses a confidence interval that consistently

[37] In our construction time is represented as the number of days from the first observation in our sample.

[38] We control for a large set of characteristics that may be relevant for hedonic pricing models, i.e. X={REGION, DEMOGRAPHIC, PROPERTY TYPE, TOTAL FLOOR AREA, CONSTRUCTION AGE, TENURE, HABITABLE ROOMS}.

[39] We use the *mgcv* package (Wood, 2017) in R to estimate the *GAM* model parameters.

includes zero prior to the policy intervention in the time period from 2015 until 2018. We interpret it as insufficient evidence to reject the hypothesis of different time trends across our two groups.

Figure 5: Fitted price trends

(a) Price trend for treated properties $\hat{f}(\text{Time})$ (b) Price trend for the difference $\hat{f}_{\Delta}(\text{Time})$

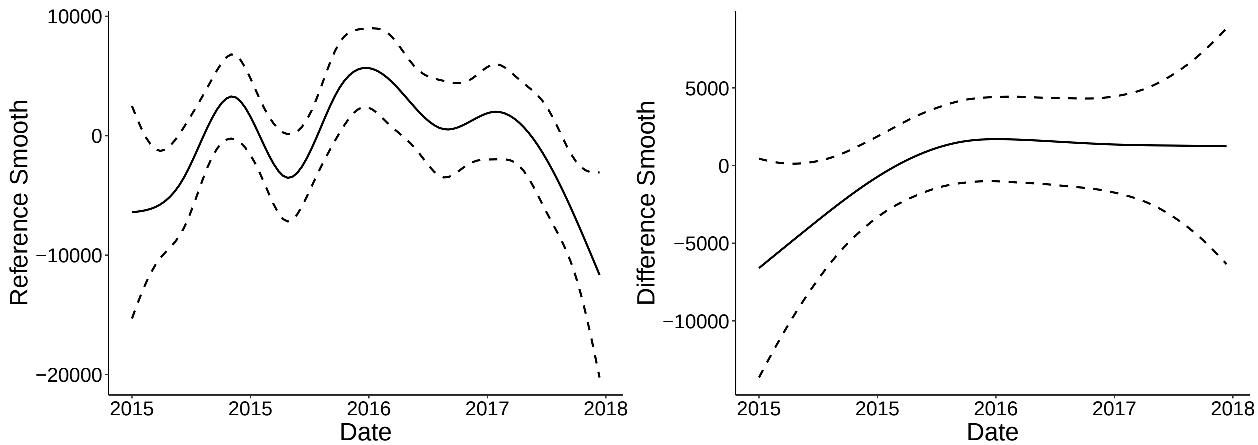

Note: Panel a) shows the fitted price trend for treated properties $\hat{f}(\text{Time})$ from the *GAM* of Eq. 5. Panel b) shows the estimated smooth for the difference between treatment and control groups $\hat{f}_{\Delta}(\text{Time})$. Dashed lines represent the estimated 95% confidence interval. Estimates are obtained from the PSM-matched samples.

We provide more formal inference for the difference smooth $\hat{f}_{\Delta}(\text{Time})$ which sheds light on the validity of the parallel trend assumption. Using our construction the F-statistic reported for this smooth is calculated under the null hypothesis: H_0 the difference between the smooth trends of treated and non-treated properties is equal to zero, with the alternative H_1 being that it is different from zero. An F-statistic that corresponds to a small p-value suggests that the trends for the two groups are different, in our case, we obtain a p-value of 0.234 (F-statistic 1.617) for the significance of $\hat{f}_{\Delta}(\text{Time})$, and a p-value of 0.001 (F-statistic 3.385) for the reference smooth $\hat{f}(\text{Time})$.[40] Overall, looking at evidence from our *GAM* plots when coupled with the summary F-statistic, we find there is insufficient evidence to reject the null hypothesis at the 5%, or even 10% level. In this case, the parallel trends assumption does seem reasonable and gives us confidence in the validity of our Difference-in-Difference (DiD) results.[41]

We argue that this test of the parallel trend assumption indirectly tests for anticipation effects of the policy. This addresses the concern that the legislation was passed in 2015. Buyers might have reacted to this event rather than the implementation of the policy in 2018.[42]

[40] It is worth noting that there is a difference in property prices based on treatment/control ($\hat{\beta}_{\Delta}$ =10,500, p-value < 0.001), however, such constant differences are allowed under the parallel trend assumption.

[41] We acknowledge that lack of evidence against a null does not imply we can accept it.

[42] To verify the robustness of our conclusions we also estimated a *GAM* model which included reference and difference smooths for each region, e.g. London, North-East, etc. While different regions possessed different reference functions, the difference smooths were never significantly different from zero, thus the parallel trends assumption appears reasonable even at a regional scale. Furthermore, visual inspection, and analysis of a linear trend model (not reported) verify the conclusions from the *GAM* model presented here.

5.3 External validity of the sample

Our empirical DiD strategy comes with the advantage of high internal validity as it reduces the relevance of confounding factors. Here, we discuss if our estimates generalise to the entire UK property market.

In a first step, we examine if the selection of our repeated sales sample is relevant. To that end, we compare the characteristics of properties in our sample of repeated sales properties with a random sample of all properties in England and Wales, which should be representative of the full population.[43] We present this descriptive comparison in Table 5, where Column 1 shows the characteristics of the random sample and Column 2 illustrates the characteristics of our repeated sales sample. For many variables, we see close alignment, for instance the size of properties, as measured by TOTAL FLOOR AREA, and CONSTRUCTION AGE. However, REGION, DEMOGRAPHIC, and TENURE, present significant differences. In terms of TENURE, the random sample of the full population of all properties, shown in Column 1, suggests that 19% of properties are socially rented. However, these properties are all but absent from our repeated-transaction sample shown in Column 2. Staying with the TENURE variable, we see that there is also a large imbalance with private rentals in the population (12.4%) compared to the random sample (21.8%). It appears that properties that are likely to be rented are transacted less often before and after the policy intervention. Moreover, we observe that our repeated sales sub-sample contains a smaller fraction of properties located in London. This suggests that properties in the capital were not sold as frequently, an attribute which may be correlated with either the presence of more social housing, and/or privately rented houses being transacted less in this region. Regarding the demographic variables, the imbalance in these is largely expected given the regional discrepancies, specifically our relative under-sampling of London.

In a second step, we then examine whether the PSM is relevant. To that end, we present summary statistics of the PSM-matched sample in Column 3 of Table 5. As we employ PSM on the repeat sales data set, we expect our final sample to inherit some patterns from the repeated sales sample. Indeed, there is evidence of under-sampling of rented accommodation, a pattern that also exists in the repeated sales data set presented in Column 2. In addition to that, matching might introduce further discrepancies compared to the repeat sales data set. Indeed, we observe such additional differences: The PSM-matched data set undersamples properties of specific ages, particularly built after 1967 and before 1900.

Comparing our final sample of PSM-matched properties with the repeated-sales sample and the random sample, we observe some difference in property characteristics. It suggests our results may not be fully representative of the entire population, e.g. for properties located in the London region. Most importantly, our results should not be projected onto the value of social housing stock, which due to the requirement for repeated sales is very hard to assess empirically via market pricing. Last, it is interesting to observe that our PSM-matched sample does not include as many new nor very old properties.[44] With these discrepancies in mind, the analysed PSM-matched sample seems not fully but reasonably representative for a large fraction of properties in most of the regions in England and Wales.

[43]By the Law of Large Numbers, the means of our random sample should converge against the true expected value.

[44]New properties will be unlikely to be of a EPC rating E or below, whereas old dwellings are very likely energy-inefficient with ratings below E.

Table 5: Comparison of property characteristics (population vs. repeated sales vs. matched)

Sample	Population	Repeated sales	PSM-matched
	(1)	(2)	(3)
PROPERTY TYPE		(SMD=0.272)	(SMD=0.456)
Bungalow	9.4%	9.6%	11.7%
Flat	25.3%	15.1%	9.2%
House	62.7%	73.5%	77.8%
Maisonette	2.6%	1.8%	1.3%
Park home	0.0%	0.0%	0.0%
CONSTRUCTION AGE		(SMD=0.106)	(SMD=0.816)
before 1900	11.7%	12.6%	0.0%
1900-1929	15.2%	16.4%	27.2%
1930-1966	31.7%	27.8%	41.6%
1967-1995	29.4%	28.8%	31.3%
1996-2006	9.6%	11.7%	0%
2007 onwards	2.5%	2.7%	0%
TENURE		(SMD=0.749)	(SMD=0.796)
owner-occupied	59.0%	86.8%	89.1%
rental (private)	21.8%	12.4%	10.2%
rental (social)	19.2%	0.9%	0.8%
DEMOGRAPHIC		(SMD=0.380)	(SMD=0.520)
Constrained city dwellers	9.0%	5.5%	5.0%
Cosmopolitans	7.0%	4.7%	4.0%
Ethnicity central	7.6%	2.7%	2.0%
Hard-pressed living	17.7%	17.8%	15.4%
Multicultural metropolitans	14.2%	9.9%	9.8%
Rural residents	9.8%	10.3%	23.6%
Suburbanites	15.9%	21.7%	16.0%
Urbanites	18.8%	27.5%	24.2%
REGION		(SMD=0.287)	(SMD=0.387)
East Midlands (England)	8.3%	10.5%	11.6%
East of England	9.6%	11.2%	10.8%
London	15.2%	6.8%	5.8%
North East (England)	5.1%	4.2%	2.8%
North West (England)	13.2%	13.8%	13.4%
South East (England)	15.3%	16.5%	13.2%
South West (England)	8.3%	10.2%	13.1%
Wales	5.5%	5.6%	7.8%
West Midlands (England)	9.8%	10.5%	10.9%
Yorkshire and The Humber	9.8%	10.7%	10.6%
TOTAL FLOOR AREA		(SMD=0.056)	(SMD=0.079)
Floor area	$87.34m^2$	$84.87m^2$	$91.03m^2$
HABITABLE ROOMS		(SMD=0.188)	(SMD=0.301)
less than 4	35.3%	28.4%	22.3%
4 to 5	45.6%	54.8%	58.0%
more than 5	19.1%	16.8%	19.6%
Number of properties	(N=443,574)	(N=147,842)	(N=20,064)

Note: This table shows a random sample of the population of all properties (column 1). It also shows the sample of repeated property sales (column 2) and the PSM-matched sample employed in the main analysis (column 3). For the continuous covariate TOTAL FLOOR AREA the mean is displayed. For each categorical covariate, percentages of each level are reported. Source: House price data are from HM Land Registry. EPC ratings are from the MHCLG. Demographic data are from the ONS.

6 Conclusion

6.1 Summary

We examine the introduction of a Minimum Energy Efficiency Standard (MEES 2018) that came into force in England and Wales on 1 April 2018. It aimed at encouraging landlords and property owners to improve the energy efficiency of their properties with the goal of reducing overall greenhouse gas emissions. It restricted the granting and continuation of existing tenancies of energy-inefficient properties with an EPC rating of F and G. Our results suggest that the prices of such affected properties decreased, on average, by about £5k-£9k relative to unaffected ones. The magnitude of this estimated decrease in price is consistent with changes in buyers' willingness to pay, when they are well informed about EPC ratings of the properties as well as details of the policy in terms fines and caps on costs for improvements. Once one considers individual risk aversion, expectations of additional policy tightening in the future as well as the higher expected costs from that, the magnitude of this effect appears realistic. As both the policy intervention as well as EPC ratings are salient to both buyers and sellers in the house market, we interpret our evidence as being consistent with the notion of semi-efficient housing markets. While our analyses suggest a strong reaction to the implementation of the policy in April 2018, we do not observe an additional effect of the adjustment of the policy in April 2019. Instead, this policy amendment reduced the magnitude of the drop in prices of energy-inefficient properties under the scope of the policy compared to energy-efficient ones. Our external validity analyses suggest these results are valid for a subpopulation of properties which are mid-age houses of average size.

6.2 Discussion

There are several further aspects of the introduction of minimum energy efficiency standards which go beyond the initial price effect. First, this estimated effect of the MEES 2018 seems to be driven by the subset of properties that stay owner-occupied. By contrast, the subset of properties remaining rented decreases in price following the implementation of the policy regardless of their EPC ratings. This could be rationalised by potential buyers perceiving this policy change as a signal that goes beyond its current scope. One explanation is that buyers of private-rented properties might believe that there is a risk of additional policy tightening in the future beyond EPC rating E. As a result, we see a price drop across all EPC ratings of private-rented properties. The estimated effect may be larger than what the examined policy would justify, although possibly not enough to align them to long-term climate policy goals.

The possibility that buyers change their expectations and react strongly to short-term changes in policy might have ambiguous consequences for transition risk of financial stability. On the one hand, one could argue that it is enhancing financial stability as it means that buyers perceive the MEES 2018 as a signal from the government of future climate policies. As a result, they become more aligned to long-term policy goals than they were before the policy change and than the policy change would justify. On the other hand, it points to transition risk as it is a reminder that the market reaction can exceed the one that could be expected from a given policy change keeping expectations about future policy tightening constant. As long as the market is not aligned with long-run policy goals, even gradual policy changes, e.g. Diluiso et al. (2020), may lead to more abrupt reactions (Carattini et al., 2021) because they may lead expectations to change.

The financial stability implications for this specific policy intervention with modest calibration and limited jurisdiction can be measureable. In theory, it decreases in property prices mean that collateral values of outstanding mortgages decrease. Once outstanding loan amounts exceed the value of collateral, mortgage lenders would incur losses if borrowers defaulted on their mortgage payments. This can be a contributing factor for financial instability depending on the share of such mortgages outstanding on banks' balance sheets and the size of the price drop. Similarly, it might have implications for the wealth distribution among home owners if property values are not evenly distributed across energy efficiency ratings. Indeed, the least energy-efficient properties are less expensive than more energy-efficient ones, as indicated in Figure 3. The MEES 2018 should have decreased the price of these properties which meant the wealth equality among UK homeowners has decreased. For this particular policy change, we expect the implications for both wealth inequality and financial instability to be limited as our estimated effect is only small, about 2% of property values, and the subset of energy-inefficient properties properties make up only a small fraction of the housing stock, in our sample about 7%.[45]

The risk of stranding of houses also depends on other factors, such as the share of properties with potential EPC ratings close to their actual EPC ratings. While some properties can be improved as a result of the introduction of this policy (mitigating the risk of stranding), others cannot be improved because their actual EPC ratings are close to their potential EPC ratings. Alternatively, some homeowners may be credit-constrained and hence unable to find funding for their home improvements. On average, such credit constraints might become more binding in geographic areas where costs of improvements are high relative to property prices.

6.3 Implications

There are aspects of climate policy and transition risks that are beyond the scope of this paper. Future research can investigate the real estate market implications of similar policies that are differently calibrated, for example higher thresholds for exemptions, higher fines, or broader coverage in terms of EPC ratings under the scope of the policy. Such policies might have stronger impact on house prices which could create a case for the intervention of central banks and financial regulation. They could intervene via macroprudential policies if such interventions indeed undermined financial stability and led to transition risk (D'Orazio and Popoyan, 2019). Indeed, there is an emerging literature analyzing the role in tackling climate change for central banks (Dafermos et al., 2018) and financial regulators in particular (Campiglio, 2016; Campiglio et al., 2018) has suggested broad measures, such as taxes or subsidies on banks' assets (Carattini et al., 2021), our analyses point to a specific type of risk which should manifest itself as adjustments to collateral values in lenders' mortgage books. This points to constraining loan-to-value ratios via quantitative limits (Carlos Hatchondo

[45] Moreover, in terms of financial instability, only a small subset of these properties would be in the critical range of the loan-to-value ratios (LTV) distribution, i.e. of at least 90%. This critical range may be quite small, but one must consider that a bank incurs costs to repossess the property and sell the property, possible below market value. There may be an additional systemic risk element if borrowers with low EPC properties are more credit constrained.

et al., 2015), changes in capital buffers (Basten, 2020) or change in business models following climate stress tests (Battiston et al., 2017). However, the answer to which of these measures would be appropriate is beyond the scope of this paper.

We also leave it to future research to examine whether the policy has achieved its primary objectives in achieving its goal to reduce energy demand and greenhouse gas emissions by improving the energy efficiency of privately rented homes, contributing to the government's climate change commitments. It would also be of interest to examine if investments in energy efficiency are indeed a function of household wealth or income or individual preferences. On the one side, financial frictions as credit constraints limit the amount of debt financing, often limited by the value of the underlying collateral (Brunnermeier et al., 2012). Hence, investing in energy efficiency might be a matter of wealth or income. On the other side, individual time preferences systematically influence willingness to invest in energy efficiency (Newell and Siikamäki, 2015), whilst political preference and media coverage may also have an impact (Bernstein et al., 2021). Finally, one may consider extensions to this study to evaluate the longer-term implications going beyond our period of study.

A Data

Table 6 provides a description of the variables from the property transaction data set, sourced from the HM Land Registry website. Each row in the data set represents a property sale that took place in England and Wales.[46]

Table 6: Description of the property transactions data set

Variable	Description
ID	Unique transaction reference number (nominal)
PAON	Primary Addressable Object Name. Typically the house number or name (nominal)
SAON	Secondary Addressable Object Name. Where a property has been divided into separate units (e.g. flats), the PAON will identify the building and the SAON will specify the separate unit (nominal)
Street	The street that the property is located upon (nominal)
Locality	Optional locality information, a more specific location of the property (nominal)
Town	The post town of the post address of the property (nominal)
POSTCODE	The postal code of the property (nominal)
Date	Date when the sale was completed, as stated on the transfer deed (date)
Price	Sale price, as stated on the transfer deed (continuous)

Note: This table provides a description of the property transactions data set obtained from HM Land Registry (HM Land Registry, 2014)

Table 7 provides a description of the variables in the data set covering energy performance certificates (EPC) which we used in our analysis. It is sourced from the Ministry of Housing, Communities & Local Government website (Ministry of Housing, Communities & Local Government, 2020b). Each row in the data set represents information regarding the EPC rating that has been issued for a specific building or building unit in England and Wales.

[46] We note that there is zero variation in the category *Park home* in Table 1 for the PROPERTY TYPE variable as there is only one row with that value. In order to resolve the issue of zero variation, this row was deleted.

Table 7: Description of the EPC data set

Variable	Description
LMK KEY	Unique individual lodgement identifier (nominal)
ADDRESS1	First line of the address (nominal)
ADDRESS2	Second line of the address (nominal)
ADDRESS3	Third line of the address (nominal)
POSTCODE	The postal code of the property (nominal)
CURRENT ENERGY RATING	Current energy rating of the property converted into a linear 'A to G' rating scale (ordinal)
CURRENT ENERGY EFFICIENCY	Current energy rating value of the property, ranging from 1 to 100 (discrete)
PROPERTY TYPE	The type of property (Bungalow/Flat/House/Maisonette/ Park home) (nominal)
INSPECTION DATE	The date that the inspection was carried out by the energy assessor (date)
LOCAL AUTHORITY	Office for National Statistics (ONS) code, giving the local authority area in which the building is located (nominal)
TOTAL FLOOR AREA	The total useful floor area (continuous)
NUMBER HABITABLE ROOMS	The number of habitable rooms, including any living room, sitting room, dining room, bedroom, study and similar (discrete)
ADDRESS	Field containing the concatenation of ADDRESS1, ADDRESS2, and ADDRESS3 (nominal)
POSTTOWN	The post town of the property (nominal)
CONSTRUCTION AGE	The age band when the building was constructed (before 1900/1900-1929/1930-1949/1950-1966/1967-1975/ 1976-1982/1983-1990/1991-1995/1996-2002/2003-2006/ 2007 onwards) (nominal)
TENURE	The tenure type of the property (owner-occupied/rental (private)/rental (social)) (nominal)

Note: This table provides a description of the EPC data set, as obtained from the Ministry of Housing, Communities & Local Government (Ministry of Housing, Communities & Local Government, 2020b).

For completeness, we also consider an overview of the 2011 Area Classification for Output Areas, as described by the Office for National Statistics (2015). We use classification at the supergroup level, forming the top tier of the clustering hierarchy. These eight groups provide the most generic descriptions of the population in the UK. Descriptions for supergroups and a discussion of their characteristics can be found in (Office for National Statistics, 2015).

B Standardized Mean Difference (SMD)

In this section, we provide the formal illustration of the SMD calculation. In the case of continuous variables, such as TOTAL FLOOR AREA, the SMD is calculated as

$$\text{SMD} = \frac{\overline{X}_T - \overline{X}_C}{\sqrt{(S_T^2 + S_C^2)/2}} \, , \tag{6}$$

where $\overline{X}_T$ and $\overline{X}_C$ are the sample means of the variable for treated and non-treated properties, respectively, while S_T^2 and S_C^2 are the sample variances for the treated and non-treated properties.

In the case of multinomial variables SMD is calculated as

$$\text{SMD}_\text{j} = \sqrt{(T_j - C_j)^\top S^{-1} (T_j - C_j)} \tag{7}$$

where $T_j = P(X = \textit{category } j | \textit{treated})$, and $C_j = P(X = \textit{category } j | \textit{non-treated})$ for $j = 1, \ldots, k$, where X is the multinomial variable with k categories, S is the cross-category covariance matrix.

C Propensity Score Matching

Classification for PSM

For all the implemented methods that are mentioned in this section, we split our data into 70% training and 30% test sets. During the split, we ensure that both sets have the original class ratio of the pre-intervention data, with the EPC rating of at least E being 93.21% of the total size, and the EPC below E of 6.79%. Before applying the classifiers on the training data using as predictors TOTAL FLOOR AREA, PROPERTY TYPE, CONSTRUCTION AGE, TENURE, DEMOGRAPHIC, REGION, and HABITABLE ROOMS, we check whether the model required tuning of any hyperparameters. Logistic regression requires none. However, random forest and boosting have a range of hyperparameter values that can be fine-tuned by implementing 5-fold cross-validation (CV) (Hastie et al., 2009, p. 241-249). Note that the boosting method that we implemented, was a more advanced, modern version called *eXtreme Gradient Boosting (XGBoost)* (Chen and Guestrin, 2016), whose usage of parallel computing makes it extremely faster, than the classical boosting algorithm. The main advantages of XGBoost over the classical boosting method are that it calculates the second partial derivatives of the loss function to get to the minimum of the loss function, and uses advanced L_1 norm and L_2 norm regularization, which has the effect of reducing variance and improving the predictive performance of the fitted model.

As a first step, we define the set of parameter values that we want to evaluate for each model. Then for each hyperparameter set of values, a 5-fold CV is implemented, and the hyperparameter set of values that maximized the CV area under the ROC curve (AUC) score, is selected. One way to interpret the AUC score, is to think that if we randomly choose a positive instance and a negative instance, then AUC represents the probability that the classifier ranks the positive instance higher than the negative instance. For the random forest, the hyperparameter set of values that maximized the CV AUC score was equal to $m = 2$, while for XGBoost was $B = 500$ trees, $d = 6$ and $\lambda = 0.01$.

Table 8: AUC values for the classifiers on the test data

	AUC
Logistic Regression	0.752
Random Forest	0.753
XGBoost	0.759

After fitting a main effects logistic regression model, a random forest with $m = 2$, and XGBoost of $B = 500$ trees, $d = 6$ and $\lambda = 0.01$, on the training set, we used the test set to check each model's predictive performance. The value 0.5 was used as a threshold to assign an observation to a class based on the predicted probability. Specifically, if the predicted probability was greater than 0.5, then the property was classified to be class *EPC_LEVEL = Below E*, i.e. belong to the treatment group. On the other hand, if the predicted probability was less than or equal to threshold 0.5, then the property was classified to be class *EPC_LEVEL = At least E*. One way to check how good each classifier separates the two classes is to plot the ROC curve for different decision cutpoints, as shown in Figure 6b. Sensitivity is defined as the true positive rate, and specificity as the true negative rate, with true positive being a correct positive prediction *EPC_LEVEL = Below E* by the classifier, and true negative being a correct negative prediction *EPC_LEVEL = At least E*. In situations when classifiers have ROC curves that intersect, one can use the AUC score, which summarizes the overall performance of each classifier. Therefore, from Table 8, we see that XGBoost is slightly superior to the other models with AUC=0.759.

Implementation of Matching Using Propensity Scores

As a next step, we then calculated the predicted probabilities of treatment assignment, i.e. the probability that the property corresponds to *EPC_LEVEL = Below E*, in the whole pre-intervention sample of 147,841 rows, using an XGBoost algorithm of 500 trees, $d = 6$ and $\lambda = 0.01$. Then, the *MatchIt* package (Ho et al., 2011) was used to conduct PSM with one-to-one nearest neighbor method, that matches each treated property to only one non-treated property based on their proximity in terms of the propensity scores. More precisely, for $i = 1, 2, ..., n_{\text{treated}}$, where n_{treated} is the number of treated properties in our sample, at the ith matching step. As such, the ith treated property was matched to the closest non-treated property that was not previously matched.

To visualize the balance of the resulting matched data set Figure 6a shows four histograms: the original treatment and control groups, and the matched treatment and control groups. The histograms for the two groups prior to matching on the left side, differ to a great degree, while the histograms after matching on the right side are practically the same. Therefore, graphically, we see that the matching was successful.

Figure 6: Propensity scores

(a) Histograms of propensity scores before and after matching

(b) ROC curve for classifiers

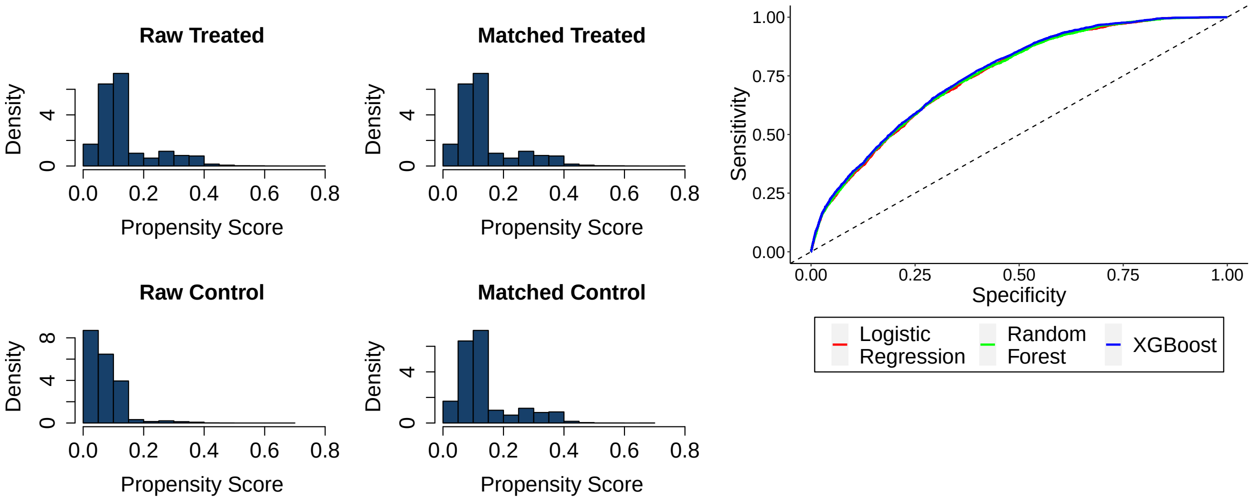

Note: Panel a) shows histograms of propensity scores before and after matching. Panel b) shows ROC curve for classifiers. Source: House price data are from HM Land Registry. EPC ratings are from the MHCLG.

Indeed, PSM matching improves the covariate balancing across treatment and control groups. Table 9 compares SMD values of each covariate across our three data sets: in the first column, we show the unmatched data set before we employ matching, and the second and third summarise the matched data sets that come from respectively using the Logistic regression and Boosting models. We note that for both the logistic and boosted methods SMD values are now lower than 0.1 for all variables. We conclude that any difference that existed between the two the treatment and control group prior to matching has been reduced dramatically after implementing PSM. In addition to the slightly higher AUC, it is now clear the boosted method provides superior alignment of covariate balance as measured through SMD. For this reason, we choose the boosted model to continue our analysis confident that the sample is far more comparable than the unmatched data set.

Table 9: Comparing PSM performance with the unmatched data

	Unmatched	Matched (Logistic)	Matched (XGBoost)
	(1)	(2)	(3)
PROPERTY TYPE	0.203	0.028	0.017
CONSTRUCTION AGE	0.924	0.013	0.008
TENURE	0.071	0.048	0.026
DEMOGRAPHIC	0.397	0.037	0.020
REGION	0.181	0.026	0.044
TOTAL FLOOR AREA	0.160	0.013	0.016
HABITABLE ROOMS	0.155	0.040	0.007
Total	2.091	0.205	0.138

Note: This table shows SMD values for three different data sets. Column 1 shows SMD values in the unmatched sample. Column 2 shows them in the matched sample using a Logistic regression. Column 3 shows them in the matched sample using XGBoost. Source: House price data are from HM Land Registry. EPC ratings are from the MHCLG. Demographic data are from the ONS.

D Generalised Additive Pricing Model

The basic GAM model enables non-linear functions of some, or all, of the covariates, while maintaining a simple additive model structure (Hastie and Tibshirani, 1986). Specifically, GAMs have the following form:

$$y_i = \beta_0 + \sum_{j=1}^{p} f_j(x_{ij}) + \epsilon_i, \quad i = 1, 2, ..., n \tag{8}$$

where y_i is the ith observation of outcome Y, x_{ij} is the ith observation of explanatory variable X_j, and f_j is a non-linear function for explanatory variable X_j. GAMs have been used to model trends as a smooth, non-linear function over time for a wide variety of data, such as for palaeoenvironmental time series (Simpson, 2018), high-frequency water-quality data Yang and Moyer (2020), or survey data for bird populations (Fewster et al., 2000). In this paper, we set up the GAM to use factor-smooth interactions, which allows us to formally test for evidence against the parallel trends assumption. Before implementing the GAM on the PSM-derived matched data set of comparable pre-intervention properties, we first discuss the underlying structure of the generalized additive model approach.

In order to move away from imposing a predetermined linear form on f_j, the GAM attempts to estimate the shape of f_j from the data whilst imposing smoothness constraints to avoid overfitting. Given these restrictions, the f_j non-linear functions in Equation 8 are referred to as *smooth functions (smooths)*. In order for the smooth to be represented in a parametric form that can be estimated, it needs to be specified using a set of *basis functions*. Let $b_{kj}(X_j)$ represent the kth basis function for the smooth of covariate X_j. Assuming that for X_j there are K basis functions, then the smooth $f_j(X_j)$ can be represented as

$$f_j(X_j) = \sum_{k=1}^{K} \beta_{kj} b_{kj}(X_j) \ , \tag{9}$$

where β_{kj} is the coefficient of the kth basis function that must be estimated (Wood, 2017, p. 162).

There are many ways to specify basis functions, such as in cubic polynomial regression, where the basis functions are: $b_{1j}(X_j) = X_j$, $b_{2j}(X_j) = X_j^2$, and $b_{3j}(X_j) = X_j^3$. We investigate a number of different basis functions for fitting the GAMs in this paper (e.g. thin plate regression splines, cubic regression splines, P-splines). They all give the same conclusions. To keep the paper concise, we only give details for the piecewise cubic polynomial model, under the constraint that the first and second derivatives of the piecewise polynomials are continuous at the knots, i.e. that the piecewise polynomial must be continuous and smooth. The parameters of the GAM model are traditionally estimated via a *backfitting algorithm* (Hastie and Tibshirani, 1986).[47]

[47] In this paper, we use the *mgcv* R package following guidance provided in the text book by Wood (2017).

References

Anas, A. and S. J. Eum (1984). Hedonic analysis of a housing market in disequilibrium. *Journal of Urban Economics 15*(1), 87–106.

Arnould, G., B. Guin, S. Ongena, and P. Siciliani (2020). (When) do banks react to anticipated capital reliefs? *Bank of England Staff Working Paper 889*.

Bach, L., L. E. Calvet, and P. Sodini (2020). Rich pickings? Risk, return, and skill in household wealth. *American Economic Review 110*(9), 2703–47.

Bardhan, A., R. Edelstein, and D. Tsang (2008). Global Financial Integration and Real Estate Security Returns. *Real Estate Economics 36*(2), 285–311.

Barkham, R. and D. Geltner (1995). Price Discovery in American and British Property Markets. *Real Estate Economics 23*(1), 21–44.

Barkham, R. J. and D. M. Geltner (1996). Price Discovery and Efficiency in the UK Housing Market. *Journal of Housing Economics 5*(1), 41–63.

Basten, C. (2020). Higher Bank Capital Requirements and Mortgage Pricing: Evidence from the Counter-Cyclical Capital Buffer. *Review of Finance 24*(2), 453–495.

Battiston, S., A. Mandel, I. Monasterolo, F. Schütze, and G. Visentin (2017). A Climate Stress-test of the Financial System. *Nature Climate Change 7*(4), 283–288.

Bernstein, A., S. B. Billings, M. Gustafson, and R. Lewis (2021). Partisan Residential Sorting on Climate Change Risk. Working Paper 27989, National Bureau of Economic Research.

Bhatia, K. B. (1987). Real estate assets and consumer spending. *The Quarterly Journal of Economics 102*(2), 437–444.

Bolton, P. and M. Kacperczyk (2021a). Do investors care about carbon risk? *Journal of Financial Economics 142*(2), 517–549.

Bolton, P. and M. Kacperczyk (2021b). Global pricing of carbon-transition risk. Technical report, National Bureau of Economic Research.

Breiman, L. (2001). Random Forests. *Machine Learning 45*(1), 5–32.

Brown, M., B. Guin, and K. Kirschenmann (2016). Microfinance Banks and Financial Inclusion. *Review of Finance 20*(3), 907–946.

Brunnermeier, M. K., T. M. Eisenbach, and Y. Sannikov (2012). Macroeconomics with financial frictions: A survey.

Caldecott, B., E. Harnett, T. Cojoianu, I. Kok, and A. Pfeiffer (2016). Stranded assets: A climate risk challenge. *Washington DC: Inter-American Development Bank*.

Caldecott, B., G. Lomax, and M. Workman (2015). Stranded carbon assets and negative emissions technologies. Working Paper.

Caliendo, M. and S. Kopeinig (2005). Some Practical Guidance for the Implementation of Propensity Score Matching. *IZA Discussion Paper Series*.

Campiglio, E. (2016). Beyond carbon pricing: The role of banking and monetary policy in financing the transition to a low-carbon economy. *Ecological Economics 121*, 220–230.

Campiglio, E., Y. Dafermos, P. Monnin, J. Ryan-Collins, G. Schotten, and M. Tanaka (2018). Climate change challenges for central banks and financial regulators. *Nature Climate Change 8*(6), 462–468.

Cannas, M. and B. Arpino (2019). A Comparison of Machine Learning Algorithms and Covariate Balance Measures for Propensity Score Matching and Weighting. *Biometrical Journal 61*(4), 1049–1072.

Carattini, S., G. Heutel, and G. Melkadze (2021). Climate policy, financial frictions, and transition risk. Technical report, National Bureau of Economic Research.

Carattini, S. and S. Sen (2019). Carbon taxes and stranded assets: Evidence from Washington state. *CESifo Working Paper*.

Card, D. and A. B. Krueger (1994). Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania. *American Economic Review 84*(4), 772–793.

Carlos Hatchondo, J., L. Martinez, and J. M. Sánchez (2015). Mortgage defaults. *Journal of Monetary Economics 76*, 173–190.

Case, K. E. and R. J. Shiller (1989). The Efficiency of the Market for Single-Family Homes. *American Economic Review 79*(1), 125–137.

Case, K. E. and R. J. Shiller (1990). Forecasting Prices and Excess Returns in the Housing Market. *Real Estate Economics 18*(3), 253–273.

Chen, T. and C. Guestrin (2016). XGBoost: A Scalable Tree Boosting System. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794.

Clapp, J. M., W. Dolde, and D. Tirtiroglu (1995). Imperfect Information and Investor Inferences From Housing Price Dynamics. *Real Estate Economics 23*(3), 239–269.

Clayton, J. (1996). Rational Expectations, Market Fundamentals and Housing Price Volatility. *Real Estate Economics 24*(4), 441–470.

Conley, T. (1999). Gmm estimation with cross sectional dependence. *Journal of Econometrics 92*(1), 1–45.

Crawley, J., P. Biddulph, P. J. Northrop, J. Wingfield, T. Oreszczyn, and C. Elwell (2019). Quantifying the Measurement Error on England and Wales EPC Ratings. *Energies 12*(18).

Dafermos, Y., M. Nikolaidi, and G. Galanis (2018). Climate change, financial stability and monetary policy. *Ecological Economics 152*, 219–234.

Degryse, H., R. Goncharenko, C. Theunisz, and T. Vadasz (2020). When green meets green. Working Paper Research 392, National Bank of Belgium.

Delis, M., K. de Greiff, and S. Ongena (2021). Being Stranded on the Carbon Bubble? Climate Policy Risk and the Pricing of Bank Loans. Swiss Finance Institute Research Paper Series 18-10, Swiss Finance Institute.

Department for Business, Energy & Industrial Strategy (2017). The Domestic Private Rented Property Minimum Standard. `https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/882957/Domestic_Private_Rented_Property_Minimum_Standard_-_Landlord_Guidance_2020.pdf`. Accessed: February 28, 2023.

Department for Business, Energy & Industrial Strategy (2018). Final Stage Impact Assessment: Amending the Private Rented Sector Energy Efficiency Regulations. `https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/760313/IA_-_Energy_Efficiency__Private_Rented_Property___England.pdf`. Accessed: February 28, 2023.

Di Maria, C., I. Lange, and E. Van der Werf (2014). Should we be worried about the green paradox? Announcement effects of the Acid Rain Program. *European Economic Review 69*, 143–162.

Diluiso, F., B. Annicchiarico, M. Kalkuhl, and J. C. Minx (2020). Climate actions and stranded assets: The role of financial regulation and monetary policy. *CEIS working paper*.

Dolde, W. and D. Tirtiroglu (1997). Temporal and Spatial Information Diffusion in Real Estate Price Changes and Variances. *Real Estate Economics 25*(4), 539–565.

D'Orazio, P. and L. Popoyan (2019). Fostering green investments and tackling climate-related financial risks: Which role for macroprudential policies? *Ecological Economics 160*, 25–37.

Edenhofer, O., M. Kalkuhl, T. Requate, and J. C. Steckel (2020). How assets get stranded: The impact of climate policy on capital and fossil fuel owners. Introduction to the JEEM special section on climate policy and political economy. *Journal of Environmental Economics and Management 100*, 102300.

Efron, B. and T. Hastie (2016). *Computer Age Statistical Inference: Algorithms, Evidence, and Data Science*. Institute of Mathematical Statistics Monographs. Cambridge University Press.

Eichholtz, P., N. Kok, and J. M. Quigley (2010). Doing Well by Doing Good? Green Office Buildings. *American Economic Review 100*(5), 2492–2509.

Englund, P., T. M. Gordon, and J. M. Quigley (1999). The Valuation of Real Capital: A Random Walk down Kungsgatan. *Journal of Housing Economics 8*(3), 205–216.

Fama, E. F. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance 25*(2), 383–417.

Fard, A., S. Javadi, and I. Kim (2020). Environmental regulation and the cost of bank loans: International evidence. *Journal of Financial Stability 51*(C).

Fatouh, M., S. Giansante, and S. Ongena (2021). Economic support during the COVID crisis. Quantitative easing and lending support schemes in the UK. *Economics Letters 209*, 110138.

Fewster, R. M., S. T. Buckland, G. M. Siriwardena, S. R. Baillie, and J. D. Wilson (2000). Analysis of Population Trends for Farmland Birds Using Generalized Additive Models. *Ecology 81*(7), 1970–1984.

Fisher, J., D. C. Ling, and A. Naranjo (2009). Institutional Capital Flows and Return Dynamics in Private Commercial Real Estate Markets. *Real Estate Economics 37*(1), 85–116.

Fowlie, M., M. Greenstone, and C. Wolfram (2018). Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Program. *The Quarterly Journal of Economics 133*(3), 1597–1644.

Freyaldenhoven, S., C. Hansen, and J. M. Shapiro (2019). Pre-Event Trends in the Panel Event-Study Design. *American Economic Review 109*(9), 3307–38.

Friedman, J., T. Hastie, and R. Tibshirani (2000). Special Invited Paper. Additive Logistic Regression: A Statistical View of Boosting. *The Annals of Statistics 28*(2), 337–374.

Fu, Y. and L. K. Ng (2001). Market Efficiency and Return Statistics: Evidence from Real Estate and Stock Markets Using a Present-Value Approach. *Real Estate Economics 29*(2), 227–250.

Fuerst, F., P. McAllister, A. Nanda, and P. Wyatt (2015). Does energy efficiency matter to home-buyers? An investigation of EPC ratings and transaction prices in England. *Energy Economics 48*, 145–156.

Gale, C. G., A. D. Singleton, A. G. Bates, and P. A. Longley (2016). Creating the 2011 Area Classification for Output Areas (2011 OAC). *Journal of Spatial Information Science 12*(2016), 1–27.

Ghysels, E., A. Plazzi, and R. Valkanov (2007). Valuation in us commercial real estate. *European Financial Management 13*(3), 472–497.

Gillingham, K., M. Harding, and D. Rapson (2012). Split Incentives in Residential Energy Consumption. *The Energy Journal 33*(2), 37–62.

Gropp, R., T. Mosk, S. Ongena, and C. Wix (2019). Banks response to higher capital requirements: Evidence from a quasi-natural experiment. *The Review of Financial Studies 32*(1), 266–299.

Hardy, A. and D. Glew (2019). An analysis of errors in the Energy Performance certificate database. *Energy Policy 129*, 1168–1178.

Hastie, T. and R. Tibshirani (1986). Generalized Additive Models. *Statistical Science 1*(3), 297–310.

Hastie, T., R. Tibshirani, and J. Friedman (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (2nd ed.). Springer Series in Statistics. New York: Springer.

Herath, S. and G. Maier (2015). Informational Efficiency of the Real Estate Market: A Meta-Analysis. *Journal of Economic Research 20*(3), 117–168.

HM Land Registry (2014). Price Paid Data [Statistical dataset]. `https://www.gov.uk/government/statistical-data-sets/price-paid-data-downloads`. Accessed: February 28, 2023.

Ho, D. E., G. King, E. A. Stuart, and K. Imai (2011). MatchIt: Nonparametric Preprocessing for Parametric Causal Inference. *Journal Of Statistical Software 42*(8), 1–28.

Ilhan, E., Z. Sautner, and G. Vilkov (2020). Carbon Tail Risk. *The Review of Financial Studies 34*(3), 1540–1571.

James, G., D. Witten, T. Hastie, and R. Tibshirani (2013). *An Introduction to Statistical Learning: With Applications in R*, Volume 103 of *Springer Texts in Statistics*. New York: Springer.

Javadi, S. and A.-A. Masum (2021). The impact of climate change on the cost of bank loans. *Journal of Corporate Finance 69*(C).

Kahn, M. E. and N. Kok (2014). The capitalization of green labels in the California housing market. *Regional Science and Urban Economics 47*, 25–34.

Lechner, M. (2010). The Estimation of Causal Effects by Difference-in-Difference Methods. *Foundations and Trends in Econometrics 4*(3), 165–224.

Lee, B. K., J. Lessler, and E. A. Stuart (2010). Improving Propensity Score Weighting Using Machine Learning. *Statistics in Medicine 29*(3), 337–356.

Lemoine, D. (2017). Green expectations: Current effects of anticipated carbon pricing. *Review of Economics and Statistics 99*(3), 499–513.

Luellen, J. K., W. R. Shadish, and M. H. Clark (2005). Propensity Scores: An Introduction and Experimental Test. *Evaluation Review 29*(6), 530–558.

MacKinnon, G. H. and A. Al Zaman (2009). Real Estate for the Long Term: The Effect of Return Predictability on Long-Horizon Allocations. *Real Estate Economics 37*(1), 117–153.

Mandelbrot, B. (1963). The Variation of Certain Speculative Prices. *Journal of Business 36*(4), 394.

Matikainen, S. (2022). What are stranded assets? https://www.lse.ac.uk/granthaminstitute/explainers/what-are-stranded-assets/. Accessed: 2022-10-20.

McCaffrey, D. F., G. Ridgeway, and A. R. Morral (2004). Propensity Score Estimation with Boosted Regression for Evaluating Causal Effects in Observational Studies. *Psychological Methods 9*(4), 403–425.

Ministry of Housing, Communities & Local Government (2020a). Energy Performance of Buildings Data: England and Wales [Statistical dataset]. https://epc.opendatacommunities.org/?/. Accessed: February 28, 2023.

Ministry of Housing, Communities & Local Government (2020b). Glossary: Domestic EPCs. https://epc.opendatacommunities.org/docs/guidance#glossary_domestic. Accessed: February 28, 2023.

Myers, E. (2019). Are Home Buyers Inattentive? Evidence from Capitalization of Energy Costs. *American Economic Journal: Economic Policy 11*(2), 165–88.

Myers, E. (2020). Asymmetric information in residential rental markets: Implications for the energy efficiency gap. *Journal of Public Economics 190*, 104251.

Newell, R. G. and J. Siikamäki (2015). Individual time preferences and energy efficiency. *American Economic Review 105*(5).

NGFS (2021). Progress report on bridging data gaps. https://www.ngfs.net/sites/default/files/medias/documents/progress_report_on_bridging_data_gaps.pdf. Accessed: February 28, 2023.

Office for National Statistics (2011). Output Area Classification (December 2011) in the United Kingdom [Statistical dataset]. http://geoportal.statistics.gov.uk/datasets/be648a095a9745998f6961e5ba54e01c_0. Accessed: February 28, 2023.

Office for National Statistics (2015). Pen Portraits for the 2011 Area Classification for Output Areas. https://www.ons.gov.uk/file?uri=/methodology/geography/geographicalproducts/areaclassifications/2011areaclassifications/penportraitsandradialplots/penportraits.pdf. Accessed: February 28, 2023.

Rambachan, A. and J. Roth (2019). An Honest Approach to Parallel Trends. Working Paper.

Ramelli, S., A. F. Wagner, R. J. Zeckhauser, and A. Ziegler (2021). Investor rewards to climate responsibility: Stock-price responses to the opposite shocks of the 2016 and 2020 US elections. *The Review of Corporate Finance Studies 10*(4), 748–787.

Reghezza, A., Y. Altunbas, D. Marqués-Ibáñez, C. Rodriguez d'Acri, and M. Spaggiari (2021). Do banks fuel climate change? Working Paper Series 2550, European Central Bank.

Rodnyansky, A. and O. M. Darmouni (2017). The Effects of Quantitative Easing on Bank Lending Behavior. *The Review of Financial Studies 30*(11), 3858–3887.

Rose, N. L., H. Yang, S. D. Turner, and G. L. Simpson (2012). An Assessment of the Mechanisms for the Transfer of Lead and Mercury from Atmospherically Contaminated Organic Soils to Lake Sediments with Particular Reference to Scotland, UK. *Geochimica et Cosmochimica Acta 82*, 113–135.

Rosenbaum, P. R. and D. B. Rubin (1983). The Central Role of the Propensity Score in Observational Studies for Causal Effects. *Biometrika 70*(1), 41–55.

Rosenthal, L. (2006). Efficiency and Seasonality in the UK Housing Market, 1991–2001. *Oxford Bulletin of Economics and Statistics 68*(3), 289–317.

Roth, J. (2020). Pre-Test with Caution: Event-Study Estimates after Testing for Parallel Trends. Working Paper.

Rozenberg, J., A. Vogt-Schilb, and S. Hallegatte (2020). Instrument choice and stranded assets in the transition to clean capital. *Journal of Environmental Economics and Management 100*, 102183.

Rubin, D. B. (2007). The Design versus the Analysis of Observational Studies for Causal Effects: Parallels with the Design of Randomized Trials. *Statistics in Medicine*, 20–36.

Ryan, A. M., E. Kontopantelis, A. Linden, and J. F. Burgess (2019). Now Trending: Coping with Non-Parallel Trends in Difference-in-Differences Analysis. *Statistical Methods in Medical Research 28*(12), 3697–3711.

Samuelson, P. (1965). Proof That Properly Anticipated Prices Fluctuate Randomly. *Industrial Management Review 6*, 41–49.

Sen, S. and M.-T. von Schickfus (2020). Climate policy, stranded assets, and investors' expectations. *Journal of Environmental Economics and Management 100*, 102277.

Setoguchi, S., S. Schneeweiss, M. A. Brookhart, R. J. Glynn, and E. F. Cook (2008). Evaluating Uses of Data Mining Techniques in Propensity Score Estimation: A Simulation Study. *Pharmacoepidemiology and Drug Safety 17*(6), 546–555.

Simpson, G. L. (2018). Modelling Palaeoecological Time Series Using Generalised Additive Models. *Frontiers in Ecology and Evolution 6*, 1–21.

Stuart, E. A., B. K. Lee, and F. P. Leacy (2013). Prognostic Score–Based Balance Measures Can Be a Useful Diagnostic for Propensity Score Methods in Comparative Effectiveness Research. *Journal of Clinical Epidemiology 66*(8), S84–S90.

UK Statutory Instruments (2015). The Energy Efficiency (Private Rented Property) (England and Wales) Regulations 2015. `https://www.legislation.gov.uk/uksi/2015/962/contents/made`. Accessed: February 28, 2023.

UNEP (2021). Adaptation Gap Report 2021. `https://www.unep.org/resources/adaptation-gap-report-2021`. Accessed: February 28, 2023.

Van der Ploeg, F. and A. Rezai (2020). Stranded assets in the transition to a carbon-free economy. *Annual Review of Resource Economics 12*, 281–298.

Walls, M., T. Gerarden, K. Palmer, and X. F. Bak (2017). Is energy efficiency capitalized into home prices? Evidence from three U.S. cities. *Journal of Environmental Economics and Management 82*, 104–124.

Westreich, D., J. Lessler, and M. J. Funk (2010). Propensity Score Estimation: Neural Networks, Support Vector Machines, Decision Trees (CART), and Meta-Classifiers as Alternatives to Logistic Regression. *Journal of Clinical Epidemiology 63*(8), 826–833.

Wood, S. (2017). *Generalized Additive Models: An Introduction with R* (2nd ed.). New York: Chapman and Hall/CRC.

Yang, G. and D. L. Moyer (2020). Estimation of Nonlinear Water-Quality Trends in High-Frequency Monitoring Data. *Science of the Total Environment 715*, 136686.