



Essays on the Consequences of Residential Energy Efficiency on Urban Property Markets

Thomas Luke Siddall

Thesis Supervisors:

Enrico Vanino

Alberto Montagnoli

A thesis submitted in partial fulfilment of the requirements for

the degree of

Doctor of Philosophy

University of Sheffield

School of Economics

July 2026

Abstract

This thesis combines both empirical and theoretical methods in three separate studies exploring the wider impacts of residential energy efficiency improvements on UK property markets. Data from Energy Performance Certificates (EPCs) are merged with property transactions to assess the drivers of such improvements and the response of property prices and household location decisions. I adopt spatial techniques to answer a set of policy-relevant questions. From the Net Zero transition arises a set of ambitious goals. As cities worldwide adapt to these, there has never been a more important time to understand the impacts on urban property markets.

Chapter two employs a novel instrumental variable approach to explore how improvements are capitalised into property prices. An alternate measure of improvement is proposed that measures the difference between a property's current and potential efficiency – termed as *unrealised efficiency*. Capturing the future benefits of improvements, I show that buyers value improvements undertaken by previous owners considerably more than undertaking their own.

Chapter three explores the role of neighbourhood effects in explaining why households undertake improvements. Neighbourhood effects mechanise how households learn from and imitate the behaviour of their neighbours. Leveraging a spatial regression discontinuity design, I assess whether buyers delineate between energy-efficient and inefficient neighbourhoods using information reported in property listing adverts. Buyers with preferences for energy efficiency select into these neighbourhoods. Using years of exposure to selecting households as an instrument, I show that the neighbourhood effect induces improvements among incumbent households.

Chapter four rationalises this with a theoretical model comprising two mutually exclusive groups of households – movers and never movers. Neighbourhood effects increase investment demand which increases property prices in equilibrium. Never movers take

investment decisions on the back of investments made by movers, highlighting that neighbourhood effects arise exclusively in the most-efficient neighbourhoods.

Acknowledgements

There are many people I would like to thank for their help along way.

Firstly, my supervisors, Enrico Vanino and Alberto Montagnoli. None of this would have been possible without their advice and guidance. Thank you for your support throughout what has, at times, felt like a very long journey. I thank you greatly for all you've done for me both personally and professionally. I will never forget it.

I would also like to thank the School of Economics for believing in me and my research project, and for providing me the scholarship to complete it. I thank everybody for their time and effort in shaping me and my research for the next step. In particular, I'd like to thank Jesse Matheson, Antonio Navas, Panos Nanos, Andrew Burlinson, and Jolian McHardy for their extensive feedback. A huge shoutout also to the members of the SUITE research group for the many fun and interesting research meetings that I will truly miss.

To my friends and fellow PhD students that I've met along the way, it's been a pleasure. Adriana, Michael, Maria, Argyris, Shiwen; PhDs can be a lonely business, but with you guys, there was never a dull moment. Thank you for being you. I must give a shoutout to two other special individuals, Lucy Ward and Andrew Bryce, for their unwavering friendship and support.

Finally, I would like to express my upmost gratitude to my family. My mother, Elaine, my father, Iain, and my sister, Sarah, for their support of my development; their sacrifices to get me here; and their love to guide me through. I consider myself very lucky to call them my family. Without them, I dare hazard a guess where life would've taken me. And to my uncles, John and David, my intellectual role models growing up, I thank you for inspiring the curiosity I still carry with me to this day.

Contents

List of Figures	ix
List of Tables	x
1 Introduction	xi
1.1 Motivation	xi
1.2 Background	xii
1.3 Thesis Structure	xiii
1.3.1 Chapter 2: Pricing Potential: Understanding the Value of Unrealised Energy Efficiency	xiv
1.3.2 Chapter 3: Neighbourhood Effects in the Market for Energy Efficient Housing	xv
1.3.3 Chapter 4: Household Location and Investment Decisions Under Neighbourhood Effects: A Theoretical Approach	xvi
2 Pricing Potential: Understanding the Value of Unrealised Energy Efficiency	1
2.1 Introduction	1
2.2 Literature Review	6
2.2.1 Underinvestment in Residential Energy Efficiency	6
2.2.2 Returns to Investment	8
2.2.3 Heterogeneity	11
2.3 Data	15
2.3.1 Measuring Energy Efficiency	16
2.3.2 Changes in the Standard Assessment Procedure in 2012	17
2.4 Empirical Strategy and Results	32
2.4.1 A Fixed Effect Approach to Estimating Current and Unrealised Capitalisation	33
2.4.2 A Two Stage Least Squares Approach to Estimating Unrealised Capitalisation	37
2.5 Heterogeneity Analysis	48

2.6	Conclusion	59
A	Appendix	63
3	Neighbourhood Effects in the Market for Energy Efficient Housing	70
3.1	Introduction	70
3.2	Literature Review	74
3.2.1	Underinvestment in Energy Efficiency: Poor Energy Savings	75
3.2.2	Underinvestment in Energy Efficiency: Asymmetric Information	78
3.2.3	Peer Effects and Individual Behavioural	83
3.2.4	Capitalisation of Energy Efficiency Improvements	87
3.3	Data	89
3.3.1	Energy Performance Certificates Policy	90
3.3.2	Dataset Used to Define Clusters	91
3.3.3	Defining Energy Efficient Clusters	91
3.3.4	Summary Statistics	97
3.3.5	Identification of Neighbourhood Effects	101
3.3.6	Discontinuities in Price	103
3.4	Empirical Analysis	104
3.4.1	Identifying a Premium	104
3.4.2	Premium Estimates	110
3.4.3	Robustness Checks	113
3.5	Mechanisms	117
3.5.1	Identifying Premium Timing	118
3.5.2	Premium Timing Estimates	121
3.5.3	Identifying Returns to Improvements	122
3.5.4	Returns to Improvements Estimates	127
3.5.5	Identifying neighbourhood Effects	130
3.5.6	Neighbourhood Effects Estimates	130
3.6	Conclusion	132

B	Appendix	135
B	Further Technical Information	135
B.1	Merging Energy Efficiency and Property Price Data	135
B.2	Redefining Clusters Within Borough	135
B.3	Intensive Margin Estimates	139
4	Household Location and Investment Decisions Under Neighbourhood Effects: A Theoretical Approach	160
4.1	Introduction	160
4.2	Stylised Facts	165
4.2.1	#1: Moderately Efficient Neighbourhoods Have Investment Potential	166
4.2.2	#2: Investment is Driven by Preferences and Potential	168
4.2.3	#3: Energy Efficiency and Investment Potential Draw Higher Prices	171
4.2.4	#4: Property Prices Reflect Higher Levels of Investment	172
4.3	Model	177
4.3.1	Movers	177
4.3.2	Investment Services	181
4.3.3	Goods Production	185
4.3.4	Housing and Investment Demand	186
4.3.5	Characterising Equilibrium	190
4.3.6	Never Movers	195
4.4	Conclusion	199
4.5	Derivations	202
4.5.1	Household Demand Functions	202
4.5.2	Rewriting Investment Demand	203
4.5.3	Deriving the Supply of Investment Services	203
4.5.4	Deriving the Equilibrium Price of Investment Services Inputs	205
4.5.5	Deriving the Change in Equilibrium Property Prices	205
4.5.6	Characterising Equilibrium: Investment Services Price	207
4.5.7	Characterising Equilibrium: Property Price	208
4.5.8	Characterising Equilibrium: Investment Services	209

4.5.9	Non-Movers' Demand for Investment Services	211
5	Conclusion	216
5.1	Summary of the Findings	217
5.2	Policy Implications	219
5.3	Limitations and Future Research	220
	Bibliography	237

List of Figures

2.1	Mean Energy Efficiency Score by EPC Inspection Year	19
2.2	Mean Energy Efficiency Score by EPC Inspection Year	26
2.3	Continuous Energy Ratings Distribution	28
2.4	Energy Ratings Distribution	29
2.5	Energy Ratings Distribution	30
2.6	Capitalisation by Distance From Potential Rating Band	54
2.7	Capitalisation by Distance From Potential Rating Band	56
2.8	Capitalisation by Distance From Potential Rating Band	58
A.1	Typical Format Energy Performance Certificate	63
A.2	Share of EPC Inspections by Year	64
A.3	Potential Efficiency Scores by EPC Inspection Year and Tenure	65
3.1	Getis-Ord Analysis of Property Listings: London TTWA 2012	95
3.2	Highly- and Poorly-Energy-Efficient Clusters: London TTWA 2012	97
3.3	The Distribution of Energy Ratings at the Border	100
3.4	Price Discontinuity	103
3.5	Efficient Clusters Across All Years: Sale vs No Sale	109
3.6	Event Study Estimates of the Price Premium by Period	122
3.7	Probability of Upgrading at the Border By Household Type	125
B.1	Efficient Clusters in Islington, London: All Years	137
B.2	Efficient Clusters and Administrative Boundaries	143
B.3	The Distribution of Energy Ratings at the Border: No Sale Clusters	143
B.4	Rightmove Neighbourhood Definition	144
B.5	Example Energy Performance Certificate	145
B.6	Typical Energy Efficiency Rating Distribution	146
B.7	Example Property Listing Advert from Rightmove	147
B.8	Average Property Prices: London vs UK	148
B.9	Defining The Spatial Extent of Efficient Clusters	149
B.10	Time Periods Defined in the Event Study	150
B.11	Construction of the Event Study Control Group	151

B.12 Property Prices in Efficient Clusters Relative to the Treatment Period . . .	152
B.13 Event Study Robustness Checks	152
B.14 Properties by Built-Form by Treatment Status	153
B.15 Properties by Built-Form by Treatment Status – No Sale	153
B.16 Properties by Age of Construction by Treatment Status	154
B.17 Properties by Age of Construction by Treatment Status – No Sale	154
B.18 Balancing of Neighbourhood Covariates at Cluster Border	155
4.1 Investment Potential by Neighbourhood	167
4.2 Energy Efficiency by Neighbourhood: London TTWA	168
4.3 Investment Expenditure by Neighbourhood: London TTWA	169
4.4 Investment Potential and Expenditure by Neighbourhood: London TTWA	170
4.5 Energy Efficiency Among Priciest Neighbourhoods: London TTWA	171
4.6 Investment Expenditure and Property Prices (2009-2019): London TTWA	173

List of Tables

2.1	Recommendations by Category	20
2.2	Change in Potential by Recommendation Type	22
2.3	Change in Potential: Marginal Effects	23
2.4	Average Change Between EPC Inspections	27
2.5	Summary Statistics	32
2.6	Fixed Effects: Current Capitalisation	37
2.7	Instrument Exclusion Tests	43
2.8	Unrealised Efficiency Capitalisation: 2SLS Results	45
2.9	Unrealised Efficiency Capitalisation: Standardised 2SLS Results	47
2.10	Unrealised Capitalisation by Recommendation Costs	50
A.1	Full Results – Change in Potential by Recommendation Type	66
A.2	Full Results – Fixed Effects: Unrealised Capitalisation	67
A.3	Full Results – Instrument Exclusion Tests	67
A.4	Unrealised Efficiency Capitalisation: 2SLS Results	68
A.5	Full Results – Unrealised Efficiency Capitalisation: 2SLS Results	69
3.1	Balancing Test at the Border of Efficient Clusters	98
3.2	Estimates of the Price Premium for Efficient Clusters	113
3.3	Placebo Test Estimates of the Price Premium for Efficient Clusters	115
3.4	Apartments: Estimates of the Price Premium for Efficient Clusters	117
3.5	2SLS Estimates of Exposure to Clusters on Property Prices	128
3.6	First Stage IV Estimates of Neighbourhood Effects in Improvements	131
B.1	Estimates of the Price Premium for the Redefined Efficient Clusters	138
B.2	Estimates of the Strength of Spatial Relationships on Prices	142
B.3	Balancing Test at the Border of Efficient Clusters: No Sale	156
B.4	Estimates of the Strength of Spatial Relationships on Prices	157
B.5	Estimates of the Price Premium for Efficient Clusters with Unrealised Ef- ficiency	158
B.6	Changes in Unrealised Efficiency on Prices – Whole Sample	159
4.1	Drivers of Investment Expenditure and Property Prices	175

1 Introduction

1.1 Motivation

As climate events become more frequent and severe, improving and adapting the residential stock should be high on the global agenda. Recent data out of the UK demonstrates the need for action, as built environments now account for 20% of national greenhouse gas emissions – 65% of which came from fuel combustion in residential settings alone (The Department of Business, Energy and Industrial Strategy, 2025). Mass retrofitting is required to lower the sector’s carbon footprint and the impact of future climate events. Improving the energy efficiency of our homes in the short-term is important to lower our use of polluting energy sources, but households also benefit from improved thermal comfort. Yet, research in this field finds that the uptake of residential energy investments are far lower than required and point to a number of reasons why: non-monetary costs arising from bureaucracy (Fowle et al., 2015); poor financial returns to investment (Allcott and Greenstone, 2012); and uncertainties around future energy prices and asymmetric information of energy savings (Jaffe and Stavins, 1994). To pin down the drivers of energy efficiency investments and understand why uptake is low is to solve the very important problem of excessive carbon emissions.

Poor uptake is not just a UK phenomenon. Estimates from the European Union find that only 25% of individuals aged over 16 live in properties that have undertaken energy efficiency improvements in the last five years – ranging from as high as 59% in the Netherlands to as low as 8% in Malta (Eurostat, 2024). High levels of inequality exist across the continent, which is problematic for all. Carbon emissions possess the properties of a non-excludable good, in that emissions created in one country are transferred to others via atmospheric conditions (in a similar way to that outlined in Heblich et al., 2021). Poor levels of residential retrofitting affect all countries simultaneously. Hitting the breaks on climate change cannot be achieved without consensus and collective action.

However, the UK faces a significant barrier in that it has some of the oldest housing stock

on average across Europe. Homes built prior to the end of WWII comprise 38% of the stock – the highest across Europe. Furthermore, a larger-than-average 15% of all homes fail to meet the decent homes standard (Eurostat, 2023). Older, poorer quality housing threatens to increase the costs of retrofitting across the board and pile further pressure on the returns to such improvements. In recognition of this, a number of schemes have been implemented by previous governments in an attempt to counteract poor financial returns. The Warm Front Scheme was the first of its kind. Running between 2000 to 2013, the scheme helped vulnerable and fuel-poor households to install energy efficiency measures such as heating and insulation by way of a government grant. During its operation, 2.3m households (11% of all households) received assistance ranging between £3,500-£6,000 in value (UK Parliament, 2013). The scheme was subsequently de-funded and replaced by the Green Deal which was provided by energy companies. Households were granted loans – repayable through their utility bills – to make energy efficiency improvements. The change in structure led to poor take up and the scheme was soon removed. Subsequent funding commitments have been made by ruling political parties with most efforts and finances decentralised to district governments. The drop-off in uptake with the introduction of private financing provides a stark message – that the returns to making energy efficiency improvements either do not exist, or are not common knowledge to households.

By combining robust empirical analysis with theoretical modelling, I aim to assess this landscape to understand whether households really do value energy efficiency improvements, what specifically they value, and whether this influences the uptake of improvements. Too much of the literature focuses on the bigger picture of capitalisation without much recognition of the mechanisms driving it. In what is to follow, I set out the institutional background guiding the residential energy efficiency framework in the UK. I then provide an overview of each of the three chapters.

1.2 Background

The barometer for energy efficiency in the UK comes in the form of Energy Performance Certificates (EPC). By law, properties that wish to sell on the market must be in pos-

session of a certificate that provides detailed information regarding a property's characteristics, including its structural quality, energy costs, energy consumption and energy efficiency. Energy efficiency scores are calculated in accordance with the Standard Assessment Procedure (SAP). The SAP works by estimating energy consumption – accounting for property characteristics like size, age and so on – from which a measure of energy cost is obtained. This measures the energy requirement per m^2 of floor area which is then charted on a scale from 0-100 to form a property's energy efficiency score, with zero being the lowest and one-hundred being the highest. Energy efficiency scores are therefore calculated on the basis of a property's energy requirement measured through its cost. Criticisms of this method arise from the fact that 'efficiency' in a tangible sense does not factor into the measure, that it's more a measure of energy usage. Controversy over the SAP has seen its revision over the years. Introduced in 1998, updates have taken place in 2001, 2005, 2009, 2012 and most recently in 2022. The introduction of the 2009 update coincided with new Energy Performance of Buildings regulation (2008) introduced in the UK that made it a legal requirement for households to acquire an up-to-date EPC prior to selling a property. This piece of legislation was required under EU law with the aim of increasing transparency in the selling process, as buyers are supplied with the certificate upon transacting. As a by-product, EPCs also create a signal of quality which was intended to incentivise households to make energy efficiency improvements. To comply with the requirements, households wishing to sell a property must be in receipt of an EPC that was issued no more than ten years age. Those not wishing to sell are not obliged to acquire an EPC. Importantly for this thesis, popular property search portals began displaying energy efficiency scores within property listing adverts in 2010 to provide buyers with up-to-date property information. The accessibility of this information makes it a well understood phenomenon for buyers across the UK.

1.3 Thesis Structure

The thesis contains two empirical studies presented in chapters two and three, and one theoretical study presented in chapter four. Both empirical chapters exploit property-level data for the period 2009-2019. I provide an outline and a summary of the findings from

each chapter below. Concluding remarks are presented in Chapter five.

1.3.1 Chapter 2: Pricing Potential: Understanding the Value of Unrealised Energy Efficiency

In this chapter I explore how energy efficiency improvements are capitalised into property prices by proposing a new approach. The literature predominantly estimates how changes in current energy efficiency scores are capitalised into property prices. I argue that this approach is restrictive and fails to capture heterogeneous capitalisation rates originating from differences in potential energy efficiency scores between properties. Energy efficiency in the present matters, but property bids on those with greater investment potential will likely differ from those with limited potential. I therefore propose a new measure which I term unrealised efficiency. This captures the difference between a property's current and potential energy efficiency score. Changes in the Standard Assessment Procedure (SAP) – used to calculate energy efficiency scores – in 2012 created a discontinuous increase in potential efficiency scores, creating a proportional increase in unrealised efficiency scores. I circumvent endogeneity issues arising from unobserved property quality by exploiting this change in an IV-RDD approach to estimate how unrealised efficiency is capitalised into property prices. Doing so captures only the change in unrealised efficiency arising from the change in SAP, not from unobserved aesthetic improvements. While the current capitalisation rate captures how buyers value improvements made by previous owners, the unrealised capitalisation rate captures how buyers value making their own improvements – providing an estimate of the value buyers place on future energy efficiency.

Three key findings arise from the chapter. Firstly, the capitalisation of unrealised efficiency improvements into prices is around 9x smaller in magnitude compared with current efficiency. The unrealised capitalisation rate is estimated at -0.15% compared with 1.3% for current efficiency; whereby the former represents a price discount following a 10% increase in the unrealised efficiency score. Buyers then pay a premium for improvements already installed, but do not discount their property bids at the same rate when required

to install the improvement themselves – pointing to an undervaluation of the benefits and under-adoption of retrofitting. Secondly, I show that price discounts are greater for more invasive improvements that require a household to adapt their daily lives during installation. Buyers dislike the interruption in lifestyle and require greater compensation to undertake the improvements. Such improvements are less likely to be undertaken. Finally, while potential efficiency is capitalised into prices, it is the distance from potential that really matters. Distance matters for reasons related to uncertainty. Being too close to achieving potential limits the financial returns to improvements, while being too far from potential increases the number of improvements required (producing a disincentive) and the uncertainty of financial return. Potential therefore matters, only so long as it is attainable.

1.3.2 Chapter 3: Neighbourhood Effects in the Market for Energy Efficient Housing

Building on the previous chapter, and given the limited spatial evidence within the field, I argue that heterogeneous capitalisation rates across space lead to the emergence of neighbourhood effects in home improvements as buyers chase returns on their investment. It seems reasonable to think that certain neighbourhoods deliver greater investment returns. I test two hypotheses as potential explanations: i) that neighbourhoods with a high proportion of energy-efficient properties signal greater potential returns to investment, and ii) that buyers with preferences for residential energy efficiency can observe these neighbourhoods and choose to select into them. I test these two hypotheses by leveraging data from property listing adverts across the London Travel To Work Area (TTWA). These data are used to identify property's listed on the market in any given year. Using these properties, I employ an algorithm to identify clusters of properties with high energy-efficiency scores which deviate from the average for the area. These property clusters are spatially bounded to create a border which I then exploit to estimate discontinuities in property prices. Property clusters are created by exploiting random variation in listings and are exogenous in size and location. I then exploit these in an IV strategy to estimate neighbourhood effects in energy efficiency improvements arising from the exposure

of incumbent households (existing residents) to buyers that choose to reside within these property clusters. I argue that it is exposure over time to these individuals that shape household preferences.

Findings from the chapter highlight a positive price discontinuity, suggesting i) that buyers are attentive to neighbourhood energy efficiency scores, and ii) that this premium ought reflect a perceived benefit to residing within such clusters. I show that this premium reflects a higher expected return on property investments in the area as buyers pay more for properties of poorer quality. This creates two mutually exclusive groups: those selecting into efficient clusters, and incumbent households. I then show that, while the property clusters themselves provide intrinsic value to selecting households, subsequent buyers are actually paying for the upgrades made by this group. In other words, returns to residence within efficient clusters becomes insignificant and are partialled out by the returns from improvements. Finally, I show that exposure to these property clusters over time and the churn of selecting households within them induces a greater uptake of energy efficiency improvement among incumbent households through neighbourhood effects. Retrofitting activity is encouraged by households with such preferences, is driven by reporting of energy efficiency characteristics within listings, and that the effects are very localised.

1.3.3 Chapter 4: Household Location and Investment Decisions Under Neighbourhood Effects: A Theoretical Approach

In this chapter I develop a simple theoretical model that explains the mechanisms identified in chapter three – specifically, the finding that buyers pay a premium to locate in neighbourhoods with higher expected returns to investment, and that exposure to this group increases the likelihood of wider home improvements. The model is comprised of two mutually elusive groups: movers, defined as those searching for property, and non-movers, which constitute the incumbent households introduced in the previous chapter. Movers gain utility from consuming a numeraire good, housing and investment services used to upgrade their property. Preferences follow a Stone-Geary functional form in that each

mover must obtain a minimum requirement of housing to survive. Movers also face an investment incentive determined by the energy efficiency and preferences of the wider neighbourhood which is allowed to differ by household on the basis of household's attentiveness to energy efficiency scores. I model this such that higher levels of energy efficiency crowd out investment in the most-efficient neighbourhoods owing to lesser investment potential. Given the preference structure, this term enters the demand functions of the consumption goods/services. This allows me to express property prices in terms of investment demand and the neighbourhood's energy efficiency (which acts as the neighbourhood effect facing movers from the previous chapter). I then turn my attention to modelling the response of incumbent households to more investment activity locally. Non-movers are exposed to investments made by movers (which acts as the neighbourhood effect facing non-movers). Exposure to this activity enters their utility function through a variable that captures the spatially weighted investment expenditure undertaken by movers, which was modelled previously. Preferences of non-movers follow a Cobb-Douglas utility function and range over consumption of a numeraire good and investment services (housing is already chosen).

Solving the model returns three interesting findings. Firstly, the investment demand of movers is greater in more energy-efficient neighbourhoods. Solving the model for the equilibrium consumption of investment services makes it a function of the neighbourhood effect that increases (decreases) the investment undertaken in the most-efficient (inefficient) neighbourhoods. The model enforces households to conform to the investment level of their chosen neighbourhood, but it is in the most-efficient neighbourhoods that investments beyond this level are undertaken. Secondly, property prices are shown to be a function of this neighbourhood effect and the household's demand for investment services. Households increase their property bids i) for the most efficient neighbourhoods and ii) when they demand investment services beyond those required to conform with their chosen neighbourhood. Greater demand for investment in these neighbourhoods comes from the fact that buyers chase higher returns, meaning they bid more in expectation of profiteering. Finally, I show that exposure to a greater number of energy efficiency improvements locally increases the demand for investment among non-movers. Similar

to before, this enters the model as a neighbourhood effect which dictates the demand for investment services after observing those consumed by movers.

Chapter 2: Pricing Potential: Understanding the Value of Unrealised Energy Efficiency

2.1 Introduction

In the pursuit of ambitious Net Zero targets, economies the world over are rethinking the way they produce and consume energy in efforts to reduce carbon emissions. The residential sector has an important role to play in these efforts. Recent estimates from the UK show that built environments account for 20% of total greenhouse gas emissions, from which 65% are produced in residential settings (The Department of Business, Energy and Industrial Strategy, 2025). This looks unlikely to fall in the medium-term as the shift to home-working following the Covid-19 pandemic has transferred energy-use from highly-retrofitted office settings to poorly insulated residential environments – increasing both the energy consumed and carbon emitted.¹ Increasing energy-use at home highlights the need to reignite the debate into housing retrofit and the challenges threatening its mass uptake. The UK government recognises the need for action, legislating stricter energy efficiency standards in the rental sector and rolling out subsidy programmes to incentivise home improvements.² While this reflects progress, other barriers to uptake remain unaddressed, such as underutilised financing options and uncertainty around the financial returns to investment. This leaves homeowners facing a trade-off: accumulate savings to protect against the rising cost of living (short-term), or invest in improving their property’s energy efficiency to reduce energy bills (long-term). Such a decision calls into question the myopia of households and whether or not they take measures to future-proof their homes.

Research into myopia and the value buyers place on future-proofing their homes is limited. Studies primarily focus on the capitalisation of current efficiency into property prices, capturing only how buyers value energy efficiency in the present and neglecting heterogeneity

¹For EPCs submitted in 2011, just 36% of commercial buildings had an energy rating between A+ and C. In 2021, this stood at 55% Department for Levelling Up et al., (2022).

²Legislation now requires newly rented properties to be rated at least a C by 2035, with schemes such as the Green Deal implemented to subsidise improvements The Department of Business and Strategy, (2020)

with regards to a property’s investment potential.³ This chapter proposes a new approach to measuring capitalisation by calculating the difference between a property’s current and potential energy efficiency score. This measure defines the energy efficiency that is attainable, but not yet realised by the household. I refer to this hereafter as *unrealised efficiency*. Exploiting unrealised efficiency allows me to explore how buyers value energy efficiency improvements in the future – providing direct comparison with how buyers value energy efficiency in the present. I deviate from the traditional repeated sales approach used in the literature because of its inability to control for the correlation between aesthetic property improvements, energy efficiency scores, and property price.⁴ Furthermore, the approach makes it difficult to identify exogenous variation in unrealised efficiency scores.⁵ I instead adopt a novel instrumental variable approach that exploits changes in the Standard Assessment Procedure (SAP) methodology. The SAP sets out the approach to calculating energy efficiency scores by converting information on energy usage into energy cost that translates into a measure of efficiency. Changes in the methodology in 2012 gave rise to a set of new improvements recommended by assessing engineers at the time of an EPC, which subsequently created a positive discontinuity in potential efficiency scores. Using a panel dataset on energy efficiency characteristics drawn from property transactions across the UK between 2009-2019, I instrument for unrealised efficiency using the change in SAP methodology. Identification comes from properties requesting EPCs either side of 2012, and hinges on two unique features of the data: i) there are no changes in current efficiency scores around the time, isolating change in unrealised efficiency to changes in potential alone, and ii) changes in potential efficiency scores after the SAP do not correspond to changes in unobserved property quality. Both features ensure that changes in potential efficiency are exogenous, creating a signal to the market that is unrelated to underlying changes in property quality. Estimates from this specification are supported by robustness checks and heterogeneity analyses that prescribe capitalisation to changes in the cost of implementing recommended improvements.

³See Fuerst et al., 2015; Aydin et al., 2020; Deng et al., 2012; Cajias and Piazzolo, 2012.

⁴For example, unobserved installations of carpets, interior doors, or flooring affect both a property’s value and its energy efficiency, leading to bias.

⁵Unrealised efficiency is a function of two endogenous variables: current and potential efficiency. In a traditional repeat sales approach, both would need instrumenting for.

The chapter makes two contributions to the literature – the first being the proposal of an alternative definition of capitalisation. A shortcoming of the existing literature is that it considers only how energy efficiency improvements are valued in the present. Doing so fails to explore the forward-looking nature of buyers, missing a potentially important explanation as to why energy efficiency investments are lagging behind the level required. This chapter fills the gap by estimating the capitalisation rate of both current and unrealised efficiency. The former can be interpreted as the price buyers pay for improvements already undertaken by the previous owner. The latter instead represents the price buyers pay for the improvements yet to be installed (and which need to be undertaken by the buyer). Comparing the two estimates, we see how buyers value the same improvements when they are already installed versus when they are yet to be installed. Evidence that buyers value the improvement less when it is not yet installed would point to myopia through a lower willingness to pay. Another drawback of the literature exploring current capitalisation is its failure to frame capitalisation in the context of a property’s potential. McCord et al., (2020) provide one of the few studies exploring the relationship between property prices and energy efficiency potential, uncovering positive price effects. The paper fails to account for a number of contextual factors, namely the distance a property lies from their potential. While potential is found to be capitalised, this should only be true if this potential is realistically attainable. Using the approach of unrealised efficiency captures this in a more robust way, accounting for how the incentives (or disincentives) to make improvements are affected by the distance one has to travel to attain potential. Using this definition, I also build on an existing literature suggesting that capitalisation differs with information about the present (Sejas-Portillo et al., 2020; Szumilo, 2020; Hilber et al., 2019) by suggesting that a proportion of this effect comes from the forward-looking nature of those undertaking improvements. This chapter argues that it is not potential itself that drives premiums, but one’s distance from it. I explore this later in a set of heterogeneity tests, whether buyers are disincentivised by high levels of unrealised efficiency.

I provide a methodological contribution to the literature in the form of a novel instru-

mental variable. The majority of studies employ repeat-sales or fixed effects approaches to deal with the endogeneity in efficiency scores, with the only 2SLS application coming from Aydin et al., (2020) in a study in the Netherlands. I construct an instrument which exploits changes in the SAP methodology in 2012 that gave rise to a new set of recommendations. These recommendations represent technological improvements on the former and which increase potential efficiency scores discontinuously. Exogeneity arises from the fact that recommendations are idiosyncratically determined by the assessing engineer, irrespective of household influence, and that this change affects all property types within the dataset. This allows me to identify the exogenous component of changes in unrealised efficiency between EPCs and estimate the true causal capitalisation rate. This chapter is the first in the literature to acknowledge changes in potential efficiency scores arising from the change in SAP. Results from previous studies conducted in the UK are called into question if one believes that potential efficiency affects current capitalisation. I use this unique setting to explore whether potential efficiency does in fact affect capitalisation, and whether we observe heterogeneous effects either side of the change in SAP that explain this capitalisation. A further benefit of the choice of instrument is that it relies on exogenous changes in potential efficiency, not current efficiency. The latter is prone to greater endogeneity given that improvements are household-specific – changing as a result of household decisions to make improvements. By exploiting changes in potential efficiency, which cannot be manipulated by households, I am able to identify changes in unrealised efficiency that are exogenous to households. I therefore provide a different empirical approach to estimating capitalisation that bypasses the endogeneity concerns that arise when using changes in current efficiency.

Baseline results comprise the estimates of two capitalisation rates: one following changes in current efficiency, and one following changes in unrealised efficiency. I adopt the standard fixed effects approach to replicate the results found in the literature for current capitalisation, using this as a benchmark for comparison. Doing so yields a capitalisation rate of 1.3%; this is the amount by which property prices increase following a 10% increase in current efficiency. This lies between the rate of 1% estimated in Fuerst et al.,

(2015) using a repeat sales approach, and the rate of 2.2% estimated in Aydin et al., (2020) using a 2SLS approach. The estimated premium reflects the value buyers place on the energy efficiency investments made by the previous owner. Using the 2SLS approach to instrument for changes in unrealised efficiency, I estimate a much smaller unrealised capitalisation rate of -0.15% – the amount by which prices fall following a 10% increase in unrealised efficiency. Estimates are robust to different specifications. Unrealised efficiency increases with an increase in potential efficiency scores following the SAP, with findings showing that buyers discount property bids in response. This discount reflects the value buyers place on making their own improvements following an increase in the investment required to attain potential. Comparing the two capitalisation rates, the current capitalisation rate is 9x larger in absolute terms, suggesting that buyers undervalue making their own improvements relative to purchasing those already made. A set of heterogeneity tests are undertaken to explore the drivers of these findings. Firstly, I find that the price discount following an increase in unrealised efficiency is fully offset by the cost of installing the recommended improvements. While buyers place little value on undertaking these improvements, findings suggest that they do accurately price them. Furthermore, this finding suggests that buyers actually overbid for properties with existing improvements. Secondly, I find that buyers need extra compensation to undertake improvements that are invasive (that require the household to adapt their daily lives while renovations occur) in the form of greater discounts. Finally, I show that households discount properties in the expectation of undertaking recommended improvements only when their potential efficiency score is attainable, or when there is the potential for return on investment; properties requiring too many improvements, or with little room for improvement, are not discounted.

From a policy perspective, if the UK is to reach Net Zero by 2050, the primary incentive of improvements should be to reduce unrealised efficiency – not to profiteer from improvements in the present. This first comes with understanding the behavioural components at play in the sector and how buyers respond to information presented differently. This motivates the chapter’s focus on parsing improvements that specifically improve energy

efficiency. A wealth of literature exists to define how buyers value housing, but not so much about how they value the energy-enhancing components. Comparing estimates under both definitions of capitalisation incorporates a behavioural component to ascertain how different information within the same EPC are considered. I therefore frame this chapter within this policy debate in order to speak to how behaviour is determined by reporting, and to speak to the challenges that arise from this.

The rest of this chapter is structured as follows. Section 2.2 reviews the relevant literature. Section 2.3 introduces data and the policy context surrounding the introduction of EPCs in the UK. The empirical strategy is presented alongside the results in Section 2.4. Concluding remarks and discussion around the policy implications of this chapter's findings can be found in Section 2.6.

2.2 Literature Review

The capitalisation literature is in general agreement that increases in energy efficiency result in property price appreciation. More debated is the magnitude of such increases and the best methods used to estimate them. I review the literature in this field, starting with studies exploring why underinvestment occurs and then moving to papers that explore the value buyers place on these investments.

2.2.1 Underinvestment in Residential Energy Efficiency

Research in this field began in the 1980s and introduced methods still used to this day. At the root of this research is the literature on the energy efficiency gap – the difference between the current energy efficiency and the socially optimal level of energy efficiency of homes. Research in this field explores human behaviour as an explanation for underinvestment. The seminal papers of Hausman, (1979) and Dubin and D. L. McFadden, (1984) give rise to the energy efficiency capitalisation literature with research on the energy efficiency gap. Their consumer choice models paved the way for explaining the derived demand conditions for energy. Both studies find that consumers with lower incomes tend to have higher discount rates for energy efficiency investments compared with

those on higher incomes. Hausman, (1979) estimates the discount rate to be on average 20%, while for those on the lowest incomes (\$6,000), this is around 89%. Lower earners are less likely to invest in energy efficiency even in the presence of rising energy costs due to credit constraints and the associated gambles around future payoff. This may be due to limited knowledge of the benefits of retrofit and high disruption levels during work (Fowlie et al., 2015). Energy demand is more elastic for those on lower incomes compared with wealthier households – meaning they are less likely to invest in retrofit, instead choosing to just reduce supply (Urban and Ščasný, 2012). The under-realisation of one’s potential efficiency gave birth to the notion of the energy efficiency gap: referring to the difference between the current and the potential energy efficiency of a property or population. Research has expanded on this to define it as a societal measure. Jaffe and Stavins, (1994) suggest defining the gap as the difference between the current and socially optimal rate of efficiency. But the efficiency gap itself is inherently subjective and difficult to estimate under this definition, as one person’s small gap may be another’s large gap. The definition above speaks to the current public efforts to reduce carbon emissions and achieve Net Zero, as energy efficiency improvement are thought of as socially-beneficial. In this sense, households must consider improvement as progress towards achieving their potential efficiency and not only as a means for capitalisation. However, efforts to close the gap do not ultimately translate into environmental gains. The presence of the “rebound effect” – higher derived demand for energy-based products resulting from energy savings – may actually reduce energy efficiency gains (see Fowlie et al., 2018; Dubin et al., 1986; Gillingham and Palmer, 2014). Fowlie et al., (2018) point to weak welfare gains as a result of the rebound effect as people increase their consumption of energy services. These gains may be positive from an economic and social perspective, but environmentally, they are defeatist with respect to CO₂ emissions. Gillingham and Palmer, (2014) argue that the rebound effect is actually welfare enhancing as the positive outweigh the negative welfare effects. Households achieve a better level of comfort in their home for the same cost as before. The topic of CO₂ emissions is explored in more detail later, as limited research is framed with the main goal of reducing such emissions.

2.2.2 Returns to Investment

The capitalisation literature builds on this research and explores the role of financial returns as incentives to close the gap. Research started in the 1980s and introduced models still applied widely today – like the hedonic pricing models in Rosen, (1974) and Roback, (1982). Publications at the time encountered a common drawback: poor data availability, both in terms of the number and quality of observations. Early papers in the literature employed 2SLS models to isolate exogenous improvements in energy efficiency and to overcome the limited data (Dinan and Miranowski, 1989; Johnson and Kaserman, 1983; Longstreth et al., 1984; Bound et al., 1995). These shortfalls resulted in limited research into capitalisation in the following years, with the exception of the paper by Nevin and Watson, (1998). They explore an important hypothesis of Longstreth et al., (1984): do market valuations for energy efficiency improvements follow the rational market hypothesis? The answer is a test of whether the price premiums following improvement outweigh the marginal cost (i.e energy cost savings vs higher mortgage cost). This explores the degree to which the benefits of energy efficiency improvement are discounted. The authors find evidence that energy efficiency improvements lead to positive price premiums: a circa \$20.73 increase for every \$1 reduction in energy expenditure. Improvements increase home value by an amount that reflects a discounted rate of energy savings in line with that of mortgage interest rates. This suggests that households can profit from improvements, regardless of time between the sale. The exploration of the timings of improvements versus the sale introduces an interesting dynamic – as little research exists into the financial viability of conducting improvements in properties with short occupation.

The availability of better data and more advanced econometric techniques reignited the debate around the capitalisation of housing energy efficiency investment. These papers can be categorised into two themes: research estimating capitalisation rates of energy efficiency into house prices, and research into the heterogeneity in these rates. New methods were also applied in the search for causal effect. The newer literature makes use of tested energy efficiency scores as opposed to proxies, and report capitalisation rates as elasticity's. In their widely cited paper, Fuerst et al., (2015) explore capitalisation with the use

of a cross-section hedonic pricing model and a panel data repeat sales model. The paper finds that a 1% increase in current energy efficiency is associated with a 0.1% increase in property prices. Included in estimation is a price trend that controls for regional and local market dynamics, as well as accounting for the time elapsed between sales. This captures well two things: the associated benefits of living (versus not living) in areas that values energy efficiency, and the occupancy for each property that determines the payback period. Secondary to this, the paper estimates the price premiums associated with achieving a higher energy rating. Compared to a hold-out category (band D), the authors find significant price premiums associated with the highest rated properties (5% for A- and B-rated) and price discounts for the lowest rated properties (7% for G-rated and 1% for E- and F-rated – dubbed by McCord et al., (2020) as "brown discounts"). What is surprising is the similarity in price premiums associated with the two lowest and highest bands. This finding is important because there appears little financial incentive to improve homes further in band B, and little incentive for homes in band F to make minor improvements. Aydin et al., (2020) extend on this paper by estimating the capitalisation rate with both a cross section and panel data using a 2SLS model. Their study uses data from the Netherlands between 2008-2011 in the former, and 2008-2015 for the panel data model. They find a 0.22% capitalisation rate following a 1% improvement in current energy efficiency. The model used here extends on Fuerst et al., (2015) by controlling for quality adjustments through the IV and a control variable accounting for property maintenance. Different to the previous paper is the smaller sample size (seven versus seventeen years). The smaller occupancy length in Aydin et al., (2020) likely biases results when not controlling for length, given that the improvements made may be representative of property investors, not long-term homeowners. This paper also extends the analysis by estimating the causal effect of moving up a rating bands. Using a Regression Discontinuity (RD) design, they find limited evidence of a price premium resulting from climbing bands. Although there is an associated premium, it is not caused by the crossing of a band. This lies in direct contrast with the paper by Sejas-Portillo et al., (2020) who find a premium of between 0.8%-2.5% from crossing a threshold. The authors find that the threshold effect is greatest for G-F rated properties (most inefficient) and lowest for C-D

rated properties (highest proportion of properties). This suggests that energy efficiency improvement is still capitalised at lower rating bands. They expand this analysis further to estimate whether rating bands reduce buyer attentiveness to energy efficiency scores. They find that attentiveness is reduced on average by 25%, 34% in E-F and 12.5% in C-D rated properties. The price premiums associated with higher-rated properties likely explains the lower degree of inattentiveness. One concern with the RD model comes from the assumption that properties are similar either side of the threshold, but the model does not control for the effect of postcode wards in urban settings. While properties may be similar in characteristics, prices are likely not. The model could be extended to account for this and illustrate the effect. These papers identify positive capitalisation rates following improvements, but highlight that these are heterogeneous. Fuerst et al., (2015) identify capitalisation, but do not offer causal estimates of the premiums associated with achieving higher ratings. There is likely some bias in their estimates, as they cannot distinguish between what constitutes an energy efficiency versus a cosmetic home improvement. As such, they have to remove a portion of their sample. My chapter extends on this by isolating the price impact of energy efficiency investment through 2SLS. I also test whether the estimates are robust to removing properties at the threshold of the next band – as these properties have greater incentive to invest, and benefit from higher capitalisation rates.

The paper by Sejas-Portillo et al., (2020) extends on an interesting – and under-researched – area around uptake of energy efficiency improvements. Moving above or towards the next rating-band requires incentive, and Comerford et al., (2018) explore this in their paper. The authors test the hypothesis that energy rating labels themselves act as incentives to improve by comparing clustering around the EPC rating D pre- and post-2008, when EPCs became compulsory for residential transactions. Their findings suggest that EPC requirements led to a clustering of households within rating band D, but find no evidence of further improvements beyond the initial. This creates a repeated incentive problem, whereby owners do not invest when far away from the next band. Similar to the study of rating bands and threshold effects, this suggests that people are inattentive

to the SAP scores. However, there is evidence that information provision is a key driver of energy efficiency improvements and that tailoring this provides incentive. This builds on a wealth of literature in behavioural economics around information absorption and heuristics (see Stigler, 1961; Pope, 2009; Lacetera et al., 2012). However, the authors point to a limitation of the paper – that the incentive to improve may be driven by the potential score. This cannot be ruled out in the paper, threatening the causal interpretation. My chapter extends this by looking at the impact of reporting potential energy efficiency scores on property prices. If indeed buyers respond to a property’s potential, then the current rating becomes less important to certain buyers.

2.2.3 Heterogeneity

The growing literature around heterogeneity is symptomatic of the need for better policy design. Low uptake – like is found in Comerford et al., (2018) – requires policymakers to better understand how buyers view energy efficiency differently. Two notable papers in this literature come from Hilber et al., (2019) and Szumilo, (2020). Both explore heterogeneity in capitalisation that comes from differences in neighbourhood characteristics. Hilber et al., (2019) explore the effects on energy consumption and energy efficiency installations of an increase in national energy prices. This is explored for homes in and outside of conservation areas (those areas with specific protections on buildings). The low improvement potential and high associated costs limit the incentives to improve energy efficiency. The authors find that households respond by making improvements and lowering their energy consumption, but this is significantly lower for households in conservation areas. This is an important extension of the literature as it accepts that previously explored incentives to make improvements do not always hold (information provision in Comerford et al., 2018; or threshold effects in Sejas-Portillo et al., 2020). Further research is needed to understand the role of market structure here. This is the topic of the paper by Szumilo, (2020). He explores the impact of market structure on property-price spillovers through peer effects following property renovation. Though not concerned with energy efficiency improvement per se (only the data from EPCs), the paper highlights the

interconnectedness between local property prices. The paper finds that property prices rise by 0.3% in response to a 1% increase by one's neighbours. Internal renovations only impact local prices when the property is listed on the market, creating an information channel, and not through intra-neighbourhood channels. This may take place through an *anchoring effect*, whereby residents change their pricing behaviour based on observed local prices and characteristics. Importantly, one cannot rule out the impact of the market in estimating capitalisation – and models should control for such characteristics. Conceivably, buyers with preferences for energy efficiency and environmental characteristics may be more likely to pay higher prices to locate near like-minded households – representing a spillover effect. This creates almost a free riding problem. These two papers represent interesting developments in the literature by exploring heterogeneity in external conditions, something that I expand on in another chapter.

Fuerst et al., (2020) and Kholodilin et al., (2017) explore heterogeneity further by considering the impact of energy efficiency improvements on the rental market. Both papers find a significant increase in rental prices following improvement, but the former paper suggests this is only small. The former applies a similar methodology to Fuerst et al., (2015) and find that properties rated B/C command a price premium of 3-5%, compared with D-rated properties. Yet, the monetary value of 5% higher rents is minimal. On the other hand, Kholodilin et al., (2017) find a higher associated premium in Berlin – around €23 for every €1 reduction in energy savings. This represents a much larger capitalisation rate. Both papers agree on the problem facing improvement uptake in the market – the landlord-tenant dilemma, also termed the split incentive problem. This posits that under-investment occurs due to market failures and products benefits that the landlord does not receive. For example, landlords do not benefit from energy savings and tenants have a low willingness-to-pay for energy efficiency investment given their limited stake in the property and (likely) short-term tenancies. This culminates in a discounting of the future; as Kholodilin et al estimate that future energy cost savings to exceed tenant's willingness to pay for investments by a factor of 2.5. This finding is corroborated by Cajas et al., (2019), who note that rental housing outweighs home ownership in Germany.

This may explain the larger capitalisation rate found in Kholodilin et al., (2017), but the authors here find a low rate. They suggest that energy efficiency improvement from rating band D-A only results in a rental premium of 1.4%. While this chapter does not explore the rental market, it is interesting to see that tenants value energy efficiency – albeit at a value less than the benefits. A flaw in the capitalisation literature is the sole consideration of the impact of improvement on prices. Little research exists on the intended benefits of improvement: the environmental impact. Clara et al., (2022) attempt to fill this gap. The authors focus on the rental market (contrasting this with the owner-occupied market) and the impact of Minimum Energy Efficiency Standards (MEES) on energy efficiency improvement. They find that the MEES regulation triggered energy efficiency investments in the private rented sector, but the environmental gains were minimal. This was driven by low-expenditure, high score-improving investments that may be considered more as cosmetic improvement. On the contrary, investments among homeowners translated into better environmental gains. This is driven by a split-incentive problem. Energy efficiency measures are calculated based on "cost to heat" type measures, but do not reflect environmental contribution. A caveat to the findings are that private-rented properties rely more heavily on electricity whose production is worse for the environment. A more econometrically robust estimation that controls for fuel-type may find that characteristics between rented/owned properties are not too dissimilar, and neither is environmental impact. This paper does not identify a causal effect on improvement of a change in the MEES regulation, but it does identify a spike. This spike is also present in owned properties, who may be incentivised to improve in anticipation of future changes in regulations. While not testing the environmental benefit of improvement directly, I frame my research in a way that stresses the salience of environmental impact. This chapter argues that energy efficiency improvement ought to be considered as improvement towards potential, and capitalisation rates ought to reflect this, not just as improvement in the current rating.

Another strand of the literature explores heterogeneity by property characteristics. One such characteristic is property prices. Papers by Evangelista et al., (2022) and McCord et al., (2020) test the degree to which energy efficiency ratings impact house prices. McCord

et al., (2020) look to test the impact of energy performance labelling on property prices using both current and potential energy ratings/scores (study conducted in Belfast). In their sample, the potential refers to the achieved score following uptake of cost-effective measures. The authors claim that this is a better information signal to buyers. The paper finds significant price premiums at only the highest quantiles and price discounts at the lowest quantiles when using energy efficiency scores. But the effects are driven by the highest and lowest rating band for changes in both current and potential energy ratings. It is hard to draw the causal link here between what is driving the price change: whether it is wealthier people's decisions or the property's characteristics. In the case of potential ratings, wealthier people may value highly-efficient properties, but they may also value poorly-efficient properties, driven by the vintage effect (see Hilber et al., 2019). However, this paper finds price discounts for low-rated, low-potential properties. Evangelista et al., (2022) extend on this by testing for energy efficiency valuation in the Lisbon market. With similar results, the authors find that higher-priced properties have higher capitalisation rates in comparison to lower-priced properties that are discounted. The studies differ in that this paper finds significant discounting. These appreciation/discounting effects are larger for existing homes compared with new builds, presumably because of the investment (dis)opportunities. One major drawback of the paper is that the authors use a sample covering the aftermath of the financial crisis, and allude to the more-than-proportionate impact on property prices among the poorer. This likely drives some of the result, given that energy efficiency is a luxury. Both papers fail to identify causality as to the drivers of this change. Difficulties also arise from the samples used. The former uses a small sample of just 1,476 properties, while the latter uses a sample ranging between the years 2009-2013. My chapter extends on McCord et al., (2020) by incorporating potential energy efficiency into estimation. This is a valued characteristic, but the authors do not interact both the current and potential together to identify their interdependence. The *unrealised efficiency* measure I use accounts for this and reflects the value that buyers place on closing the distance towards potential. I also conduct heterogeneity tests like in Evangelista et al., (2022) to identify the causal effect of improving energy efficiency in homes of different types.

To conclude, the bulk of the modern literature exploring heterogeneity in the market for energy efficiency recognises that there exist stark differences between household’s characteristics and incentives. The recognition of such heterogeneous responses is important in the formation of policy to roll out mass-retrofit, as there is no ”one-size-fits-all” response to this challenge. The coverage of the heterogeneity literature grows wider than the themes identified here. There are papers that explore the heterogeneity in energy efficiency uptake by housing regulation (Hilber et al., 2019), through peer effect and explicit vs non-explicit information provision (Szumilo, 2020), and by levels of risk aversion (Qiu et al., 2014), to name but a few. Future policy formation should reflect these differences and should account for multiple outcomes, like both environmental and price impacts.

2.3 Data

This chapter exploits detailed property-level data created by researchers from University College London (Chi et al., 2021) combining data from two sources: the Energy Performance of Buildings database (MHCLG) and the HM Land Registry Price Paid data for the period 2011-2019 for properties across England and Wales. The former reports information from EPCs on a property’s structural characteristics, including build type, year of construction, floor space (in m^2), number of rooms, and extension count. Moreover, the database reports energy characteristics, such as energy efficiency ratings and scores, estimated energy consumption and the associated costs.⁶ The Land Registry database on the other hand records transaction prices for all dwellings sold on the open market, providing the address, build type, land holding, and selling price.⁷ The resulting sample consists of over seven million properties and includes data on both energy efficiency characteristics and transaction prices, including the date when the property was inspected and sold. I use the longitudinal nature of the data to investigate changes in transaction prices resulting from EPCs using properties with multiple sales. Doing so allows me to isolate price fluctuations that can be attributed to changes in energy efficiency scores. By

⁶The property types included in the analysis: houses, bungalows, maisonette, and park homes; and the home build type: detached, semi-detached, (mid, end, enclosed mid- and enclosed end-) terrace. Flats are exuded from the analysis. The dataset does not record information on internal aesthetic quality.

⁷The dataset records the final transaction price, not the original listing price.

limiting the analysis to these repeated sales, the sample size is narrowed down to roughly 293,000 observations consisting of properties with two unique EPCs. I exploit changes in the Standard Assessment Procedure (SAP), the approach used to calculate energy efficiency scores, in 2012 that produced a discontinuous increase in potential energy efficiency scores to estimate the capitalisation of unrealised efficiency. In this section, I present the mechanisms driving this effect and motivate the empirical strategy with descriptive statistics.

2.3.1 Measuring Energy Efficiency

The SAP represents the methodology used to calculate residential energy efficiency scores. First established in 1993, the tool was originally developed to estimate residential energy consumption based on energy cost measures. By 2007, it became the cornerstone of EPC evaluations used to determine the extent to which homes were energy efficient. Obtaining an EPC involves scheduling a home visit with an assessing engineer who records features of the property (i.e. its structural characteristics and energy systems) to input into the SAP platform. From this, the engineer uses the information collected to develop two measures: i) energy demand, and ii) energy consumption for the household. The two differ in a subtle, but important way. Energy demand refers to the amount of energy needed for a task, while energy consumption calculates the *actual* energy consumed for the task.⁸ Properties with more efficient technologies reduce this discrepancy by reducing the amount of energy consumed. Using a set of standard assumptions with regards to the property's energy needs, information into the SAP is then used to estimate the annual energy consumption required to provide a standard level of heating, lighting and hot water relative to that demanded. An estimate of the total cost is produced, which will be greater for properties whose energy consumption grows relative to their demand. Costs estimated under the SAP are then normalised into a score that takes a value between 0-100, where larger scores correspond to lower energy costs. These scores represent the energy efficiency scores seen on EPCs.

⁸For example, a household wanting to heat a room with a boiler operating at only 80% efficiency would consume more energy than it demanded.

A caveat of the SAP approach is that properties may obtain a SAP score that exceeds 100. This occurs in instances when the household's net energy usage is negative, meaning they produce more than they consume – usually through renewable energy sources such as solar panels. These properties are few and are not of direct importance for this research; I therefore remove them from the analysis and bound energy efficiency scores between 0-100 both for current and potential efficiency. Potential efficiency then measures the lowest energy cost a property can obtain, based on the standard assumptions, after completing the improvements recommended by the assessing engineers. Given that I bound potential efficiency at 100, properties receiving high potential efficiency scores are those that are the best insulated and that require the least energy consumption relative to that demanded. In other words, higher rated properties are those that lose the least energy to waste arising from product inefficiencies. The same logic applies to the calculation of current energy efficiency scores, only these reflect the cost of energy consumed relative to that demanded in the present.

The SAP has undergone a number of changes, most recently in 2005, 2009, 2012, and 2022. Of interest to this chapter is the change in 2012 that gave rise to a new set of recommendations and re-weighting of existing calculations. I exploit this later as an instrumental variable that explains changes in unrealised efficiency.

2.3.2 Changes in the Standard Assessment Procedure in 2012

The updated SAP introduced a number of changes in energy efficiency calculations and its underlying assumptions, from changes in the assessed quality of structural materials to updates in the adjudicated importance of certain energy efficiency installations. A brief summary of those changes may be found below.


- Updating economic conditions: updates to fuel prices used in calculating energy costs
- Climate variation: factors regional climate into the measure to weight energy costs by local conditions
- Updated building specs: the effect of structural characteristics on energy consump-

tion re-weighted

These changes are tantamount to a re-weighting of energy scores arising from updated data on the cost of fuel inputs, incorporating variation in climate, and updating the efficacy of certain building codes with respect to their insulation capabilities. One would expect a re-weighting of the SAP components to affect both current and potential scores in relatively equal measure, as both calculations are subject to the same inputs. Figure 2.1 shows that this is not the case. The figure plots the average current and potential by EPC inspection year across the sample, illustrating the change in SAP in 2012. Potential efficiency scores experience a sharp increase around the time of the change, while current efficiency scores continue on the same trajectory. A-priori, it would be expected that both scores share a similar discontinuity. Given that they do not, this likely points to another explanation.

Potential efficiency is dictated by the improvements that *could* be undertaken, which differ on a property-by-property basis. Some properties have capabilities, structures, and limitations that others may not possess, and so will differ in the improvements available to them. It is plausible to expect that changes in the SAP may also bring about changes in the types of recommendations put forward.

Figure 2.1: Mean Energy Efficiency Score by EPC Inspection Year



NOTES: The figure shows the average current, potential and unrealised energy efficiency scores by EPC inspection year for all properties within the sample. The x-axis denotes the year in which the EPC inspection takes place.

I begin by exploring the types of recommendations put forward by assessing engineers both prior to, and after, 2012. These can be separated into three groups: with recommendations i) growing in frequency post-SAP change, ii) declining in frequency post-SAP change, and iii) with no notable change in frequency post-SAP change. Table 2.1 provides descriptive statistics on those growing and declining in frequency (referred to hereafter as growing and declining recommendations). I calculate the proportion of EPCs that contain each improvement prior to, and after, 2012. To constitute a growing or declining recommendation, we must observe a $\pm 30\%$ change in the frequency with which they arise in EPCs between periods.

Table 2.1: Recommendations by Category

Category	Improvement	Pre-2012	Post-2012 (%)	Change
Growing	Heat recovery system for mixer showers	0%	1.5%	>1000%
Growing	Floor insulation	0.006%	71.1%	>1000%
Growing	Flue gas heat recovery device in conjunction with boiler	0%	4.8%	>1000%
Growing	Flat roof insulation	0.001%	2%	>1000%
Growing	Draughtproof single-glazed windows	0.02%	10.4%	>1000%
Growing	Replace boiler with biomass boiler	0.02%	0.05%	131.7%
Growing	Cavity wall insulation	22%	31.8%	44.7%
Growing	Add additional 80 mm jacket to hot water cylinder	9.8%	14%	41.5%
Declining	Secondary glazing to single glazed windows	0.6%	0.1%	-84.6%
Declining	Replace heating unit with condensing unit	0.17%	0.1%	-40%
Declining	Fan assisted storage heaters	5.603%	3.56%	-36.4%
Declining	Replace single glazed windows with low-E double glazing	16.9%	10.8%	-36.2%
Declining	Change heating to gas condensing boiler	4.6%	2.9%	-35.7%
Declining	Replacement warm air unit	0.5%	0.33%	-33.3%

NOTES: This table shows the improvements that are increasing or decreasing in prevalence after the 2012 change in SAP. Pre-2012 and Post-2012 (%) refer to the proportion of EPCs each recommendation appears in over the period. Like in the empirical analysis, I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP.

The table shows that growing recommendations are predominantly related to insulation and heat recovery – improvements that would substantially affect the retention of heat and likely increase one’s potential by a large amount. In particular, measures that look to improve windows and better insulate the walls and flooring. Recommendations in this group appear more often on EPCs after 2012. This points towards a concerted effort to push these types of recommendations. On the other hand, declining recommendations appear to be related to older technologies, judging by their similarity in nature to those in the growing category (i.e. heat storage, window glazing, boiler replacements). It could be argued that the recommendations being pushed after 2012 were improvements on older technologies and represent technological innovations capable of significantly improving potential efficiency above-and-beyond the exiting technologies.

If the new recommendations that were growing in popularity after 2012 were technological improvements on their predecessors, then we would expect that they deliver greater energy efficiency improvements. Given that these are recommendations and are yet to be installed within the property, this should affect potential efficiency scores exclusively. I explore whether the presence of these recommendation did in fact cause potential efficiency

scores to increase after 2012. To do this, I split recommendations into those growing and declining in popularity post-2012 according to Table 2.1, as well as an *everything-else* category that captures the remaining recommendations. Using the universe of certificates recorded in the EPC register, I then undertake a series of regressions to estimate the effects on potential efficiency according to the following specifications

$$E_i^P = \alpha_k \mathbb{1}.\mathbb{D}_k + X_i' \alpha + \epsilon_i \quad (1)$$

$$E_{it}^P = \Phi \text{recs}_{it}^G + \theta \text{recs}_{it}^D + X_{it}' \alpha + \psi_i + \epsilon_{it} \quad (2)$$

$$E_{it}^P = \Phi \text{recs}_{it}^G \times \text{Post-SAP} + \theta \text{recs}_{it}^D \times \text{Post-SAP} + X_{it}' \alpha + \psi_i + \epsilon_{it} \quad (3)$$

In all specifications, the dependent variable is the potential efficiency score measured in levels. The first specification exploits a cross-section within the panel data structure. I estimate the difference in potential efficiency among properties that have recommendations with growing and declining frequency (captured in the dummy variable $\mathbb{D}_k$). I condition on a set of controls (X_i) to guard against omitted variable bias, including the EPC inspection year, total floor area, current efficiency, and a set of structural characteristics capturing age, build type, and property type. Estimates in α_k capture the differences in potential efficiency between properties with different types of recommendations. In the second specification, I exploit the panel structure to calculate the number of growing/declining recommendations included within each EPC. Including property fixed effects (ψ_i) to control for time-invariant characteristics, I identify within-property changes in potential efficiency arising from a change in the frequency of recommendations – under the a-priori that an extra growing recommendations should increase potential efficiency by more than an extra declining recommendation. In these specifications, controls are only included for the current efficiency score, total floor area, and a dummy indicator for the change in SAP. Finally, in specification three, I interact the number of growing and declining recommendations with a dummy variable signifying the change in SAP to see whether changes in potential efficiency scores are greater, or even heterogeneous, between the types of recommendations. If the discontinuity in potential efficiency scores truly comes from changes in the types of recommendations put forward following the change in SAP, then findings from this regression should show that properties with grow-

ing recommendations experience much greater increases in potential efficiency between EPCs compared with other properties after 2012. Results from all three specifications are presented in Table 2.2, with the marginal effects from the interaction terms in Eq.(3) presented in Table 2.3.

Table 2.2: Change in Potential by Recommendation Type

	(1) XS	(2) Panel	(3) Panel×Sap	(4) Panel×Sap
1.Growing	0.814*** (0.011)			
1.Declining	-2.317*** (0.030)			
Growing		3.824*** (0.018)	2.821*** (0.030)	3.700*** (0.018)
Declining		-2.511*** (0.052)	-2.564*** (0.052)	-3.345*** (0.053)
Growing×Post-SAP			1.298*** (0.019)	
Declining×Post-SAP				3.287*** (0.073)
Property FE		✓	✓	✓
R ²	0.693	0.700	0.702	0.704
Observations	3,410,871	733,398	733,398	733,398

NOTES: Standard errors clustered at the property-level in parentheses. Column 1 uses dummy variables to signify if a property possesses recommendations from the increasing/decreasing groups and uses all other improvements as the hold-out category. Columns 2-4 instead use the number of recommendations in each group for improved accuracy. Further controls are included for the current efficiency score and total floor area, while Column 1 is augmented further with controls for the property’s construction age, build type, and property type. This regression represents a linear specification, whereby controls are measured in levels. Full results are presented in Appendix A.1. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Columns 1-2 respond to Eqs.(1-2), while columns 3-4 correspond to the interaction terms presented in Eq.(3). Column 1 presents the results from the cross-sectional analysis, where recommendations from the “*everything-else*” group represent the hold-out category, and where we identify only whether or not a property receives a growing or declining recommendation. Estimates show that, relative to properties with recommendations that do not change in frequency, properties with declining recommendations have lower levels of potential.

Table 2.3: Change in Potential:
Marginal Effects

	(1) Growing	(2) Declining
Pre 2012	2.821*** (0.031)	-3.345*** (0.053)
Post 2012	4.120*** (0.019)	-0.058 (0.072)
Observations	733,398	733,398

NOTES: Standard errors clustered at the property-level in parentheses. The analysis takes the marginal effects from the estimates presented in Table 2.2; please refer to the notes for this table. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

On average, potential efficiency scores for these properties are 2.3 points lower. Properties with growing recommendations have relatively higher levels of potential efficiency, on the other hand. On average, potential efficiency scores for these properties are .8 points higher. These estimates only represent a static analysis of the differences in potential energy efficiency scores. Column 2 exploits the panel structure to identify how changes in the number of recommendations in the increasing and decreasing groups between EPCs change potential efficiency scores. Like in the previous column, the sample is restricted to properties requesting a new certificate either side of 2012. Estimates show that acquiring an extra recommendation that is growing in frequency on EPCs after 2012 increases potential efficiency scores by around 3.8 points on average. Conversely, acquiring an extra recommendation that is declining in frequency after 2012 decreases one's potential efficiency by about 2.5 points. These results point to differences in the efficacy of the recommendations in improving one's efficiency. Recommendations of increasing frequency deliver greater gains, providing support to the notion that these improvements represent advances in technologies. However, estimates from Column 2 do not explore the heterogeneous changes in potential efficiency following the change in SAP. I explore this in Columns 3 and 4 by interacting the number of recommendations in both groups with a dummy variable representing the change in SAP. Estimates from the interaction term in the former column show that acquiring an extra recommendation that is growing

in frequency after 2012 increases a property's potential by a greater amount following the change in SAP. The same is true of acquiring an extra recommendation declining in popularity after the change in SAP – though, the marginal effects present this more clearly. The marginal effects in Table 2.3 illustrate the effects on potential efficiency scores of an extra recommendation prior to, and after, 2012. The marginal effect of an extra growing recommendation prior to 2012 is an increase in potential of around 2.8 points, whereas after the change in SAP this increases to 4.1 points. Increasing by nearly double, we see that the different recommendations put forward after the change in SAP led to larger gains in potential efficiency. On the other hand, the marginal effect of an extra declining recommendation is negative in both periods. Prior to the change in SAP, an extra recommendation would reduce potential efficiency scores by around -3.3 points. After the change in SAP, the change effectively becomes zero and statistically insignificant. One would expect this result if these recommendations were believed to represent older technologies, as properties do not unlock extra potential from installing them. Evidence presented in this table suggests that the recommendations growing in popularity after the change in SAP represent technological advances, and that these advances lead to increases in potential efficiency scores over-and-above those of properties without access to said advancements.

Results from Tables 2.2 and 2.3 show that the average property experiences a significant increase in potential efficiency scores following an increase in the number of growing recommendations available to them. Conversely, the average property with declining recommendations experiences a decrease in potential efficiency scores. Estimates show that the SAP change delivered larger gains in potential to those properties receiving growing recommendations – owing to the fact that these recommendations represent technological innovations. However, Figure 2.2 illustrates the stagnant nature of current efficiency scores, highlighting that the change in SAP produced an isolated increase in potential efficiency scores alone.


This chapter exploits variation in energy efficiency scores affecting a property's unrealised

efficiency, defined as the difference between a property's potential and current energy efficiency scores

$$E_{it}^U = E_{it}^P - E_{it}^C \quad (4)$$

Unrealised efficiency defines the improvement required for the property to attain its potential – that being the ultimate goal of the EPC initiative. Variation is expected to arise predominantly from households undertaking home improvements that increase their current efficiency. However, given that the sample used in this chapter coincides with a change in the SAP methodology, it seems reasonable to assume that energy efficiency scores might also change as a result. I begin by plotting the average energy efficiency scores by years of the sample in Figure 2.2. Evident from the graph is the spike in potential efficiency in 2012 at the time of the change in SAP. Potential energy efficiency score experiences a structural break, with the average score increasing by more than ten points. Current efficiency remains relatively constant, increasing slowly through households improving their energy efficiency. Unrealised efficiency increases as a result of the increase in potential, providing reasonable variation through which to explore any resulting price effects. The sharp discontinuity in unrealised should act as an information shock to future buyers that a property i) has greater investment potential, but also ii) that this investment requires greater financing; both of which should have opposing effects on price. Identification in this chapter rests on the differential impacts from changes to the SAP on current and potential efficiency scores. Identification becomes cleaner as unrealised efficiency is the product of changes in potential efficiency alone around 2012.

Figure 2.2: Mean Energy Efficiency Score by EPC Inspection Year



NOTES: The figure shows the average current, potential and unrealised energy efficiency scores by EPC inspection year for all properties within the sample. The x-axis denotes the year in which the EPC inspection takes place.

Table 2.4 explores this change in efficiency scores between EPCs for three sub-samples of the dataset: for i) households that request both EPCs prior to the 2012 change in SAP, ii) households that request EPCs either side of 2012, and iii) households that request both EPCs after 2012. Reading along the row for current efficiency we notice that the average change in current efficiency scores remains relatively constant between EPCs across all three sub-samples. In-fact, households whose EPC request coincides with the change in SAP obtain smaller gains in current efficiency, suggesting that the SAP change did not discontinuously re-weight efficiency scores in the present. For potential efficiency scores, we notice that the change between EPCs for households requesting both before or after 2012 are not too dissimilar, with changes of 4.2 and 2.7 points respectively. These changes are small and could be attributed to measurement error or idiosyncratic variation in the assessing engineers' assumptions. Households requesting EPCs either side of 2012 experience a much larger change. The average household experiences a 16.2 point increase

in their potential efficiency score – an increase that far outweighs any change in current efficiency and one that would be unlikely to arise from a re-weighting of components within the SAP.

Table 2.4: Average Change Between EPC Inspections

Period	Pre-2012	Pre- & Post-2012	Post-2012
Current	10.7	8.2	11.1
Potential	4.2	16.2	2.7


NOTES: The figures represent the unconditional means of the change in current and potential scores between different subsamples of properties: properties whose both EPCs were requested prior to 2012 (Pre-2012), properties whose both EPCs were requested either side of 2012 (Pre- & Post-2012), and properties whose both EPCs were requested after to 2012 (Post-2012). Like in the empirical analysis later, I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP.

A discontinuity of this nature arising from the change in SAP would likely impact the retrospective distribution of unrealised efficiency scores. Figure 2.3 illustrates the distribution of current and potential efficiency scores either side of the change in SAP. I compare two samples: households requesting a new EPC prior to 2012, and households requesting a new EPC after 2012. Plotting the distribution of current efficiency scores shows no discernable differences between households requesting EPCs either side of 2012. The distribution is similar in shape, but for a larger proportion of properties obtaining a rating of D after 2012. Such differences could be attributed to households making energy efficiency improvements, but the change in SAP does not appear to generate any significant effects. Changes in potential energy efficiency are more profound. Households requesting a new EPC after 2012 have much higher potential, highlighted by the rightward shift in the distribution. The distribution peaked at a rating of C prior to 2012 compared with a rating of B after 2012. This corroborates the findings at the mean from Figure 2.2, suggesting that the effects of the change in SAP affected the underlying distribution of potential efficiency scores. Such widespread effects provides a good degree of variation in unrealised efficiency.

A concern across the literature is that properties bunch around energy rating bands (depicted by the red dashed lines); Sejas-Portillo et al., (2020) most recently observe this,

finding this to create price premiums for properties in the next rating band. Bunching is potentially problematic for empirical studies given the endogeneity in household decisions to obtain and remain at efficiency scores close to certain rating bands. In other words, household decisions to make improvements are likely correlated with unobservable factors, which further endogenises efficiency scores. Observing the distributions shows that bunching is not particularly prevalent in either score, from which I conclude that households are not upgrading for the purposes of just attaining the next rating band.⁹

Figure 2.3: Continuous Energy Ratings Distribution




NOTES: This figure shows the distribution of current and potential efficiency score for all properties in the sample for the period prior to the change in SAP (Pre-2012) and for the succeeding period (Post-2012). Like in the empirical analysis later, I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP.

A similar degree of variation in unrealised efficiency is also evident when observing the effect of the change in SAP on energy ratings. Figure 2.4 plots all unique combinations of current-potential ratings within the sample for properties requesting an EPC both prior to, and after, 2012.

⁹Bunching in this respect would threaten the validity of the results.

Figure 2.4: Energy Ratings Distribution




NOTES: The figure highlights the proportion of properties within the sample that fall into each current-potential rating band pair. This represents all properties within the repeat sales framework, *not* just the latest certificate. Like in the empirical analysis later, I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP.

Each cell contains the percentage of properties in the sample corresponding to that current-potential rating combination. Properties in the left-hand figure are those requesting a new EPC prior to 2012, from which we observe a large spread in the ratings. The majority of households possess a potential rating within the same band as their current rating, or in the band above. Large shares of households along the lead-diagonal (low unrealised efficiency) represents limited ability for properties to make significant energy efficiency improvements. We also observe a non-negligible proportion of properties that have potential ratings in bands E and F. In the post-2012 sample, following the change in SAP, we observe fewer properties in these bands, as the distribution of potential efficiency scores shifts upwards. A greater proportion of properties obtain a new level of unrealised efficiency spanning two and three rating bands, providing a greater amount of variation to be exploited. Econometrically, this is also an important source of variation as energy ratings provide a stronger, more noticeable signal of a property's quality at the time of purchasing – thereby influencing buyer's price offers.

Finally, I explore how the distribution of unrealised efficiency scores change between EPCs for two sub-samples of the data: for properties who request an EPC either side of 2012 (SAP Sample), and for properties requesting both EPCs before, or after, 2012 (Non-SAP Sample). If the change in SAP had disproportionate effects on energy efficiency scores,

the former would be expected to observe a greater change in unrealised efficiency. Plotting the result in Figure 2.5 shows this to be the case.

Figure 2.5: Energy Ratings Distribution



NOTES: This figure shows the distribution of the change in unrealised efficiency between EPCs for all properties in the sample. Negative (positive) values mean that unrealised efficiency has decreased (increased). Like in the empirical analysis later, I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP.

Households that request new EPCs either side of the 2012 change in SAP exhibit greater changes in unrealised efficiency relative to households that did not. A greater proportion of those in the Non-SAP sample experience a reduction in their unrealised efficiency, denoting that they are closer to achieving their potential. This is likely due to increases in current efficiency achieved by undertaking energy efficiency improvements. The distribution shifts rightward for the SAP sample as households experience smaller reductions in unrealised efficiency or larger increases. In the majority of cases, this increase in unrealised efficiency is greater than 10 points. What is illustrated in this figure matches the descriptives presented beforehand, that the increase in potential efficiency scores following the change in SAP brought about a widening of the gap between current and potential

efficiency scores which ultimately affected the signals to the market.

I present summary statistics in Table 2.5 firstly for the whole sample, but also splitting the sample between households requesting EPCs before and after 2012. Columns 2 and 3 are of particular interest given the clear distinction between EPCs either side of the change in SAP. For EPCs after 2012, I restrict the time window to four years, meaning I only capture properties that upgrade up to and including 2015. The table shows that the average adjusted price, deflated to 2009 prices and accounting for time-invariant property characteristics, differs between the two samples by roughly £10,000 in favour of the pre-2012 sample. The lower average transaction price following the change in the SAP suggests the presence of a price discount attributable to the change in SAP methodology. Estimates of the post-2012 adjusted price are accompanied by considerably larger standard deviations, suggesting that some properties were gaining significantly more than others. This could be explained by the growth in potential efficiency between the two sample. Potential efficiency scores are much larger after 2012, growing by an average of 15 points. Changes in current efficiency between the two samples are minimal, as expected from the evidence in the descriptive statistics. Together, this provides the variation outlined in the previous subsection whereby unrealised efficiency grew between samples at an average of 12 points – with a much larger spread across the distribution. Growth in unrealised efficiency could produce an opposing effect on prices, whereby buyers discount their property bids in response to financing extra energy efficiency improvements. That unrealised efficiency grew at such a rate makes me confident that I can use this variation in the estimation of price effect, which have been shown here to be ambiguous.

Reassuringly for the empirical strategy, there appears to be no economically significant differences in property structural characteristics between samples with respect to the size and number of rooms. I am confident that differences between the samples do not arise as a result of changes in the SAP. Furthermore, we see no differences in the energy cost reported in the EPCs between samples. Energy cost is reported separately by utility to which I sum the cost of heating, hot water and lighting. Limited differences at the

mean provides further evidence that the change in SAP had isolated effects on potential efficiency scores – otherwise we would observe changes in current efficiency and the benefit of lower energy costs.

Table 2.5: Summary Statistics

	(1) All Observations	(2) EPCs Before 2012	(3) EPCs After 2012
Price	236,379 (159,034)	215,181 (135,419)	241,856 (164,727)
Adjusted Price	291,426 (157,627)	301,948 (135,615)	282,738 (164,715)
Current	58 (13.2)	56 (14.7)	59 (12.7)
Potential	75 (11.2)	66 (11.9)	81 (8.0)
Unrealised	17 (11.3)	10 (7.8)	22 (12.0)
Floor Area	91 (41.6)	91 (40.0)	92 (39.0)
No. rooms	4.6 (1.5)	4.6 (1.6)	4.6 (1.4)
Energy cost	1,000 (461)	994 (481)	985 (454)
Observations	293,584	92,352	91,288

NOTES: Standard deviations in parentheses. The table presents the unconditional means after removing the top and bottom percentile of properties along the price distribution, and using all properties in the sample. This serves to remove the impact of outlier observations on capitalisation rates. Floor areas is measured in m² and energy cost is measured as the GBP value of hot water, lighting, and heating utilities. Like in the empirical analysis later, I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP.

2.4 Empirical Strategy and Results

The main focus of this chapter is to explore the effect of an increase in unrealised efficiency on property prices, under the a-priori that this leads prospective buyers to discount their bids in anticipation of financing greater improvements. To benchmark these results, I supplement the 2SLS results with a set of results from a standard fixed effects model. I further run the same fixed effects model to explore the capitalisation of current efficiency into property prices, which serves as a grounds for further comparison. I conclude the

section by undertaking a set of heterogeneity analyses.

2.4.1 A Fixed Effect Approach to Estimating Current and Unrealised Capitalisation

I begin by estimating the capitalisation of an increase in both current and unrealised efficiency scores into property prices. Previous studies have only estimated the former, using a variety of methods – the most popular being fixed effects and first differencing (see Fuerst et al., (2015); Fuerst et al., (2016); Cajias and Piazzolo, (2013)). Given the balanced panel adopted in this chapter, the two methods are equivalent. Researcher justify adopting a fixed effects approach by arguing that the model differences out common, time-invariant factors that may bias the estimation. Such factors include property construction age and build type. Furthermore, it is assumed that this differences out other factors impacting both property prices and energy efficiency scores, such as cosmetic home improvements (for example, the installation of new doors, carpets, flooring, or upgrading walls). These improvements are not observable within traditional datasets and will likely add value to a property whilst also increasing energy efficiency. Inappropriate controls for these features will ultimately introduce bias. Under the assumption that households do not make further home improvements between EPC inspections, this approach identifies the causal effect of an increase in current and unrealised efficiency scores on property prices. I estimate the model using the following specification

$$P_{it} = \beta_0 + \beta_1 E_{it}^C + \beta_2 E_{it}^U + \beta_3 \text{tfa}_{it} + \phi_i + \lambda_t + \Gamma_{kt} + \epsilon_{it} \quad (5)$$

After log-transforming the variables, I regress the log of property price on the log of current efficiency scores (E_{it}^C). The coefficient β_1 captures the impact of an increase in the current efficiency score between certificates on the property price. Estimates of capitalisation are conditional on a property-level control for the total floor area (tfa_{it}) which captures changes in the size of the property between sales (this is preferred to controlling for the number of rooms, which is often suppressed/not captured within the data). I further condition on a number of fixed effects across two different specifications. In a first specification, I control for property and year fixed effects (ϕ_i and λ_t , respectively). Captured by

the property fixed effects are those time-invariant factors affecting both price and energy efficiency, such as construction age and build type. It further provides a granular control for the property's location. Including year fixed effects controls for factors affecting the decision to sell within a given year, such as favourable market conditions or news/market shocks. In a second specification, I remove year fixed effects in favour of $\text{Location} \times \text{Year}$ fixed effects. Research in the field of urban economics recognises that neighbourhoods gentrify over time as richer households enter neighbourhoods in search of better amenities (Guerrieri et al., 2013; Couture et al., 2019). A similar trend may be present here in the form of home improvements in wealthier neighbourhoods. To the best of my knowledge, this has not been included in previous research, nor has it been acknowledged as a requirement. Including the dynamic fixed effect captures the time-variant factors affecting a location's attractiveness which may derive from a greater number of home improvements locally, and whose effects may be akin to those of gentrification. Such effects make areas more appealing for buyers, which increase local prices. Controlling for location dynamics seeks to eliminate the bias that may arise from such effects.

The regression in Eq.(5) represents a standard fixed effects model used within the literature, but for the inclusion of unrealised efficiency (E_{it}^U), which I omit in different specifications to produce replicable results. Identification of the effect arises from exploiting within-property variation in energy efficiency scores while holding other characteristics of the property constant. Though, this assumes that the unobservable, underlying structural characteristics of the property do not change between sales. If these characteristics do not change, one can assume that the property in the second sale is cosmetically and structurally similar to the first sale, meaning the only change is in the energy efficiency score. One can then attribute the causal effect of a change in current and unrealised efficiency on prices to the coefficients β_1 and β_2 , respectively. To date, the literature has not provided a credible approach to account for a property's investment potential. While two properties may have the same current, they may very different potential efficiency, and so buyers interested in adding value may pay more to secure the property with greater potential. In this case, higher price originate from the potential and not

the current efficiency, leading to bias if this is inadequately controlled for. I therefore augment the two specifications with a control for unrealised efficiency to capture a property's investment potential and comment on the effects of the resulting capitalisation rate.

Estimates from the fixed effects specifications are presented in Table 2.6. I split the estimates by specification, whereby Panel A represents the model inclusive of property and year fixed effects, while Panel B is inclusive of property and Location $\times$ Year fixed effects. Column 1 presents results from a regression of property price on the change in current efficiency scores alone. Inclusive of the respective set of fixed effects, estimates of the capitalisation rate between the two specifications are identical, suggesting that a 10% increase in current efficiency scores between sales generates a 1.7% increase in property prices. I recreate the same specification in Column 2 but for unrealised efficiency scores. Estimates of the capitalisation rate are also similar between specifications, ranging between -0.54% and -0.51%. However, they differ in that the coefficients are negative, suggesting that buyers discount their property bids in response to an increase (of 10%) in unrealised efficiency. As the property requires greater financing to achieve its potential, price falls. Column 3 merges the two specifications together without time-variant controls for property quality. Both capitalisation rates fall. Current capitalisation falls to around 1.5%, while unrealised capitalisation falls to around -0.3%. Upon including controls for the property's total floor area in Column 3, I arrive at the specification most closely aligned with the literature. Estimates of the capitalisation rates are most stable here between specifications, with a current capitalisation rate of 1.35% and an unrealised capitalisation rate of -0.25%. Unrealised capitalisation is more than 5x smaller in absolute terms than the current capitalisation rate. The two differ subtly in what they measure: current capitalisation captures how buyers value the improvements installed by previous owners, whereas unrealised capitalisation captures how buyers value the same improvements when they are yet to be installed. Results from this exercise show that buyers place a much lower value on the same improvements when they themselves are required to install them. This points to a central issue in the literature whereby households do not like making their own energy efficiency improvements, pointing to issues around the low

uptake of retrofitting activities.

Results from Table 2.6 represent the first estimates of unrealised capitalisation, capturing how buyers value energy efficiency gains in the future. This is significantly lower than the estimate of current capitalisation, which is similar to that of the literature and represent a midpoint of that estimated in two leading studies: Fuerst et al., (2015) estimate a capitalisation rate of 1% using a similar fixed effects approach, while Aydin et al., (2020) estimate the rate at 2.2% using an instrumental variable approach. That the result aligns with previous studies is reassuring and provides this chapter with a good comparator when estimating capitalisation from unrealised efficiency.

An issue with the fixed effect approach is that one must make a number of assumptions – mainly, that property quality remains constant. If this were true, then the property fixed effect would perfectly capture property quality. However, it is much more realistic to assume that buyers make other property improvements between EPCs that affect the cosmetic appearance of the property. This introduces endogeneity into the estimation, as cosmetic improvements are likely capitalised into property prices whilst inherently contributing to the property’s energy efficiency score. Changes such as these are unobservable in traditional datasets and cannot be adequately controlled for, making it difficult to identify causality within a traditional two-way fixed effects approach. I instead opt for a 2SLS to estimate the unrealised capitalisation rate that I then compare with estimates for current capitalisation presented above according to the standard literature definition.

Table 2.6: Fixed Effects: Current Capitalisation

	(1)	(2)	(3)	(4)
Panel A: Baseline Fixed Effects				
Current	0.172*** (0.003)		0.154*** (0.003)	0.134*** (0.003)
Unrealised		-0.054*** (0.001)	-0.030*** (0.001)	-0.026*** (0.001)
Floor Area				0.261*** (0.006)
R ²	0.621	0.592	0.629	0.658
Observations	98,262	98,262	98,262	98,262
Panel B: Location×Year Fixed Effects				
Current	0.168*** (0.003)		0.150*** (0.003)	0.133*** (0.003)
Unrealised		-0.051*** (0.001)	-0.028*** (0.001)	-0.024*** (0.001)
Floor Area				0.241*** (0.006)
R ²	0.986	0.985	0.987	0.988
Observations	98,262	98,262	98,262	98,262

NOTES: Standard errors clustered at the property level in parentheses. The baseline model in Panel A includes property and year fixed effects, while Panel B includes property and Location×Year fixed effects. All terms are log-transformed, with both sets of regressions including a further control for property floor area. I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

2.4.2 A Two Stage Least Squares Approach to Estimating Unrealised Capitalisation

I now set out the approach to estimating the true causal effect of an increase in unrealised efficiency on property prices. I adopt a Two Stage Least Squares approach that accounts for the endogeneity present in unrealised efficiency scores (detailed in the above sections). The exogenous component of the change in unrealised efficiency between EPCs is obtained by exploiting the exogenous change in SAP methodology and the differential effects this had on current and potential efficiency scores.

Identification of Changes in Unrealised Efficiency As previously mentioned, fixed

effect models suffer from an inability to control for the time-variant improvements made between sales – especially those that impact on both a property’s value and its energy efficiency. Without sufficient data on said improvements, a fixed effect model will always be an inadequate causal estimator. A further challenge in obtaining causality arises from the way in which unrealised efficiency is measured. We know from Eq.(4) that unrealised efficiency is a function of two endogenous variables, current and potential efficiency scores. Traditional approaches to isolating exogenous variation in unrealised efficiency would involve finding exogenous treatments for both variables. To circumvent this problem, I exploit the change in SAP methodology in 2012 as an exogenous shock that had differential impacts on both scores. Figure 2.2 illustrates the sudden increase in potential efficiency scores, coupled with stagnation in current efficiency scores. Identification of unrealised efficiency arises exclusively from changes in potential efficiency – which I show in the previous section to arise from technological advancements in energy efficiency measures. We can therefore express changes in unrealised efficiency purely as a function of changes in potential $E_{it}^U : f_n(E_{it}^P, E_{it}^C) \rightarrow f_n(E_{it}^P)$. This is completely exogenous to the household and is determined the the assessing engineer.

By isolating changes in current efficiency scores, I exploit the panel structure of the dataset to write post-2012 potential efficiency scores as

$$E_{i2}^P = E_{i1}^P + U_{2012}$$

For properties requesting an EPC either side of 2012, potential efficiency in the second sale is equal to that of the first sale plus any change arising from the change in SAP (U_{2012}) – which serves as the exogenous shock. Rearrange this shows that the change in SAP should be equal to the difference in potential efficiency scores between the two EPCs. The change in SAP should then account for the full change in potential efficiency between the two certificates.

$$\begin{aligned} U_{2012} &= E_{i2}^P - E_{i1}^P \\ U_{2012} &= \Delta E_i^P \end{aligned} \tag{6}$$

I corroborate this feature of the shock with Figure 2.2 that shows the change in unrealised efficiency is practically equivalent to the change in potential efficiency. Taking Eq.(6), we can express this as $U_{2012} = \Delta E_i^U$. This is important for identification, as changes in potential efficiency are exogenous to the household and cannot be manipulated, and therefore should not be correlated with household- or property-level unobservables. Identification hinges on the change in potential arising from the change in SAP, whilst being uncorrelated with changes in property fundamentals. This precludes within-property variation in property quality and limits any confounding effects. It is therefore plausible that the change in SAP is an exogenous event and can be used to instrument for unrealised efficiency scores.

The mechanics presented above represent a stylised view of the effects driving changes in potential efficiency scores. Though the instrument is unlikely to be correlated with unobservables, concerns that it may fail to capture the entire change in potential efficiency remain. Changes could arise through factors such as current efficiency, structural characteristics, or even measurement error. If this were the case, we could rewrite the identity as

$$E_{i2}^P = E_{i1}^P + U_{2012} + \varepsilon_i$$

Where ε_i captures other factors contributing to the change in potential efficiency. Rearranging the equation above shows that changes in potential efficiency scores arising from the SAP change could be confounded by this error

$$U_{2012} + \varepsilon_i = \Delta E_i^P$$

$$U_{2012} = \Delta E_i^P - \varepsilon_i$$

I use this unique setting to test for the presence of error in the estimation. Specifically, I estimate a separate 2SLS model with a different instrumental variable to explore the sensitivity of the estimates. Where U_{2012} represents the initial instrument, I interact this with another exogenous variable θ to obtain a new instrument θU_{2012} ; doing so returns

the following identity

$$\theta U_{2012} = \theta \Delta E_i^P - \theta \varepsilon_i$$

This approach tests whether the change in unrealised efficiency following the SAP is confounded by other factors, a feature that would then change the capitalisation estimates. If it were not confounded, one would expect no error $\varepsilon_i = 0$, meaning we would obtain the same results as if we estimated the original equation in Eq.(6)

$$\theta U_{2012} = \theta \Delta E_i^P$$

$$U_{2012} = \Delta E_i^P$$

One would expect the results between specifications to change very little, showing the change in SAP to be a strong instrument for identifying changes in unrealised efficiency.

The observed changes in potential efficiency following the change in SAP should be exogenous, with variation in magnitudes across properties being as good as random. I show in Appendix A.3 that this change in potential occurs across all property types with similar magnitudes at the mean, suggesting that the treatment is not exclusively determined by particular property characteristics. Furthermore, flats also experience an increase in potential efficiency, irrespective of their limited scope for energy efficiency improvement, meaning that the change in recommendations were not targetting properties with any underlying potential for improvement. Identification comes from properties that request EPCs either side of 2012. Using within-property variation around the SAP, I identify the effect of changes in unrealised efficiency on property prices against the counterfactual that unrealised efficiency remained constant. Two features of the empirical strategy allow me to achieve this. First, there is little change in current efficiency scores, meaning that the property’s characteristics in the present-day do not change. Given this, the property’s quality and efficiency should not change between EPCs either-side of 2012, conditional on the property not making their own improvements.¹⁰ Secondly, there should be no

¹⁰Controls for current efficiency scores within the regressions account for this, as energy efficiency improvements are unobserved within the dataset.

changes in a property's structure or underlying potential between certificates, given that changes in potential efficiency are driven by changes in recommendations pertaining to SAP; the instrument instead identifies a change in the *signalling* of potential between certificates. Using this empirical strategy in this context should ensure that the only difference between a property at two different EPC inspections is related to changes from the SAP. The instrument is therefore driven by exogeneity, firstly, in the timing and form of technological innovations. There are no signs that bunching occurs, whereby a mass of households flock to certify their property after 2012 (see Appendix A.2). Bunching would only be expected in the event that households could pre-empt the change in SAP and the anticipated benefits of certifying their properties retrospectively. Secondly, exogeneity arises from idiosyncratic variation in the types of recommendations put forward by assessing engineers. Individuals cannot manipulate either their potential or their underlying potential in order to benefit from specific recommendations, making the instrument exogenous to household behaviour. This can be seen by coverage of the treatment across all property types.

I now estimate how changes in unrealised efficiency are capitalised into property prices. Unrealised efficiency is instrumented for using two separate instrumental variables. Firstly, I use the sharp and discontinuous change in potential efficiency scores arising from the change in SAP in 2012. I employ this in an IV-RDD style framework, with the instrument taking a binary form ($\mathbb{1}.U_{2012}$). Secondly, as a sanity check to verify that the SAP change captures well the exogenous change in unrealised efficiency (confirming that Eq.(6) holds), I interact this binary SAP variable with the property's potential efficiency scores at the first EPC ($\mathbb{1}.U_{2012} \times E_{i1}^P$). Given that the change in unrealised efficiency should arise from changes in potential efficiency alone, this provides a good test for the validity of the SAP instrument. Both instruments should isolate the change in unrealised efficiency arising from change in potential efficiency as a result of the change in SAP. The second stage can be written as follows

$$\log(P_{it}) = \varphi \widehat{E}_{it}^U + X'_{it}\alpha + \phi_i + \lambda_t + \Gamma_{kt} + \epsilon_{it} \quad (7)$$

I regress the log of property prices for property i at time t on the exogenous component of unrealised efficiency and a set of property-level control variables (X_{it}) – including the total floor area and the current efficiency score. Controlling for the latter ensures that I account for any home improvements that may have taken place between EPCs and which I cannot observe. Failure to do so would risk biasing the estimates. Furthermore, by keeping current efficiency fixed within the regression means I’m able to explore just the effect of an increase in potential on property prices. Two separate specifications are estimated, like in the fixed effects models: one specification controls for property and year fixed effects (ϕ_i and λ_t), while the other controls for property and Location $\times$ Year fixed effects (ϕ_i and Γ_{kt}). The latter controls for any time-variant unobservables affecting a locality’s attractiveness; this may increase, given that we’re observing properties who suddenly become more attractive for investment.

The coefficient of interest is φ which tells us the percentage change in property prices following a unit-percent increase in unrealised efficiency arising from the change in SAP. Conditioning on property-level characteristics, and exploiting within-property variation, the coefficient tells us the extra value buyers place on the same property following an exogenous change in its potential efficiency score. We therefore make the assumption that property characteristics and its underlying investment potential remain the same between certificates.

Table 2.7 presents results from the exclusion test for both instruments. I refer to the respective instrument in the table as Z, which corresponds to binary SAP variable in Columns 1-3 and the binary SAP variable interacted with the potential efficiency score from the first EPC in Columns 4-6. Columns 1 and 4 represent the results naive cross-sectional specifications without fixed effects. Both estimates show the instruments to be correlated with property prices, returning negative and statistically significant coefficients. The instrument on its own is not excludable from the second stage regression, as it directly affects price. Property prices are notoriously sensitive to fluctuations within the market, and with regards to property aesthetics – the latter, I expect, drives these

results. Augmenting this specification with property and year fixed effects in Columns 2 and 4 removes the correlation – although the instrument is borderline statistically significant in Column 2. This approach controls for property-level unobservables like property aesthetics, age, build type, and property type, capturing features that also affect price. I argue that this combination of fixed effects omits important dynamics that come with an increase in potential efficiency, like those akin to gentrification. In this vein, I include Location $\times$ Year fixed effects to account for time-variant unobservables that make an area more attractive over time. If properties upgrade en-masse within an area, this would create a local amenity that likely causes local prices to appreciate. I present the results in Columns 3 and 6 and take the estimates as the preferred specifications moving forward into the estimation. Both columns show that both instruments do not have a structural relationship with property prices and therefore should not bear significance in the second stage equation. I therefore control for property-level time-invariant characteristics, as well as neighbourhood-level time-variant characteristics within the main IV specification.

Table 2.7: Instrument Exclusion Tests

	(1)	(2)	(3)	(4)	(5)	(6)
		U_{2012}			$U_{2012} \times E_{i1}^P$	
Z	-0.0736*** (0.0130)	-0.0113* (0.0040)	-0.0102 (0.0052)	-0.0438*** (0.0030)	-0.0020 (0.0011)	-0.0013 (0.0012)
Observations	98,262	98,262	98,262	98,262	98,262	98,262
Property FE		✓	✓		✓	✓
Year FE		✓			✓	
Location $\times$ Year FE			✓			✓

NOTES: Standard errors clustered at the property level included in parentheses. All specifications are inclusive of second-order polynomials of the running variable (EPC inspection year), with the dependent variable being property price $\log(P_{it})$. All terms are log-transformed, with both sets of regressions including a further control for property total floor area and the current efficiency score. I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP. Full results are presented in Appendix A.3. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

I argue that the instrument is both valid and excludable conditional on controlling for these time-invariant unobservables, as some properties benefit more from the change in SAP than others due to reasons unknown. For example, property age is captured by the property fixed effect, but would otherwise determine the extent to which better tech-

nological improvements are recommended.¹¹ Capturing such features ensures that the probability of receiving the treatment is not endogenous. Failure to control for these features would otherwise create heterogeneous effects that become baked into the error term. As well as ensuring the instrument is excludable, we must also be confident that it is exogenous. Earlier results from Table 2.2 highlight the exogeneity of the change in potential efficiency, showing that the change arose from variation in recommendations put forward by assessing engineers. This is out of the control of households and is not driven by property characteristics – as shown in Appendix A.3 by the discontinuity in potential efficiency scores across *all* property types. I am confident then that the empirical strategy adopted here is a robust causal estimate of the capitalisation of unrealised efficiency.

Results from the 2SLS specifications are presented in Table 2.8. Estimates presented in even-numbered columns include dynamic controls for neighbourhood quality and represent the preferred specifications. In Columns 1-2 I instrument for unrealised efficiency using the binary indicator for the change in SAP. Estimates from the preferred specification in Column 2 show that a 10% increase in unrealised efficiency leads to -0.155% decrease in property prices. This is statistically significant at all conventional levels, but with limited economic significance owing to the small price effects. Column 1 differs in its omission of Location×Year fixed effects. Estimates do not differ greatly, placing the capitalisation rate at a marginally higher -0.165%. This provides strong evidence that the change in potential efficiency scores is not correlated with neighbourhood-level factors. Both estimates benefit from high F-statistics, suggesting that the relationship between the SAP change and potential efficiency scores is strong – satisfying the IV’s relevance criteria. I now turn my attention to the results using the second instrumental variable. I use the results in this specification to explore whether the estimates are confounded due to the choice of instrument. The point estimates in the preferred specification in Column 4 show that the results barely change, moving from a capitalisation rate of -0.155% in Column 1 to -0.153% (I re-estimate both specifications including and excluding current efficiency scores in Appendix A.4, showing the importance of controlling for current efficiency). That the two are identical provides reassurance that the choice of instrument is not driving the

¹¹For example, older properties may be less likely to support more modern technologies.

results. The capitalisation rate following an increase in unrealised efficiency is the result of changes in the SAP methodology.

Table 2.8: Unrealised Efficiency Capitalisation: 2SLS Results

	(1)	(2)	(3)	(4)
	U ₂₀₁₂		U ₂₀₁₂ × E _{i1} ^P	
Second Stage				
Unrealised	-0.0164*** (0.004)	-0.0155*** (0.005)	-0.0188** (0.006)	-0.0153* (0.007)
First Stage				
Z	1.110*** (0.023)	1.113*** (0.028)	0.171*** (0.006)	0.172*** (0.007)
F Statistic	2,660	3,833	1,101	1,598
Property FE	✓	✓	✓	✓
Year FE	✓		✓	
Location×Year FE		✓		✓
Observations	98,262	98,262	98,262	98,262

NOTES: Standard errors are clustered at the property level. All specifications are inclusive of second-order polynomials of the running variable (EPC inspection year). All terms are log-transformed, with both sets of regressions including a further control for property total floor area and the current efficiency score. I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP. Full results presented in Appendix A.5. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Unrealised efficiency captures the distance a property is from achieving its potential. This represents the central goal of the energy performance directive, yet buyers attribute little value to it. Capitalisation is negative, suggesting that buyers discount their property bids in response to an increase in unrealised efficiency. In essence, a property with increasing unrealised efficiency (post change in SAP) requires additional energy efficiency improvements to attain potential at an extra cost to the buyer. Buyers wanting to undertake improvements therefore bid less for the property in order to unlock greater finance. The estimate of unrealised capitalisation of -0.15% represents φ in Eq.(7), which I show to be robust across specifications. This estimate is significantly lower in absolute terms than the estimate of the current capitalisation rate of 1.3% estimated in the Table 2.6.¹² Current capitalisation captures how buyers value the improvements *already undertaken*, as

¹²The current capitalisation rate is estimated to be as low as 0.45% in Cajias and Piazzolo, (2013) and as high as 2.2% in Aydin et al., (2020). Both still represents a stark deviation from the unrealised capitalisation rate estimated here.

buyers bid on properties with recent energy efficiency improvements. On the other hand, unrealised capitalisation captures the value of work *yet-to-be undertaken*, as we observe properties with a greater divergence between current and potential efficiency. This is an important distinction that explains the results. Buyers place a greater value on energy efficiency improvements when they do not have to install the measures themselves. If this were not the case, we would see an equal, or even larger, unrealised capitalisation rate that would account for the monetary and non-monetary costs associated with installing the same improvements.

To see how small the capitalisation rate is, the average property in the sample with an average price of £261,349 would be discounted by just £392 following a 10% increase in unrealised efficiency. The same property experiencing a 10% increase in current efficiency would be capitalised by £3,400. While the results suggest that buyers place higher value on existing improvements, they also suggest that buyers are less willing to undertake their own improvements, unless they finance through means other than discounting the price of the property.

A drawback of the standard log-specified models is that point estimates are difficult to interpret when observations belong to heterogeneous groups. In the case of EPCs, each property has an energy efficiency score that assigns them to an energy rating, but each band varies by size with respect to its range (see Appendix [A.1](#)). Therefore, a one-percent increase in the energy efficiency score of a property in rating band B is much larger than for a property in rating band E – leading to a non-standard interpretation of the baseline estimates. To ease interpretation, I standardise the regression inputs and dependent variable in Eq.(7) such that they have a mean of 0 and a standard deviation of 1. Doing so makes the point estimates easier to interpret. I employ the same specification of control variables and fixed effects and present the results in Table [2.9](#).

Table 2.9: Unrealised Efficiency Capitalisation: Standardised 2SLS Results

	(1)	(2)	(3)	(4)
	U ₂₀₁₂		U ₂₀₁₂ × E _{i1} ^P	
Second Stage				
Unrealised	-0.043*** (0.009)	-0.028** (0.009)	-0.089*** (0.022)	-0.047* (0.021)
First Stage				
Z	1.208*** (0.019)	1.112*** (0.023)	0.119*** (0.005)	0.119*** (0.006)
F Statistic	3,259	4,625	616	902
Property FE	✓	✓	✓	✓
Year FE	✓		✓	
Location×Year FE		✓		✓
Observations	98,262	98,262	98,262	98,262

NOTES: Standard errors are clustered at the property level. All variables have been standardised, including the dependent variable which captures the property price. All specifications are inclusive of second-order polynomials of the running variable (EPC inspection year). All terms are log-transformed, with both sets of regressions including a further control for property total floor area and the current efficiency score. I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Like in the previous 2SLS results, even-numbered columns include Location×Year fixed effects. Estimates using the SAP instrument alone are captured in Columns 1-2, while estimates using the SAP instrument interacted with potential efficiency at the first sale are captured in Columns 3-4. Across the board, the coefficients remain negative and statistically significant; buyers discount their property bids when unrealised efficiency increases. However, the estimates change significantly in magnitude both between fixed effects specifications and IV specifications. In the preferred specification in Column 2, estimates show that a 1 s.d increase in unrealised efficiency leads to a 0.028 s.d decrease in property prices, whereas in the Column 1, this decrease is around 0.04 s.d upon including property and year fixed effects. A 1 s.d increase in unrealised efficiency is equivalent to 10.8 points – equivalent to the size of rating band C – generating a £4500 and £6,700 decrease in property prices, respectively. These results are indicative of the discounts buyers apply when their distance from potential increases by close-to a whole rating band. Estimates from Column 4 are very similar, though the interacted instrument is not as strong in this example – explained by the significantly lower F-statistic. That being said, for the pre-

ferred specification, we estimate generally the same discount of around £6,700 following an increase in unrealised efficiency by nearly an entire rating band. If the market was perfect and buyers priced in the entire cost, this is expected to be equivalent to the cost of lifting the property up an entire band (i.e. the cost of undertaking the recommended improvements).

A benefit of this approach is that it provides an estimate of the monetary value of the price discount following an increase in unrealised efficiency that is standard across the entire distribution (in this case, a 10.8 point increase). Column 2 provides estimates for the preferred IV and preferred fixed effects specification, suggesting that prices are discounted by around £4,500 following the increase in potential efficiency after the change in SAP. This is a non-negligible amount, and one that calls into question the validity of estimates from other studies that fail to acknowledge the heterogeneous effects arising from this change. Such a large price discount following a large increase in potential begs further questions, particularly around the price effects of larger changes in unrealised efficiency. I now move onto explore the heterogeneous effects.

2.5 Heterogeneity Analysis

Heterogeneous Effects by Recommendation Type: Estimates of the unrealised capitalisation rate are low and show that buyers place little value on making their own energy efficiency improvements. However, these results are estimated at the mean and likely mask a lot of heterogeneity – particularly with respect to the types of improvements and their respective costs. I first explore the extent to which variation in the cost of implementing different recommendations affects the overall unrealised capitalisation rate. I create three different measures of recommendation costs, capturing the cost of i) implementing all recommendations, ii) implementing invasive vs non-invasive recommendations, and iii) implementing recommendations requiring work within the property’s interior vs exterior. The first type is a catch-all measure identifying the cost of undertaking all recommended improvements. The second defines invasive recommendations as those requiring the household to alter their way of living while the improvements are

installed. Non-invasive recommendations capture everything else. Finally, the third type captures the same thing to an extent, but offers a looser definition whereby all interior work is classified as invasive. Estimating unrealised capitalisation following a change in the respective costs allows me to explore how buyers perceive different types of improvements and the value they place on different types of recommendations. I estimate an augmented version of Eq.(7), whereby I interact unrealised efficiency scores with the cost of each recommendation $\mathbb{C}_{\tau it}$ of type τ .

$$\log(P_{it}) = \varphi \widehat{E}_{it}^U \times \mathbb{C}_{\tau it} + X'_{it}\alpha + \phi_i + \lambda_t + \Gamma_{kt} + \epsilon_{it} \quad (8)$$

Results can be interpreted as the unrealised capitalisation rate following a unit-increase in the respective recommendation cost. Results for the unrealised capitalisation rate following a change in recommendation costs are presented in Table 2.10 under the 2SLS specification. All recommendation costs are measured in £000s. In Column 1 I interact the unrealised efficiency score with the cost of installing all recommendations. Estimates show that for the same level of unrealised efficiency, a £1000 increase in recommendation costs leads to a price discount of around -0.17%, suggesting that a greater financing requirement to reach potential leads buyers to bid lower. Intuitively, this is to be expected as households reserve finances in order to undertake energy efficiency improvements. Turning to the marginal effect following an increase in unrealised efficiency, the increase in total recommendation costs dampen the marginally positive capitalisation rate. The marginal effect is effectively zero upon accounting for recommendation costs, suggesting that the increase in recommendation costs following an increase in potential efficiency is actually driving the price discounts. Evidence from Column 1 suggests that the baseline effect in Table 2.8 is driven by a shock in the types of recommendations put forward in EPCs and is not the result of some form of arbitrage.

Table 2.10: Unrealised Capitalisation by Recommendation Costs

	(1)	(2)	(3)
Unrealised	0.0031 (0.0080)	-0.0260** (0.0093)	-0.0260** (0.0094)
Unrealised×Total	-0.0017*** (0.0003)		
Unrealised×Invasive		-0.0058*** (0.0002)	
Unrealised×Non-Invasive		0.0025*** (0.0005)	
Unrealised×Indoor			-0.0040*** (0.00026)
Unrealised×Outdoor			0.0028*** (0.00056)
Property FE	✓	✓	✓
Location×Year FE	✓	✓	✓
Observations	75,206	75,206	75,206

NOTES: Standard errors are clustered at the property level. All specifications are inclusive of second-order polynomials of the running variable (EPC inspection year). Unrealised captures the log of unrealised efficiency. Total (i), Invasive and Non-Invasive (ii), and Indoor and Outdoor (iii) all refer to the costs of recommendations across the three variations listed in the text above. Costs are measured in £000s. All terms are log-transformed – with the exception of recommendation costs – with both sets of regressions including a further control for property total floor area and the current efficiency score. I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Column 2 explores the effect of an increase in recommendation costs for invasive and non-invasive recommendations – the former being those that require a household to adjust their daily life in order to accommodate the improvement. The estimated coefficient for unrealised efficiency is negative and statistically significant like in the baseline capitalisation estimates. Failing to account for the total cost like in the previous column makes this so, suggesting once more that price discounts are driven by this cost. An increase in the cost of invasive recommendations extenuate the price discount, meaning buyers further lower their price bids in anticipation of undertaking improvements that disrupt their daily lives. A component of this likely comes from a non-monetary cost associated with this disruption, whereby buyers require compensation to install such improvements. However, the coefficient on non-invasive improvements is positive and statistically significant, sug-

gesting that an increase in the cost installing such improvements actually reduces the discount. Both results point to the willingness of buyers to implement the easier and less disruptive recommendations, but a lack of willingness to implement those that are more disruptive to their lifestyle. The difference between the two coefficients is tantamount to the non-monetary costs of the lifestyle change associated with the installation of invasive energy efficiency improvements (which is negative). Finally, in Column 3, I explore whether capitalisation varies by whether the recommendations can be installed externally or whether it requires entry into the property. I employ this as a less rigid categorisation of invasiveness to see whether the results change. A similar pattern to before holds for the estimate of unrealised capitalisation, with a negative and statistically significant coefficient which increases after accounting for the cost of indoor (invasive) recommendations. That buyers further discount properties upon observing greater costs associated with indoor improvements is expected and suggests that they require further compensation. A similar result holds for outdoor (non-invasive) improvements, whereby an increase in their cost leads to lower discounts. Like before, the difference between the coefficients on the two interactions is tantamount to a non-monetary of the lifestyle change associated with the installation of energy efficiency improvements within the property – which is again negative.

A number of conclusions can be drawn from the results presented above: i) buyers discounting properties with greater unrealised efficiency do so by the value of the improvements required, ii) these discounts are larger when a larger proportion of the recommendation costs are for more invasive improvements, and iii) buyers appear to factor into the price discount the non-monetary cost of installing home improvements that are associated with a brief change in lifestyle. These invasive changes, alongside the non-monetary costs, exacerbate the price discounts for properties with greater unrealised efficiency. Even after accounting for these factors, buyers do not value installed vs yet-to-be installed improvements in the same way. Estimates from this analysis suggest that buyers would still underinvest in energy efficiency improvements moving forward.


Heterogeneous Effects by Distance to Potential Rating: Evidence presented in Section 2.3 shows that potential efficiency scores increased by an average of 15 points after the change in SAP. This may have led to properties changing potential efficiency rating band in some instances. I show in the previous results that increasing the unrealised efficiency score causes households to drive down their property bids as the difference between current and potential efficiency scores diverge. I now explore whether the same relationship holds when energy ratings diverge, and whether the effect is heterogeneous with the magnitude of divergence. Research into the heuristics finds that simpler labelling strategies are often the most noticeable – this extends to the energy efficiency literature (Sejas-Portillo et al., 2020). If we observe properties whose potential moves further away from their current rating band after the change in SAP, this could have implications on property prices – particularly if energy ratings are more salient than the scores. I estimate the following extension of the 2SLS regression in Eq.(7)

$$\log(P_{it}) = \varphi \widehat{E_{it}^U} \times N_{it} + X'_{it}\alpha + \phi_i + \lambda_t + \Gamma_{kt} + \epsilon_{it} \quad (9)$$

Where $N_{it} = \sum_{P=A}^G \sum_{C=A}^G \mathbb{1}[E_{it}^P > E_{it}^C]$ captures the difference between the property's current and potential in terms of the number of rating bands, which I interact this with the exogenous component of unrealised efficiency. A benefit of exploiting within-property variation means that I am able to estimate whether an increase in the number of bands between current and potential leads to significant price discounting, conditional on constant property quality after the change in SAP. A property whose distance to potential grows between EPCs would be expected to experience greater price discounts. I plot the marginal effects from the above regression in Figure 2.6. Plotted on the x-axis is the number of rating bands between a property's current and potential rating. The plot shows that having a current and potential rating in the same band, or having a current rating one-band below potential, leads to statistically insignificant price effects. Such properties have fewer improvements to make, and so buyers do not discount their prices in anticipation of undertaking improvements. Fewer improvements likely translates into fewer profitable opportunities, and so buyers seem to place little value on investing in these properties. Being two rating bands away from one's potential is more important, as buyers discount

property price by an extra 1%. The effect is non-linear, with discounts growing at a faster rate beyond this point. Properties whose current rating is three and four rating bands below their potential experience extra discounts of 2.5% and 3.25%, respectively. Buyers appear more willing to shave pounds off the price as unrealised efficiency grows, but four rating bands appears to be a minimum point – as by five rating bands below, the price discount appears to become statistically insignificant. Results from the marginal effects suggest that there exists a sweet spot in which unrealised efficiency encourages buyers to shave money off the asking price. Being too close to one's potential rating suggests that little can be done to improve efficiency, with fewer prospects of achieving returns to investment. Being too far away would also make buyers question their returns on investment but for different reasons – namely, the sheer volume of improvements that would be required. Buyers seem more willing to discount their bids in anticipation of undertaking work when the distance from potential is not too large and when achieving potential is more realistic.

Figure 2.6: Capitalisation by Distance From Potential Rating Band



NOTES: Estimates represent the unrealised capitalisation rate for each rating band a property lies away from their potential and are the result of taking the marginal effects following a unit-percent change in unrealised efficiency. The capitalisation rate is estimated on properties of varying distance from their potential efficiency, measured as the distance in rating bands between the current and potential efficiency rating. Unrealised efficiency is log-transformed and includes the same transformation on controls for property total floor area and the current efficiency score. I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP.

Heterogeneous Effects by Size of Unrealised Efficiency: The previous analysis measured unrealised efficiency as the number of rating bands between current and potential efficiency, exploring how a growing divergence between the two is capitalised into prices. A similar analysis can be undertaken using the distribution of unrealised efficiency scores themselves. While moderate levels of unrealised efficiency may incentivise investment into energy efficiency improvements, large or even small levels may instead act as a deterrent for different reasons explained above. This hypothesis is consistent with the results presented in Figure 2.6. To explore this, I deviate slightly from the standard 2SLS due to the structural break in potential efficiency scores arising from the change in SAP in 2012. Unrealised efficiency scores changed as a result, making them difficult to compare across the whole sample, requiring a less restrictive sample. In other words, the top-end of the distribution of unrealised efficiency scores are likely entirely made up of

properties with EPC inspections post-2012, thereby providing a similar setup to the 2SLS analysis. It is therefore difficult to produce a reliable distribution for the whole sample that captures the effects of a change in unrealised efficiency. Instead, I employ a fixed effects regression using the same specification of control variables and fixed effects as the previous analyses. This is detailed below.


$$\log(P_{it}) = \varphi_k[\mathbb{U}_{2012} \times \mathbb{Q}_k \times E_{it}^U] + X'_{it}\alpha + \phi_i + \lambda_t + \Gamma_{kt} + \epsilon_{it} \quad (10)$$

The regression differs from the previous analyses in two ways: i) I do not instrument for unrealised efficiency, and ii) I include the term $\mathbb{Q}_{it}$ which splits unrealised efficiency scores within the sample into quintiles. Splitting into quintiles ensures that we have consistent bands for use within the regression, meaning we're comparing properties within the same intervals of unrealised efficiency scores. I use this variable within a triple interaction term that interacts with an indicator for the change in SAP $\mathbb{U}_{2012}$ and the unrealised efficiency score. Interacting with the change in SAP means that I am able to explore the heterogeneous effects of a change in unrealised efficiency on property across bands and across time periods (both before and after 2012). The result is an estimate of the unrealised capitalisation rate, by quintile, both before and after the change in SAP; whose estimates are reflected by φ_k . Adopting a 2SLS approach would not allow me to incorporate this flexibility and would ultimately generate imprecise estimates. I therefore explore whether large levels of unrealised efficiency due to changes in the SAP are valued differently to the large levels of unrealised efficiency prior to the change. Results from estimating the marginal effects are presented in Figure 2.7.

First comparing between the pre- and post-SAP results, we see no discernable differences in the point estimates for low-to-moderate levels of unrealised efficiency. For bands 0-5 to 15-25, we get roughly similar estimates suggesting that the change in SAP did not affect how buyers discount properties with low prospects for improvement. Estimates of unrealised capitalisation within these bands are not surprising. Properties in the 0-5 band had little unrealised efficiency both before and after the change in SAP, experiencing very little change in capitalisation. Properties in the 5-15 band observe a more noticeable,

but not significant, change whereby post-SAP capitalisation falls. This likely captures properties who were in the 0-5 band pre-2012 and whose new recommended improvements are minor with little prospect of returns. While the 15-25 group represent properties with greater development prospects, we do not observe heterogeneous effects pre- and post-SAP change. A slightly larger point estimate after 2012 might point to the introduction of recommendations that were increasing in popularity after 2012, and for which larger returns to investment were possible.

Figure 2.7: Capitalisation by Distance From Potential Rating Band



NOTES: Estimates represent the unrealised capitalisation rate for each quintile of the unrealised efficiency distribution and are the result of taking the marginal effects following a unit-percent change in unrealised efficiency. The Pre 2012 point estimate for the band 35+ is positive and significant, but very large owing to a small sample size. The estimate has been omitted from the figure for ease of interpretation. Unrealised efficiency is log-transformed and includes the same transformation on controls for property total floor area and the current efficiency score. I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP.

Properties with higher levels of unrealised efficiency observe more noticeable changes in the form of larger price discounts. Properties in the 25-35 quintile, with high levels of unrealised efficiency prior to the change in SAP, were not particularly discounted (though, estimates are noisy). However, after the change in SAP, these properties became heavily discounted. The difference here is that these properties likely had greater potential


following the change in SAP, after previously being in the lower bands. This suggests that buyers are much more likely to discount a property when they have expectations of reaching greater potential. However, if we look at the point estimates for properties with unrealised efficiency in excess of 35 points, we see that the increase in potential does not make this group more likely to discount. Like in the previous analysis in Figure 2.6, this suggests that potential efficiency is important, but only if it is attainable. Estimates from the figure point to the fact that properties with medium-to-high levels of unrealised efficiency experience greater discount post-SAP. Such large changes in unrealised efficiency likely come from the introduction of new recommendations put forward by assessing engineers post-2012 – potentially explaining the increase in discounts.

If we now turn to the distribution within the Pre- and Post-2012 groups, we see a relatively similar pattern again for the bands 0-5 to 15-25. We observe an increase in the discount as we move further from potential, though not at a fast rate. Buyers do appear to discount for larger values of unrealised efficiency at the same rate between the two periods. The story is different for properties with larger levels of unrealised efficiency. Discounts continue to increase, but only in the Post-SAP group, where the coefficient increases significantly. This suggests that properties within this band are discounted by larger levels than in other bands – with a clear heterogeneous effect. Like in the analysis in Figure 2.6, we observe a zero effect on properties with unrealised efficiency greater than 35 points – alongside a similar pattern of discounts across the distribution. Results presented in this heterogeneity analysis show that potential is important, but that one's distance to potential is equally important in determining buyer's valuation of property. The latter is overlooked within the literature, with estimates here showing that large levels of unrealised efficiency deter households from undertaking improvements and which limits the capitalisation of potential into property prices.

Heterogeneous Effects by Property Type: Beyond differences in the underlying energy efficiency of properties and the prescribed recommendations, one final channel through which we would expect heterogeneity is through property type. One may expect

heterogeneous effects owing to the ability of different property types to accommodate energy efficiency improvements. For example, households living in flats are likely more hampered in their ability to improve their energy efficiency compared with households living in detached properties. As such, if one believes that discounts following an increase in unrealised efficiency are attributable to buyer's desire to conduct energy efficiency improvements, then discounts should differ between property types based on the ability to undertake these improvements. I explore this by adopting the same specification in Eq.(9), differing only in that I interact unrealised efficiency with the property type. The specification retains the same control variables and fixed effects. I take the marginal effects of the unrealised capitalisation rate by property, with results presented in Figure 2.8.

Figure 2.8: Capitalisation by Distance From Potential Rating Band



NOTES: Estimates represent the unrealised capitalisation rate for each property type and are the result of taking the marginal effects following a unit-percent change in unrealised efficiency. Unrealised efficiency is log-transformed and includes the same transformation on controls for property total floor area and the current efficiency score. I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP.

I omit flats from the analysis given their limited ability to make energy efficiency improvements.¹³ Results from the marginal effects show that the average household in

¹³Flats are included in the pooled estimates, but including here is more problematic. Property type is

each property type discounts price following an increase in unrealised efficiency. This is expected, given that each property type should have the ability to undertake improvements. Point estimates show that terraced properties receive the largest discounts, while detached properties receive the lowest discounts. However, the estimates are noisy with large confidence intervals, meaning I am unable to identify any heterogeneous effects between property types.

2.6 Conclusion

As the UK economy adapts to meet stringent Net Zero targets by 2050, we cannot neglect the need for energy efficiency improvements at mass scale. By better insulating homes and taking measures to reduce energy consumption, we can reduce emissions generated within residential settings – which currently account for around 14% of total emissions. Yet, these actions are taken at the household-level and require private investment. Combining property-level data on transactions and energy efficiency characteristics across the UK between 2009-2019, I show that the appetite for such investments remains low. I define a new approach to exploring capitalisation, termed unrealised efficiency, that captures the energy efficiency not yet realised by the household. Measured as the difference between current and potential efficiency scores, I explore how buyers value future energy efficiency improvements. This chapter produces a novel instrumental variable which exploits the change in SAP (2012) and the subsequent random variation in types of recommendations put forward by assessing engineers during EPC inspections. I show that the change in SAP introduces new recommendations that represent technological improvements on those previously, and which deliver greater energy efficiency improvements. Recipients of a technologically improved recommendation experience a 4 point increase in their potential efficiency score, compared with just a 1 point increase for those receiving a recommendation for the older technology. I exploit the change in these recommendations as an exogenous treatment which arises due to the idiosyncratic variation present in the assessing engineer’s assessment of a property. Such variation generates a discontinuous increase in potential efficiency scores, creating exogenous variation in unrealised efficiency scores

included as an interaction with the endogenous and exogenous variable, creating separate estimates for each type. This deflates the estimates for other property types, and so I omit flats in this analysis.

that leads households to diverge from their attainable potential. This chapter then explores the price effects of this change.

Estimates show that an increase in unrealised efficiency, driven by this discontinuity in potential efficiency, leads buyers to discount their property bids. Following a 10% increase in unrealised efficiency, property prices fall by around -0.15%. Comparing this with an increase in current efficiency, we observe an increase in property prices of around 1.3%. The two measures differ subtly: the former captures the value buyers place on making their own improvements, while the latter captures the value buyers place on the same improvements that have already been made by the previous owners. The findings (taken in absolute value) show that buyers place less value on undertaking their own energy efficiency improvements, which is akin to a willingness to pay. A set of heterogeneity analyses are undertaken to explore the mechanisms driving this effect. I explore whether these effects are heterogeneous by the cost of undertaking recommended improvements. Findings from the analysis suggest that the price discount is fully offset by this cost, showing that buyers accurately account for the investment required to raise potential. This is an important finding as it shows that buyers price into their property bids the cost of making energy efficiency improvements. However, buyers place a lower value on installing the improvement compared with adopting it from the previous owner. Though, it is not just the cost of undertaking improvements that buyers price into their property bids, it is also the nature of the improvements. Estimates from the same heterogeneity analysis show that buyers further discount properties that require more invasive improvements to increase their energy efficiency. These are the improvements that require households to adapt their daily lifestyle during the installation, and which incur a non-monetary costs to the household. Such costs generate dis-utility and explain why some households do not make certain types of improvements. In a separate set of heterogeneity tests, I show that it is not just the potential that is capitalised, but the degree to which this is attainable and to which it delivers returns on investment. Properties with low levels of unrealised efficiency are not discounted owing to their limited room for improvement and limited scope for financial gains. Similarly, properties with large levels of unrealised efficiency

are also not discounted, owing to the need for too many improvements and uncertain returns. Buyers prefer properties with mid-ranging levels of unrealised efficiency that provide greater certainty around returns with the need for fewer improvements.

Creating meaningful progress with regards to climate change involves a concerted effort towards meeting the end goal – lowering energy use whilst making its production more sustainable. Short-term financial incentives raise barriers to achieving this goal and alter the balance of preferences. In the residential energy efficiency arena, this arises in the form of high costs associated with making home improvements. Create real environmental change requires households to look beyond these costs and towards the ultimate goal of attaining their property’s potential energy efficiency. This chapter provides a starting point in thinking about how buyers value energy efficiency in the future and how preferences for undertaking improvements affect this. To the best of my knowledge, this is the first study to provide causal estimates of this. Greater research is needed to explore the forward-looking nature of buyers in this field; this can take a number of forms. An important one is the extent to which individual characteristics affect the results of this chapter, particularly with respect to heterogeneity arising from characteristics such as age, education, voting behaviour. A limitation of this chapter is that I can only access data at the property-level which is insufficient to explore such topics; research into these heterogeneous effect could provide useful context to the results found in this chapter. Similarly, one may explore how this is affected by housing tenure or even by the length of occupation. Renters and owners have been shown to have different preferences of making improvements and will attribute different valuations to energy efficiency. A wealth of literature exists that explores this relative to current capitalisation, but not unrealised capitalisation (see Fuerst et al., 2020; Kholodilin et al., 2017; Cajias et al., 2019). These are but a few avenues where more research could be undertaken.

The policy implications of this research are potentially important. I show that buyers value the improvements undertaken by the property’s previous owners far more than those they have to undertake themselves. This threatens the uptake of retrofitting, and

calls into question our ability to meet Net Zero targets. Buyers are shown to favour greater certainty in the returns to investment, shying away from retrofitting properties that require too many improvements, or those that require too few to deliver financial return. Financial uncertainty arises from buyers needing to increase their financial commitment in the face of higher unrealised efficiency. Moving forward, policy needs to account for these heterogeneous effects in order to incentivise the retrofitting of the most inefficient properties who stand to gain the most from energy efficiency improvements.

A Appendix

Figure A.1: Typical Format Energy Performance Certificate

Energy Performance Certificate (EPC)



17 Any Street, District, Any Town, B5 5XX

Dwelling type: Detached house
Date of assessment: 15 August 2011
Date of certificate: 13 March 2012

Reference number: 0919-9628-8430-2785-5996
Type of assessment: RdSAP, existing dwelling
Total floor area: 165 m²

Use this document to:

- Compare current ratings of properties to see which properties are more energy efficient
- Find out how you can save energy and money by installing improvement measures

Estimated energy costs of dwelling for 3 years	£5,367
Over 3 years you could save	£2,865

Estimated energy costs of this home

	Current costs	Potential costs	Potential future savings
Lighting	£375 over 3 years	£207 over 3 years	
Heating	£4,443 over 3 years	£2,073 over 3 years	
Hot water	£549 over 3 years	£222 over 3 years	
Totals:	£5,367	£2,502	

These figures show how much the average household would spend in this property for heating, lighting and hot water. This excludes energy use for running appliances like TVs, computers and cookers, and any electricity generated by microgeneration.

Energy Efficiency Rating

Very energy efficient - lower running costs

(92 plus) **A**

(81-91) **B**

(69-80) **C**

(55-68) **D**

(39-54) **E**

(21-38) **F**

(1-20) **G**

Not energy efficient - higher running costs

Current	Potential
49	76

The graph shows the current energy efficiency of your home.

The higher the rating the lower your fuel bills are likely to be.

The potential rating shows the effect of undertaking the recommendations on page 3.

The average energy efficiency rating for a dwelling in England and Wales is band D (rating 60).

Top actions you can take to save money and make your home more efficient


Recommended measures	Indicative cost	Typical savings over 3 years	Available with Green Deal
1 Increase loft insulation to 270 mm	£100 - £350	£141	✓
2 Cavity wall insulation	£500 - £1,500	£537	✓
3 Draught proofing	£80 - £120	£78	✓

See page 3 for a full list of recommendations for this property.

To find out more about the recommended measures and other actions you could take today to save money, visit www.direct.gov.uk/savingenergy or call 0300 123 1234 (standard national rate). When the Green Deal launches, it may allow you to make your home warmer and cheaper to run at no up-front cost.


NOTES: The figure presents an example of an EPC passed over to the buyer at the time of the property transaction.

Figure A.2: Share of EPC Inspections by Year



NOTES: The figure shows the EPCs requested by EPC inspection year as a proportion of the total number of EPCs requested across the sample window.

Figure A.3: Potential Efficiency Scores by EPC Inspection Year and Tenure



NOTES: The figure shows the average current, potential and unrealised energy efficiency scores by EPC inspection year for all property types and using all properties within the sample. The x-axis denotes the year in which the EPC inspection takes place.

Table A.1: Full Results – Change in Potential by Recommendation Type

	(1) XS	(2) Panel	(3) Panel×Sap	(4) Panel×Sap
1.Growing	0.814*** (0.011)			
1.Declining	-2.317*** (0.030)			
Growing		3.824*** (0.018)	2.821*** (0.030)	3.700*** (0.018)
Declining		-2.511*** (0.052)	-2.564*** (0.052)	-3.345*** (0.053)
Growing×Post-SAP			1.298*** (0.019)	
Declining×Post-SAP				3.287*** (0.073)
1.SAP		8.210*** (0.024)	7.358*** (0.027)	7.823*** (0.023)
Current	0.512*** (0.0007)	0.650*** (0.002)	0.641*** (0.002)	0.641*** (0.002)
Floor Area	-0.02*** (0.0001)	-0.002*** (0.0006)	-.002*** (0.0005)	-.002*** (0.0005)
Property FE		✓	✓	✓
R ²	0.693	0.700	0.702	0.704
Observations	3,410,871	733,398	733,398	733,398

NOTES: Standard errors clustered at the property-level in parentheses. Column 1 uses dummy variables to signify if a property possesses recommendations from the increasing/decreasing groups and uses all other improvements as the hold-out category. Columns 2-4 instead use the number of recommendations in each group for improved accuracy. Further controls are included for the current efficiency score and total floor area, while Column 1 is augmented further with controls for the property’s construction age, build type, and property type. This regression represents a linear specification, whereby controls are measured in levels. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.2: Full Results – Fixed Effects: Unrealised Capitalisation

	(1) Location×Year	(2) Fixed Effects	(3) Baseline	(4) Fixed Effects
Unrealised	-0.051*** (0.001)	-0.024*** (0.001)	-0.054*** (0.001)	-0.026*** (0.001)
Current		0.133*** (0.003)		0.134*** (0.003)
Floor Area		0.241*** (0.006)		0.261*** (0.007)
Property FE	✓	✓	✓	✓
Year FE			✓	✓
Location×Year FE	✓	✓		
R ²	0.985	0.988	0.592	0.658
Observations	98,262	98,262	98,262	98,262

NOTES: Standard errors clustered at the property level in parentheses. The baseline model includes property and year fixed effects. The estimation is limited to the same sample window as later estimates. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.3: Full Results – Instrument Exclusion Tests

	(1) U_{2012}	(2) U_{2012}	(3) $U_{2012} \times E_{it}^P$	(4) $U_{2012} \times E_{it}^P$
Z	-0.0102 (0.0052)	-0.0113* (0.0040)	-0.0013 (0.0012)	-0.0020 (0.0011)
Current	0.166*** (0.003)	0.170*** (0.003)	0.166*** (0.003)	0.169*** (0.003)
Observations	98,262	98,262	98,262	98,262
Property FE	✓	✓	✓	✓
Year FE		✓		✓
Location×Year FE	✓		✓	

NOTES: Standard errors clustered at the property level included in parentheses. All specifications are inclusive of second-order polynomials of the running variable (EPC inspection year), with the dependent variable being property price $\log(P_{it})$. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table A.4: Unrealised Efficiency Capitalisation: 2SLS Results

	(1)	(2)	(3)	(4)
	U ₂₀₁₂		U ₂₀₁₂ × E _{i1} ^P	
Second Stage				
Unrealised	-0.0202*** (0.0045)	-0.0155*** (0.0045)	-0.0637*** (0.0054)	-0.0153*** (0.0069)
First Stage				
Z	1.113*** (0.028)	1.110*** (0.023)	0.228*** (0.007)	0.172*** (0.007)
F Statistic	3,833	2,660	2,579	1,598
Current		✓		✓
Property FE	✓	✓	✓	✓
Location×Year FE	✓		✓	
Observations	98,262	98,262	98,262	98,262

NOTES: Standard errors are clustered at the property level. This Table recreates the results from Table 2.8 after including and excluding controls for current efficiency scores. All specifications are inclusive of second-order polynomials of the running variable (EPC inspection year). All terms are log-transformed, with both sets of regressions including a further control for property total floor area and the current efficiency score. I restrict the Post-2012 sample to include four years of data such that it matches the period prior to the change in SAP. Full results presented in Appendix A.5. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table A.5: Full Results – Unrealised Efficiency Capitalisation: 2SLS Results

	(1)	(2)	(3)	(4)
	U ₂₀₁₂		U ₂₀₁₂ × E _{i1} ^P	
Second Stage				
Unrealised	-0.0155*** (0.005)	-0.0164*** (0.004)	-0.0153* (0.007)	-0.0188** (0.006)
Current	0.137*** (0.004)	0.140*** (0.004)	0.137*** (0.006)	0.138*** (0.006)
Floor Area	0.239*** (0.007)	0.257*** (0.007)	0.239*** (0.007)	0.257*** (0.007)
First Stage				
Z	1.113*** (0.028)	1.110*** (0.023)	0.172*** (0.007)	0.171*** (0.006)
Current	-0.736*** (0.018)	-0.747*** (0.018)	-0.698*** (0.018)	-0.712*** (0.018)
Floor Area	-0.374*** (0.026)	-0.388*** (0.022)	-0.362*** (0.026)	-0.377*** (0.026)
F Statistic	3,833	2,660	1,598	1,101
Property FE	✓	✓	✓	✓
Year FE		✓		✓
Location×Year FE	✓		✓	
Observations	98,262	98,262	98,262	98,262

NOTES: Standard errors are clustered at the property level. All specifications are inclusive of second-order polynomials of the running variable (EPC inspection year). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Chapter 3: Neighbourhood Effects in the Market for Energy Efficient Housing

3.1 Introduction

The residential sector is set to play an important role in achieving Net Zero targets. By insulating properties to retain heat and adopting cleaner energy sources, the sector could significantly reduce its carbon footprint – yet, household investment in these areas remains low. Challenges facing the property market may in-part explain this, as rising property prices coupled with stagnant wage growth stifle property investment.¹⁴ Literature has explored the capitalisation of improvements into property prices only to find modest growth (Costa and Kahn, 2011; Fuerst et al., 2015; Fuerst et al., 2016; Aydin et al., 2020; Sejas-Portillo et al., 2020), calling into question the returns on investment. A limitation of the existing literature is its failure to account for the spatial nature of property markets. Instead, the decision to make property investments is treated as a private decision with isolated price effects. In reality, both are likely location-specific and are determined by the characteristics of one’s surroundings – affecting both the decision to invest and the subsequent returns. Evidence from the housing literature supports the latter. Price spillovers between neighbours have been found to arise from information reported in listings (Szumilo, 2020) and by listers who price their properties after learning from their peers (Anenberg, 2016). A separate body of literature finds that peer effects are present in the adoption of energy efficient technologies (Bollinger and Gillingham, 2012; Bollinger et al., 2020; Wolske et al., 2020). In failing to account for spatial factors within property markets, we potentially overlook an important channel through which households decide to invest and through which we observe differing price effects.

In this chapter I define a set of energy-efficient property clusters and explore how property prices and home improvements differ at their border. I leverage energy efficiency scores reported in London property listing adverts between 2009-2019 – along with random variation in their quality and location – to define these clusters in a way that replicates the

¹⁴Average property prices in the UK have tripled since the turn of the millennium as housing costs account for one-fifth household income (Ministry of Housing and Government, 2024).

radii search function used by buyers in popular property portals. Spatial variation in the location and timing of property clusters on the market allows me to identify two research questions: firstly, whether buyers selecting into a cluster pay a price premium, and secondly, whether subsequent exposure to these clusters alters the likelihood that incumbent households – those already residing within the neighbourhood – make home improvements. Households may enter or exit the treatment based on the energy efficiency of their surrounding neighbours. By employing property fixed effects, I exploit household’s entry and exit to test whether any premium arises due to the reporting of energy efficiency scores. Efficient clusters are bounded in space, providing a discontinuity at which energy efficient properties are no longer statistically clustered. I exploit this in a spatial regression discontinuity design (RD) to test the whether property prices are discontinuous at the border of a cluster. Identification comes from the fact that properties either side of the border do not differ in terms of structural observables, only in their energy efficiency. A premium paid to purchase within a cluster following an efficient listing would constitute a selection from households with preferences for energy-efficiency. Incumbent household’s are exposed to different households that chose to purchase within these neighbourhoods over time, exogenously changing their peer group. I exploit the temporal nature of the property clusters to explore whether this affects the likelihood that incumbent households make energy efficiency improvements following exposure to households with such preferences. Exposure is measured as the number of years a property falls into an efficient property cluster and the subsequent property churn. I use this to instrument for a household’s decision to make improvements. This is expected to lead to a higher probability that an incumbent household undertakes energy efficiency improvements. Using the exogenous change in such household’s peer group over time and limiting my sample to households that do not knowingly purchase in an observed property cluster, I estimate neighbourhood effects for households with no demonstrable preferences for energy efficiency, ex-ante.¹⁵

I find that energy efficient property clusters create residential selection, as buyers draw information on surrounding properties. This generates a premium for these properties

¹⁵I use the term *peer group* synonymously with one’s neighbourhood.

of 1.6% that buyers pay to reside within a cluster, translating to roughly £6,100 for the average property. Selection occurs because buyers perceive a benefit to residing within a cluster, as buyers anticipate higher returns to their investments in these locations. Estimates are robust to changes in the selected bandwidth and upon accounting for access to local amenities, with robustness checks showing that the premium is also not present in the years when clusters are not on the market.¹⁶ Evidence of a premium highlights that buyers are attentive to energy efficiency scores across neighbourhoods and are able to distinguish between high- and poorly-efficient neighbourhoods. These findings provide evidence that the premium results from contemporaneous information spillovers through property listings. An event study analysis confirms this by identifying a premium for properties within a cluster that disappears once the clusters is no longer visible on the market. This would suggest that buyers of such properties have preferences for energy efficiency. I refer to these as selecting households. Exposure to these selecting households over time increases the likelihood that incumbent households make energy efficiency investments. I find that each extra year of exposure increases the probability of upgrades among incumbents by 6%. By restricting the sample to incumbent households that did not knowingly purchase at the time of an observed property cluster, this probability falls to 1.1%.¹⁷ After accounting for neighbourhood effects, residence within a cluster bears no significance on the uptake of improvements once exposure is accounted for. It is therefore not the neighbourhood that buyers value, but the neighbours.

This chapter provides a methodological contribution to the literature. I propose an approach that allows households to select their own peer group, thereby defining property clusters. The housing and environmental literature commonly considers peers as those within proximity (see Bollinger et al., 2020; Szumilo, 2020; Bigler and Janzen, 2023). Some studies have gone further, seeking to bound peer groups spatially according to characteristics in order to control the dispersion of interactions, but often-times this is done using administrative geographies (see Bollinger and Gillingham, 2012, Nauze, 2023, Sokołowski, 2023). Such approaches fail to capture the granularity required and likely

¹⁶in years for which listings do not identify a cluster as being energy efficient.

¹⁷Correlated effects, a la Manski, (1993), are shown to account for a large proportion of the effect – providing a good sense check for the adopted approach.

overstates the true extent of who households believe their peers to be. Doing so risks failing to adequately capture the neighbourhood in which the effects originate. In this chapter I propose a data-driven approach that adopts a measure of spatial autocorrelation in energy efficiency scores to capture similarities in both energy efficiency and proximity. By measuring similarity in energy efficiency, I am able to capture potential spatial interactions (where households imitate each others behaviour). I employ a clustering algorithm with distance decay to capture similarity in a way that identifies tipping points in energy efficiency scores, thereby allowing proximity to define itself. The approach allows households to choose their own peer group by uncovering underlying spatial interactions that derive from the property's similarity with others nearby. In this way, two households would be classed as peers if they were located nearby and shared a similar energy efficiency score that differs significantly from the local average. Bounding these properties according to the tipping point in energy efficiency scores produces spatially bounded peer groups which represent the geographical area from which I explore neighbourhood effects. Examples from the literature above consider how selection into energy-efficiency adoptions influences nearby households in a static sense, whereby one's peer group remains constant. However, little to no evidence explores how the selection of households into neighbourhoods generates neighbourhood effects. I adopt a data-driven approach to capturing peer groups in order to fill this gap.

Building on a wealth of literature that captures the drivers of residential sorting (Couture et al., 2023; Koster et al., 2014; van Duijn and Rouwendal, 2021; Guerrieri et al., 2013), I further contribute to the literature by exploring this through the lens of residential energy efficiency. I aim to show how households choose their residential locations based on spatial variation in energy efficiency scores, as reported within property listings. To the best of my knowledge, I produce the first study exploring this. Furthermore, I use this setting to investigate the effects of this selection. Residential selection represents an exogenous change in neighbourhood demographics which then affects the peer group of incumbent households (i.e. those that do not move). Using the churn of households within a neighbourhood to identify the change in peer group preferences for energy efficiency, I

identify the effects this has on home improvements locally. The literature considers the decision to make improvements as a household decision and fails to account for spillovers of neighbourhood influence. However, this omits outside pressures arising from one's neighbourhood, such as the effects of social norms, or heterogeneity in the returns across neighbourhoods. This chapter aims to be the first to capture such effects in the housing market.

The rest of the chapter proceeds as follows. I review the relevant literature in Section 3.2. The approach taken to define energy efficient clusters and the characteristics of the dataset are presented in Section 3.3. The empirical strategy used to identify a premium is presented alongside the results in Section 3.4. I then turn my attention towards exploring the mechanisms that drive these results in Section 3.5. The chapter finishes with concluding remarks in Section 3.6.

3.2 Literature Review

This paper straddles a number of different fields across the discipline by studying how residential energy efficiency reporting influences behaviour within property markets. In this section, I review the related literature from the fields of environmental, urban and public economics. I group the literature among three distinct lines. I first review papers that explore the underinvestment in residential energy efficiency. Researchers put forwards many explanations as to why this occurs and evaluate the effectiveness of corrective policy measures. I then go on to explore the literature pertaining to peer effects in both the housing and energy literature – owing to the limited evidence of neighbourhood effects in the latter – to identify the gaps in which this chapter sits. A review of the methods used also places into context the novelty of the approach taken within this study. I finish the section by reviewing the literature that explores the capitalisation of energy efficiency investments into property prices. An extensive literature dating back to the 1980s, there remains an underutilisation of spatial approaches resulting in limited evidence around how capitalisation differs across space. I review the literature and observe this as a distinct oversight. This represents a large gap in the literature and one which potentially explains

sub-optimal uptake in improvements.

3.2.1 Underinvestment in Energy Efficiency: Poor Energy Savings

Early papers explored the adoption of energy-efficient technologies in residential settings and the drivers of such adoption. Seminal papers from Hausman, (1979) and Dubin and D. L. McFadden, (1984) catapulted this literature by estimating the discount factors applied to the purchasing of energy-related durables. Hausman, (1979) models the adoption of durables noting the trade-off between upfront capital costs and energy savings. Buyers discount savings by roughly 20% when applied to data for air conditioning units, increasing among consumers with lower income. Dubin and D. L. McFadden, (1984) find the exact same discount factor in the context of space- and water-heating appliances, but argue that previous approaches to estimating energy consumption are lacking. The authors provide an econometric approach to show that discount rates are overstated if researchers fail to correct for exogeneity between the demand for durables and their final use. Essentially boiling down to a selection problem, durables represent a derived demand pertaining to the demand for energy, thereby absorbing features of energy consumption into estimates of the discount rate. The paper provides a rebuke to M. Parti and C. Parti, (1980) who adopt a 2SLS approach in an attempt to circumvent this issue. However, the instrument uses electricity prices which are non-excludable. A limitation of these papers lies in the assumptions they make regarding a durables lifecycle. The models cannot account for depreciation of the stock, and must assume either finite or infinite life which has implications on a household's capital expenditure at end of life. An alternative approach to circumvent this issue and that of selection arising from the purchase of durables comes from D. McFadden et al., (1977) who employ a structural approach that models decisions regarding one's portfolio of durables. Allowing consumers to adapt their durable's portfolio in the long-run affects their energy consumption and hence their savings, and also means that one does not have to make explicit assumptions about the durable's useful life. Poorer durables increase energy consumption, thereby accentuating savings after their removal – and are accounted for in this approach. While researchers have found ways to somewhat circumvent issues of selection, one cannot easily derive the demand curve for

electricity that arises from stepwise pricing present in energy markets Taylor, (1975). The limitation of this literature then is that household preferences obscure measures of energy consumption and savings, and hence the resulting discount factors for the adoption of durables.

Estimated discount factors are subject to debate. That being said, discount factors of any magnitude can be explained by the poor returns to investment. A number of papers find this, but Dubin et al., (1986) do not. The authors show this not to be the case in one of the earliest papers exploring the issue, arguing that savings are much greater than engineering estimates would have one believe (up to 15% greater), owing to the exclusion of household characteristics in these models. Realised savings are around 90% of those predicted (referred to as the realisation rate). Issues again around estimating electricity demand in the presence of heterogeneous preferences potentially obscures the resulting savings, as well as the small sample (400 households). Hirst, (1986) estimate a lower, but still relative high realisation rate of 63% in the context of space heating. The bias from early papers in the literature is all the more likely given the contradictory evidence from the modern literature. Ansar and Sparks, (2009) show this in their estimation of realisation rates, pointing to the fact that pre-1990s, the median estimate sat at 67%. While Metcalf and Hassett, (1999) estimate energy savings from attic insulation, savings are much lower than those of engineering estimates. Energy savings of 9.7% compared with that of 50% from the engineering estimates represent sizeable differences. The authors attribute this to the selection problem driving the uptake of improvements and the failure of engineering estimates to account for the differing preferences among these households. Experiments, both natural and randomised, have been applied to estimate energy savings to circumvent these issues; approaches that Allcott and Taubinsky, (2013) show to be the most accurate to perform welfare analyses of the impact of energy cost information on durable good purchases. Fowlie et al., (2018) exploit a randomised control trial (RCT) using a weatherisation programme implemented in Michigan that granted home improvements (heating) to low income households. The paper estimates energy savings savings between a control and treatment group of households that were eligible for grants, only the latter

group were randomly assigned extra help in applying for funds. Presented alongside a quasi-experimental design that compares appliers with receivers, savings range from 10% in the quasi-experimental design to 20% under the RCT. Explanations of this variation arise from measurement error owing to energy consumption, as is the case in the previous papers. While the presence of savings supports earlier estimates, upfront costs are found to be twice that of the savings which are actually only one-third of the savings predicted by engineering estimates. A natural experiment observed by L. W. Davis et al., (2014) corroborates these findings for a subsidy and replacement of durables programme in Mexico. The authors find energy savings of 8% following the replacement of refrigeration units – much lower than savings predicted by engineering estimates. No savings are realised for air conditioning units, however, highlighting instead a rebound effect – whereby energy consumption increases on the back of savings. One can explain this by the fact that air conditioning is a luxury and only comprised 10% of overall take up in the experiment, suggesting uptake by wealthier households. The weatherisation programme instead finds no evidence of a rebound effect which seems interesting. An explanation could be found in the fact that low-income households exclusively form the treatment and control groups and that these households behave differently to others and have higher demand elasticities for energy (see Alberini et al., 2011, Díaz and Medlock, 2021, Harold et al., 2017) owing to income constraints. Older papers from the literature fail to account for rebound effects which means that estimates of energy savings are likely capturing changes in both preferences and prices. Evidence from Boomhower and L. Davis, (2020), who exploit a natural experiment subsidising the installation of residential air conditioning in South Carolina, confirm similar energy savings but argue that this is much more nuanced than previously thought. The authors confirm the seasonality effects of energy savings observed in Dubin et al., (1986) (owing to the demand for air heating and cooling), but find that savings differ during the time of the day. Using smart metre data, they show that engineering estimates underestimate the true energy savings, which are higher when demand and prices are higher. Accounting for this timing increases energy savings by as much as 50% as households seek to divert their consumption. Unlike in Metcalf and Hassett, (1999) who also use billing data, this paper only explore the treated – those already in ownership of

the durable – circumventing issues of selection. Allcott and Greenstone, (2017) show that actual energy savings are also inferior to those predicted. Estimating a structural model of the efficiency of two residential retrofit programmes, aided by a randomised control trial, they show that realised savings were roughly half of those predicted. A degree of this can be explained by high ex-ante savings causing selection into the programmes by households that would not otherwise have acquired subsidies, and that were not the ones standing to benefit most. Realisation rates – the ratio of actual to predicted savings – are therefore prone to complications. Higher rates are estimated in older papers that fail to adequately capture selection.

The literature provides conflicting evidence on the magnitude of these savings, but in most cases find these savings to be less than ex-ante estimates. Findings of this nature gave rise to literature exploring the energy efficiency gap, from which multiple definitions arise. The previously analysed papers would define the gap as the shortfall of investments arising from positive discount factors – or the underinvestment arising from the failure to undertake those improvements with positive net present value. Others would define the gap as notion that private investments in energy efficiency are much lower than the rate which is socially optimal Jaffe and Stavins, (1994). Either way, this reflect rampant underinvestment, which the literature seeks to explore. A number of hypotheses are put forward. For this study, the most relevant is that of behavioural bias and uninformed agents. I explore this under the umbrella of asymmetric information.

3.2.2 Underinvestment in Energy Efficiency: Asymmetric Information

Bias in this respect arises from households being misinformed about how their daily activities translate into energy usage. The presence of bias, argue Enrich et al., (2024), is proved by buyers that change their behaviour after receiving nudges and which persists after nudging ends. They show using an experiment conducted in Spain that varying household energy tariffs during different times of the day reduced energy consumption by 1-2%, and by as much as 9% at peak times. Furthermore, they find that consumption falls even after reverting back to uniform pricing – suggesting consumers were previously

misinformed about energy prices. Allcott, (2011a) conduct a similar RCT in Chicago with similar results. Buyers receiving information about intra-day price hikes adjust their behaviour and reduce consumption by between 1-2% and by as much as 5-14% throughout different times of the day. Studying townships in Connecticut, Jessoe and Rapson, (2014) argue that the same price information reduces consumption by more, around 7%. Both studies notify households the day prior to a price hike, but the authors go further by testing the household response when releasing information alongside real-time data from Smart Metres; consumption among this group was lowered by 17%. No evidence of load shifting is found around the time of the price shocks, which the authors use to argue bias, as households deviate from their previous preferences. Conducting a similar experiment in Connecticut in the early summer months, Faruqui et al., (2012) find similar evidence of a decline in energy consumption – but of a smaller magnitude; suggesting that climate plays a role in determining the magnitude. Like the previous papers, the authors find no evidence of load-shifting, suggesting that buyers are not modifying their actions but their preferences.

While the previous papers suggest that buyers are misinformed as to their energy usage, these findings are not universal. Fabra et al., (2021) find no evidence of an effect of real-time pricing, but find that households shift consumption to off-peak hours. Importantly, participants in this study are not directly informed of price hikes while they were more informed about Time-of-Use rates (the former were reported online) – suggesting that the information provision identified in Jessoe and Rapson, (2014) is actually an important component to modifying preferences. Borenstein, (2002) confirm this using case-study evidence from California rate hikes around the time of nationalisation that did not curtail residential demand. Boomhower and L. Davis, (2020) would instead suggest that changing the time one uses energy is not evidence of bias in light of new information, but simple rationality. They find that buyers actually adjust their preferences and usage when savings (and prices) are at their highest. After removing any information provided by nudges, they still find that buyers act rationally, suggesting that buyers are not uninformed but are instead optimisers. Arimura et al., (2012) find the same rationality in the

context of demand-side management (DSM) policies, whereby households receive rebate for reducing their energy usage. Reductions in usage are around 1% in the DSM period, increasing to 1.8% afterwards. The absence of nudges, showing that buyers reduce usage when faced with demand-side management policies – promoting rationality over bias. Beyond it being rationality, Boomhower and L. W. Davis, (2014) suggest that these DSM programmes mainly target infra-marginal households; those that would alter their preferences regardless – a similar view is shared by Joskow and Marron, (1992). Though this is dispelled by Ito, (2015) in the context of energy savings – suggesting instead that buyers make rational decisions based on evidence. This makes it difficult to explain away the effects as purely causal, as the reviewed papers do not explicitly address this.

Other might simply argue that this is not a feature of limited knowledge or uninformed agents, but of behavioural bias. A seminal paper in the literature comes from Samuelson and Zeckhauser, (1988) who argue that buyers making decisions favour the Status Quo under the guise of certainty in the result. Through lab experiments and field studies, the authors present evidence that this is not a conscious choice but an implicit form of risk-aversion. This result forms the basis of why we estimate discount factors; yet, experiments from the paper were conducted under strict time pressure and were of a complex nature – inherently inducing the status quo. This serves as a metaphor for anecdotal evidence arising from the literature on energy efficiency adoptions. Behavioural bias is captured in these decisions as individuals hold onto poorly-informed beliefs that obscure them from acting optimally. Behavioural applications in the energy efficiency literature agree that bias drives the under-adoption of durables, but argue that reality is more nuanced than this. Carroll et al., (2016) and Kallbekken et al., (2013) both conduct experiments to ascertain the uptake of energy-efficient durables (in this case, clothes dryers) and find no evidence of an increase in the uptake of more efficient models following disclosure of a model's operating costs. The latter do find a significant increase in the sale of energy-efficient dryer models when this intervention is combined with staff training around the benefits of such purchases. Information that does not confirm one's biases is therefore irrelevant in these studies – which Ölander and Thøgersen, (2014) would refer to as a

framing problem which presents information in a manner that allows buyers to opt-out of making a decision. When this is not an option, buyers are more receptive to information. This calls into question the attentiveness of buyers to information as opposed to their willingness to accept it. Newell and Siikamäki, (2014) explore this in a field experiment for the adoption of water heating appliances. Presenting respondents with twelve different energy-efficiency labels for an appliance, they find that simple and more digestible information increases buyer's WTP for appliances. For example, buyers are more responsive to labels providing information on operating costs over those converting the same information for CO₂ expulsion. Unlike in Carroll et al. and Kallbekken et al., Deutsch, (2010) do find evidence of an increase in the energy efficiency of appliances sold when reporting information. They argue that reporting the lifecycle costs of an appliance is more accessible which drives the result, suggesting that the information way information is reported could be the problem. Several papers have explored such biases in the context of heuristics (see Lacetera et al., (2012) and Pope, (2009)). Information is important in that it *nudges* buyers to act in certain ways. While heuristics is important, it rests on the assumption that buyers can process complicated information. Attari et al., (2010) would suggest that this is not the case; their study finds that undervaluation of energy costs is more common among lower-educated households (Sterman and Sweeney, 2007; Gardner and Stern, 2008). Such findings are countered by Dietz et al., (2009) who suggest that it is not information per-se that is the problem, but how it is delivered; finding that social media is a more effective tool for providing information that modifies behaviour. A critique of these papers, like Samuelson and Zeckhauser, (1988), is that they place too much weight on lab experiments without real-world application. From this perspective, one misses a good deal of nuance. While Carroll et al., (2016) and Kallbekken et al., (2013) offer an analysis of store-based purchasing decisions, limited sample sizes ranging between four seventeen months call into question the power of the results and selection arising from seasonality. Papers reviewed up to this point prescribe provide solutions to these biases based on information interaction alone. Some argue that consumers – in failing to be attentive to energy efficiency information – may indeed be acting rationally. For example, inattentiveness may in fact be a best response if information is costly to

gather (Sallee, 2014), if buyers only sporadically update their information or news disperses slowly (Reis, 2006), or if information is difficult to digest (Sims, 2003).

Researchers exploring behavioural biases towards the adoption of durables have also tested the efficacy of different remedies to correcting them. Sallee, (2014) would argue that even poorly designed policies work if they reduce the cost of attention to the matter. In the context of energy efficiency, this is perhaps a slightly simplistic view of correcting bias. Allcott and Taubinsky, (2013) note the bias when observing the uptake of Compact Fluorescent versus incandescent lightbulbs – the former being much more efficient but with lesser uptake. Running two RCTs in the form of a lab and field trial to understand household's WTP in the presence of both information and subsidy-based nudges, bias is found to arise from the undervaluing of energy savings. WTP only increases marginally when participants are supplied information – confirming results from earlier papers – but that effects are much larger when coupled with subsidies. Allcott et al., (2014) instead argue that this is much more nuanced, and that subsidies mainly work for those with larger internalities – those that undervalue energy costs relative to the price of more efficient durables. Such consumers need subsidies, as taxes (a la pigouvian) on energy costs only increase one's undervaluation. Allcott and Greenstone, (2012) formalise this approach with a conceptual model, showing that taxes affect a durable's adoption rate while subsidies affect its utilisation rate. However, this is a simplified approach to the understanding the impact of subsidies. Heutel, (2011) would argue that this literature does not account adequately for time-inconsistency in preferences, however. Instead, they would advocate for two-stage policies that target consumers at time of purchase and time of use. Yet, even in the designing effective policy responses, unintended consequences still remain. Houde and Aldy, (2017), like Boomhower and L. W. Davis, (2014), suggest that these policies favour infra-marginal consumers who would have purchased durables regardless of subsidy. Fowlie et al., (2015) argue instead that subsidies simply ease financial burdens, and that they do not adequately increase the uptake of durables owing to the non-monetary costs associated with their adoption.

The major drawback of these studies – which is also their strength – is the reliance on RCT and natural experiments. While these methods are preferred to identify causal estimates, they shroud the influence of other factors external to the participants – such as externalities and peer-to-peer influences. The papers adopting these techniques reviewed previously find evidence of behavioural bias, whereas those using quasi-experimental methods do not (Fabra et al., 2021; Boomhower and L. Davis, 2020). This raises questions as to the extent to which the literature can accurately account for external influences in modifying individual behaviour and in diagnosing under-investment.

3.2.3 Peer Effects and Individual Behavioural

A common source of external influences overlooked by the previous literature comes in the form of peer effects. The literature has explored this along a number of different avenues – most importantly for this studies, in relation to energy consumption and housing investment. In this subsection, I review both strands of this literature.

Peer effects play a key role in encouraging households to change their energy consumption patters. Allcott, (2011b) show that this operates through a conformist approach whereby households adapt their behaviour to meet a social norm. In the context of energy consumption, they find that household energy consumption falls by 2% in response to information on the consumption of their neighbours. The author calls these non-price interventions, which, in terms of magnitude, seem equally as effective in correcting bias as the studies reviewed earlier. This also speaks to targetting non-monetary costs as documented in Fowlie et al., (2015). The study exploits data from the Home Energy Report (HER) framework conducted by a utility provider that mails out the usage information of one’s peers to customers. Ayres et al., (2013) use the same data, but focus their natural experiment on proximate households who receive the same home reports. This differs from Allcott who assumes that behaviour changes are a personal choice; Ayres instead estimates the effects of interactions on behaviour modification as treated households discuss the reports, tying the effects to social interactions. Adopting the same treatment and time frame, the authors find larger reductions in energy consumption of 2.5%. But the

papers only explores effects in the short-run. Allcott and Rogers, (2014) later adopt these data in a much larger sample to explore the long-run implications of conformism. They find that social norms spur short-run action in energy conservation, but that this wanes once the information channel is removed. Habitual change requires a long and persistent channel through which peer effects are reported, with the effects being more persistent if buyers adopt durables that force them to reduce consumption. This ties social norms together with habit formation. Brandon et al., (2017) expand on this approach using the same data and contradict these findings. The authors use moving households in conjunction with energy consumption data to estimate the probability that a household adopted durables. They use this then to quantify the long-run (post-occupier) reductions in consumption that are attributable to the durable and find a very loose link between social norms and habit formation. As much as 55% of the reduction comes from adoption of the durable, instead arguing that reductions in energy consumption come from adoptions as opposed to behaviour modifications. Such findings drive a wedge between the short- and long-run mechanisms. A shortcoming of these studies is that one's peer group is only defined based on proximity – not on characteristics. Schultz et al., (2008) argues that normative messaging is important when applied to a particular reference group. They apply a model of resource conservation in hotel rooms, in which they provide normative messaging about conservation efforts of previous guests – which should be of similar characteristics. They find that this delivers greater changes in short-run behaviour than other, less specific messaging. Peer effects induce conformism to social norms, but one that dies out over time according to Ito et al., (2018). They argue that this is a form of moral suasion that plays on individual's ethics. The authors find that, like Brandon et al., (2017), such effects do not form habitual reductions in energy consumption – adding that economic incentives are much more effective. Bollinger et al., (2020) produce a similar result in the context of water conservation – finding that peer effects in conservation are more present when accompanied by strong price signals.

The literature is therefore mixed in the conclusions about the strength of peer effects in modifying behaviour. Evidence suggests that short-term effects are present, but that

long-term effects are much more reliant on other factors to permanently modify behaviour – such as market forces or persistent information channels. The context of energy conservation however is a choice that households take regularly based on real-time information and norms; in the context of housing improvements, however, this is a one-time decision taken by buyers. One would expect peer effects to be more effective in facilitating investments as a result. A number of papers explore peer effects in the adoption of energy efficient technologies as a result of adoptions locally; a strand of the literature applies this approach to specifically explore adoptions of solar PV, with differing results. Bollinger and Gillingham, (2012) find that every extra solar PV adoption within a zip code increases the likelihood of another by c.0.8 percentage points; Bigler and Janzen, (2023) explore more local effects from adoptions within 100m and estimate a much lower rate of around 0.02 percentage points; whereas Mundaca and Samahita, (2020) estimate an increase in adoptions of .44 panels for every PV adopted within half a mile. Subtle differences exist in the approaches adopted, however. Bollinger and Gillingham, (2012) circumvent the issues around simultaneity by exploiting the lag between ordering and installation of solar PV owing to regulatory constraints (i.e. bureaucracy). They argue that this avoids simultaneous decision-making between adopter and peer such that they can identify the cause-effect relationship in adoptions. Using location-time fixed effects, they further control for unobservables that drive adoption. However, conducting this analysis across a zip code area means that the study likely masks a large degree of heterogeneity and risks overestimating the true peer effect – as households cannot be guaranteed to interact across such great distances. Mundaca and Samahita, (2020) follow this same approach, but argue that the zip code is too large a geography. They instead estimate the probability of peer adoption at different distance intervals within a zip code. However, the analysis fails to employ dynamic fixed effects that account for temporal variation in unobservables over time in estimation. Bigler and Janzen, (2023) on the other hand restrict peer effects within 100 metres to capture the true interactions between household. The authors instead exploit lagged adoptions to get around issues of simultaneity and argue that, while plausible, a fixed effects regression risks failing to capture time-varying unobservables; In contrast, they employ an instrumental approach that exploits the local benefits of solar generation

to explain uptake. The methodologies applied here rely on a complex set of assumptions on both the timing and exposure of interactions – leaving one to question the plausibility of the methodologies. The results differ quite dramatically, suggesting that peer effects are much more nuanced – and may even follow that of (Schmidheiny, 2006) who argues that conformism relies on similarity in characteristics between peers.

A separate, albeit small, strand of the literature explore the role of peer effects in the demand for housing. The difficulties in identifying a household's peers means that very few papers adopt such studies. Two papers that do so are Patacchini and Venanzoni, (2014) and Szumilo, (2020). Patacchini and Venanzoni, (2014) explore how one's school peers influence their household's demand for housing investment. The authors exploit a network of school-based social interaction that – following an investment in the network – they argue induces further investments. They tie this to the mechanism of *keeping up with the Joneses* as information spreads across the network. Peer effects are estimated using a peers-of-peers approach that abstracts from information exchanges between direct friends, as this introduces the problem of correlated effects (a la Manski, 1993). The characteristics of these indirect peers is then used as an instrument for the behaviour of one's direct friends. A similar approach is adopted in Szumilo, (2020) who explore peer effects in property prices originating from learnings about home investments locally. The author exploits information asymmetries to explore whether property bids are a learned behaviour in the absence of information. A linear-in-means model is adopted that includes the mean characteristics of one's peers within the regression to account for both exogenous and correlated effects. This represents a variation of the approach adopted in Patacchini and Venanzoni, (2014). While networks are represent the gold standard in identifying peers, one could argue that the peers-of-peers methodology provides a tenuous link between households. While a novel strategy, it requires the reader to make a number of strong assumptions regarding how information is traded. So to does the adoption of the linear-in-means which prevents one from incorporating into the model heterogeneity in one's peer's.

3.2.4 Capitalisation of Energy Efficiency Improvements

While the majority of researchers are in agreement that energy efficiency improvements are capitalised into property prices, disagreements arise in other areas – one of which being the magnitude of such returns. Early papers in the field explore the signalling effect arising from possession of a green label or certificate. Yet, the decision to obtain a certification is non-random, thereby causing issues of selection bias. Brounen and Kok, (2011) were one of the first papers to recognise this problem in their efforts to explore the price effects of green certification in the Netherlands. Using a two-stage Heckman procedure to control for selection into obtaining certification, they find that possession of a green label increased prices 3.6% relative to related properties. A similar study of the extensive margin was undertaken in the US state of California by Kahn and Kok, (2014) who place the premium at around 2.1%. The authors instead apply propensity score matching to compare certified with non-certified properties that provide a strong counterfactual. The estimates likely differ due to their locational context. Bypassing issues of selection, Zalejska-Jonsson, (2014) estimates the willingness-to-pay (WTP) for low-energy dwelling using a natural experiment. Using survey data from Swedish households, they find the WTP to be around 5%. Data for the study came in the form of subjective property price valuations; coupled with a limited sample size, the estimates should be taken with a pinch of salt. Estimates of the extensive margin vary. Brounen and Kok, (2011) – like the UK – define energy-efficient dwellings as those rated C or above. This range is tested in Fuerst et al., (2015) who find that property prices in the UK are positively related to energy efficiency ratings. A/B-rated properties sell for around 5% more than those that are D-rated; C-rated properties sell for 1.8% more. Estimates of the extensive margin in the previous papers fall within this range to provide a credible midpoint estimate. Looking at this for scores, the authors find a 1% increase to translate to 0.1% capitalisation rate using a repeat sales approach. Estimates of the intensive margin exploit changes in prices between sales as a way to partial out unobserved factors influencing price, and also factors affecting one's certification (influencing selection on unobservables). The same approach is adopted in Cajias and Piazzolo, (2013) using residential data from Germany but find differing estimates, placing the capitalisation rate at 0.45%. This seems abnormally high

in comparison with the literature. The paper also uncovers a non-monotonic relationship between prices and energy ratings – suggesting lower-rated properties sell for more. This differs from estimates out of the UK; it also differs from estimates around the same time frame from Singapore; where Deng et al., (2012) confirm the findings from the UK, that prices increase monotonically with energy performance.

These papers were part of the emerging capitalisation literature when uptake of energy reporting was low; public knowledge and perception also differ to that of the present day, leading to variation in findings. While these papers tackle selection arising from acquiring certification, they do not address selection arising from its disclosure at time of sale. This is a valid concern as the aforementioned papers all draw on samples from years prior to – or years close to – the implementation of the EU Energy Performance of Buildings Directive requiring disclosure of energy certificates. Hyland et al., (2013) raise this concern early in the literature; the authors find evidence of households selecting into reporting energy certificates based on observables. Households of larger size and closer to urban areas were more likely to disclose, which were those households with higher property values ex-ante. Hyland et al. replicate the study from Fuerst et al., (2015) on a dataset of property transactions between 2008-2012 (post-EU-directive). Controlling for sample selection, they find much higher premiums: 11% for A-rated properties and 5.8% for B-rated. This places under scrutiny the estimates of studies using data around this time frame.

Recent papers from the field have access to better data with longer time frames. A critique of the earlier literature is the failure to appreciate the dynamic nature of energy efficiency improvements. These papers consider only a property’s possession of a certificate, its energy ratings, or its efficiency score at one point in time – omitting information from energy efficiency scores that call into question their causal nature. Sejas-Portillo et al., (2020) highlight this as a source of bias by showing that property price premiums vary discontinuously for those at the threshold of the next rating bands. Applying a Regression Discontinuity Design, they show that properties at the threshold but in the

next rating band receive premiums ranging between 0.8-2.5% reflecting the value of the improvements undertaken. They also show that properties at the threshold but in the lower rating band are much more likely to make quick energy efficiency investments to move into the next band – reflecting that seller’s behaviour varies discontinuously at the border. Failing to control for the score and hence the distance to the next band biases estimates of capitalisation in two ways: i) as those in the higher band obtain price premiums, and ii) those in the lower band may in fact adjust upwards their asking prices. A recent study by Aydin et al., (2020) show that measurement error also biases estimation when it is not corrected for. Such error arises from differences in engineering assumptions between assessing engineers and even from errors in data entry. The authors provide the one application of a 2SLS approach from the literature. They apply this to correct for measurement error and to obtain causal estimates of the capitalisation rates present throughout the Netherlands; finding that energy efficiency improvements are capitalised into prices at a rate of 2.2% for every 1% increase in the score.

3.3 Data

I acquire the data used within this chapter from two open-source databases: the Energy Performance of Buildings database (supplied by MHCLG) and the HM Land Registry Price Paid database.¹⁸ Both property-level datasets, the former records data on energy efficiency scores, consumption, and structural characteristics, allowing this chapter to control for property size (floor space, number of rooms, bedrooms, reception rooms), quality (discrete values of quality for a property’s roof, walls, windows), and energy consumption and environmental impact (energy usage, efficiency, and CO₂ emissions). Detailed data on transaction prices are obtained from the Land Registry to track price changes between energy performance certificates (EPCs). Contained within the dataset are the exact transaction dates for each sale, which I use to isolate the properties listed on the market in a given year. Both datasets lack a common property identifier, requiring the use of a multi-rule matching process to merge the two. More details can be found in the

¹⁸MHCLG is the Ministry of Housing, Communities and Local Government.

3.3.1 Energy Performance Certificates Policy

From 2008, households looking to sell their property were legally required to supply prospective buyers with an up-to-date EPC, as per EU legislation. Mandating EPCs was expected to increase transparency in the selling process and provide households with an incentive to differentiate their property on the market by improving their energy efficiency.¹⁹ A by-product of the legislation meant that buyers could now distinguish a property's quality, sending signals to potential buyers. A key measure of quality comes in the form of energy efficiency. Energy efficiency measures are reported within an EPC in two ways: i) in a categorical manner that provides the property with a rating between A-G (A being the most and G the least efficient), and ii) in a continuous manner providing a score ranging from 0-100.²⁰ Both are produced in line with the The Standard Assessment Procedure (SAP) which calculates a fuel-cost based measures that is subsequently used to calculate energy efficiency scores. After accounting for property size, properties that maintain ambient temperatures at higher (lower) running costs receive lower (higher) energy efficiency scores. The SAP procedure has undergone changes in 2009, 2012 and recently in 2022 to reflect changes in best practice – most recently to account for climate information in the calculation of energy efficiency scores across the country.

As the rollout of EPCs continued, basic energy efficiency reporting became integrated into popular search platforms from 2010, allowing buyers to form expectations regarding a property's efficiency well in advance of a sale.²¹ More in-depth information, such as energy costs and structural characteristics, usually exchange closer to the time of sale, factoring less into buyer's purchasing decision. A drawback for most studies is that EPCs are valid for a period of ten years, irrespective of recent refurbishment. This raises concerns around the accuracy of their information. For example, their effectiveness in capturing home improvement activity due to the non-necessity of updating the EPC, and the lack of real-time information arising from the depreciation of existing property features. Both

¹⁹A detailed example of an EPC can be found in Appendix B.5.

²⁰With properties rated A and properties scoring closer to 100 being the most energy efficient.

²¹Popular platforms in the UK include Rightmove and Zoopla.

may lead buyers to make uninformed pricing decisions. For this chapter, I make two assumptions: i) rational homeowners request new EPCs after improvements have been undertaken, and ii) buyers are attentive to EPCs across their chosen neighbourhood at the time of purchasing. The timeframe of the sample (2009-2019) means there are only limited EPCs that may reach the ten-year expiry date; using EPCs as a proxy for whether or not a property upgraded then seems reasonable.

3.3.2 Dataset Used to Define Clusters

I exploit a dataset of unique sales – each accompanied by an EPC containing energy efficiency information – across the London Travel To Work Area (TTWA). Using the London TTWA as a target area has both benefits and limitations. The large population provides a healthy degree of property churn unmatched by other cities. However, property in London is comparatively more expensive, impacting on household’s valuation of energy efficiency characteristics. I show in Appendix B.8 that prices in London peak at twice the national average before declining from 2016 – meaning the trend does not worsen across the sample. Households facing higher housing costs are more likely to be credit constrained, less able to undertake home improvements, and will rank other property features above energy efficiency. To overcome this, I remove flats from the dataset and keep only free-standing properties for the econometric analysis.²² Households living in flats are more likely to belong to this constrained group whilst also encountering greater difficulty installing home improvements, owing to their modern construction and constrained designs. Their inclusion in the analysis would likely bias any estimate of a neighbourhood effect. I collect a cross-section of 550,000 unique observations from across 395,000 unique properties over a period of 11 years which I use to identify energy efficient property clusters from property listings in a given year.

3.3.3 Defining Energy Efficient Clusters

In this section I set out the approach used to define the spatial extent of the efficient clusters. Clusters capture similarities in energy efficiency scores between nearby prop-

²²Flats are used to assign treatment/control groups but are subsequently dropped from the econometric analysis given their limited scope for energy efficiency improvements – making them a poor case study for exploring neighbourhood effects.

erties and define the extent of the neighbourhoods attractive to those with preferences for efficient properties. The terms cluster and neighbourhood are used synonymously and are akin to peer groups in this chapter. Clusters then provide a strong test-bed on which to evaluate the presence of neighbourhood effects in energy efficiency investments. Previous attempts in the housing literature only consider proximity when defining peer groups (Szumilo, 2020, Bollinger and Gillingham, 2012, Bollinger et al., 2020). I argue that this is an oversimplification of a complex decision-making process taken by buyers, as it i) fails to provide reasoning as to who households perceive their peers to be, and ii) defines peer/neighbourhood effects in a static way – failing to account for the influence of new movers. This chapter adopts the view that the influence of one’s peers is driven by the household’s choice of peer group as much as their proximity. The same is concluded by Anenberg, (2016) who finds that sellers set prices in accordance with similar properties listed nearby (neighbours that are proximate and characteristically similar).

Neighbourhoods with a high degree of similarity should be those with historically high levels of spatial interactions, as households imitate each other. This makes such locations a good test-bed for identifying neighbourhood effects. The approach to creating clusters aims to capture these exact locations and identify tipping points in energy efficiency scores that define the boundary of such interactions. At this border, I estimate the premium buyers attach to locating within a highly-efficient property cluster over-and-above a poorly-efficient property cluster. Leveraging energy efficiency scores reported within property listings, I capture the information visible to buyers at the time of their property search. Similarity between scores is measured on a continuous scale and is calculated using the clustering algorithm developed by Ord and Getis, (1995). The algorithm runs across each property within the sample and estimates the degree of spatial autocorrelation in the energy efficiency score with proximate neighbours. Similarity is measured as the degree of spatial interconnectedness between the scores. Running the algorithm over properties listed on the market in each year returns a set of clusters identifiable from listings and allows me to test whether the information observed by buyers informs their pricing decisions related to residential energy efficiency.

The Getis-Ord algorithm has been applied across multiple disciplines, mainly in fields related to geography and health. Its most notable deployment in the field of environmental economics comes from Graziano and Gillingham, (2014) who explore clusters of Solar PV adoption and its diffusion across space in Connecticut, US. The authors only use the outputs in a set of summary stats without employing the within the empirical strategy. The algorithm has also been employed to analyse the clustering of economic activity, productivity and employee remuneration across the UK (Patacchini and Rice, 2007). While the algorithm is a strong tool for producing data and descriptive evidence, its application in econometric studies remains limited.

The flexibility of the algorithm allows me to construct clusters based on the timing of information released to the market through property listings. In a later robustness check, I re-estimate the algorithm after removing the listing (information) channel (see later in the section). The algorithm is presented mathematically in Eq.(11).

$$G_{it}^* = \frac{\sum_{j=1}^{n_t} w_{i,j} x_{it} - \bar{X} \sum_{j=1}^{n_t} w_{i,j}}{S \sqrt{\frac{n_t \sum_{j=1}^{n_t} w_{i,j}^2 - (\sum_{j=1}^{n_t} w_{i,j})^2}{n_t - 1}}} \quad (11)$$

Where $\bar{X} = \frac{\sum_{j=1}^n x_{jt}}{n}$ and $S = \sqrt{\frac{\sum_{j=1}^n x_{jt}^2}{n} - (\bar{X}^2)}$. For each property i , the algorithm identifies the difference between the local sum and the expected sum of energy efficiency. The local sum takes the current energy efficiency score x of each neighbouring property j listed in year t and weights it relative to the spatial proximity of all $\{i, j\}$ pairs. Spatial proximity represents a measure of density and the degree to which properties could be considered to exhibit close contact. The expected sum takes the average energy efficiency of all neighbours $\bar{X}$ and multiplies by the sum of the spatial proximity. Given the degree of proximity of one's peers, this estimates a baseline energy efficiency that the locality ought to exhibit. Subtracting the expected from the local sum identifies properties belonging to clusters where energy efficiency is higher- or lower-than that expected (the numerator of Eq.(11)). This is enough to identify the sign of the difference between the local and expected sum, but not its statistical significance. Using the denominator, the

algorithm normalises this difference by the standard deviation of the energy efficiency scores and spatial weights which outputs a Z-score to be used for significance testing. One identifies a significant cluster when the local sum differs from the expected sum by such a large amount that the difference cannot be the result of random chance. Those with large (small) differences are identified as significant (insignificant) clusters. Two types of clusters are identified under this approach: clusters of properties with greater-than-average and clusters with lower-than-average energy efficiency scores. Using this approach achieves the goal of allowing peer groups to form on the basis of characteristics (energy efficiency) as well as proximity (defined by the weights).


I place a set of restrictions on the algorithm to ensure the proximity is captured accurately. Neighbourhood effects are likely stronger across shorter distances, so I first restrict the proximity of the weights w . I allow each property to find neighbours within 0.5km and impose an inverse distance weighting that assigns lower weights to more distant neighbours. Imposing distance weights allows me to choose an arbitrary search radius and means that properties at a distance greater than 0.3km are effectively excluded. After choosing neighbours and calculating the strength of spatial relationships for all properties (G_i^*), I then sort properties into clusters. Clusters are formed by grouping properties with greater similarity scores that lie within 300m of each other. The chosen grouping method aligns with Rossi-Hansberg et al., (2010) who find this to be the half-life of spatial spillovers.²³ One can think of this as a residential cluster. It is plausible therefore that each household is able to visually observe properties within its cluster and draw information – thereby acting as a peer group. Using yearly listings data allows a property and area’s treatment status to change across the sample as the area’s perceived energy efficiency varies exogenously with the properties entering onto the market. This creates an entry-and-exit style treatment that I exploit later in the chapter to estimate neighbourhood effects in energy efficiency improvements.

I provide a visual illustration of the clustering output in Figure 3.1. Each point represent

²³For example: if property A groups with property B and property B then groups with property C , all three properties then form a cluster.

a property sold within the respective year alongside its cluster classification. To avoid detecting false relationships, the algorithm is run separately for areas North and South of the river Thames such that relationships do not transcend the river.

Figure 3.1: Getis-Ord Analysis of Property Listings: London TTWA 2012




NOTES: The Getis-Ord statistics is computed on a property-by-property basis and returns hotspots (areas with higher-than-expected efficiency) and coldspots (areas with lower-than-expected efficiency). Insignificance captures properties that do not differ significantly from the average efficiency. This figure shows the G_i^* statistic calculated for each property in the sample in 2012 using Eq.(11).

Points coloured in red identify properties belonging to efficient clusters (hotspots), whereas points coloured in blue/green identify properties belonging to inefficient clusters (coldspots). *Inefficient clusters* is a catch all term used to identify locations predominantly fostering poorly-efficient properties or properties for which no significant spatial relationships exist. I omit such clusters as they are not important for the empirical strategy. Efficient clusters are identified across London with no obvious spatial pattern – which points to randomness in the spatial distribution. By grouping properties based on these categories, I am able to create clusters bounded in space; these then represent a boundary at which energy efficiency scores reach a tipping point. At the edge of these boundaries, properties become less energy efficient. I cluster properties in a way that replicates as closely as possible

the buyer's property search procedure. Properties are grouped if they are all deemed as hotspots and if they lie within 300m of each other. These are then connected in a dot-to-dot fashion to define the spatial extent of the similarity (see Appendix B.9). A minimum bounding geometry is applied to the resulting polygon designed to encapsulate it within the smallest parcel possible. Figure 3.2 provides a visual illustration. A circular bounding approach is chosen for two reasons: i) it allows for a smooth border that does not favour one side of the cluster more than another (in the way that a shape with vertices would), and ii) it provides a result similar to the radius search function available to buyers scouring popular property portals (see Appendix B.4). Buyers are expected to search for properties in a similar fashion to that replicated here. Turning attention towards the output we see even covering of efficient clusters across London – a number of which actually arise within inefficient clusters. In this case, preference is given to the efficient cluster with properties designated as such. Efficient clusters can also contain properties with poor energy efficiency scores. This is important for the identification strategy, as their inclusion allows any premium at the border to reflect the signalling arising from energy efficient properties on the market. Without energy-inefficient properties, one might attribute this instead to a general amenity associated with neighbourhoods where all properties are energy efficient.

Figure 3.2: Highly- and Poorly-Energy-Efficient Clusters: London TTWA 2012



NOTES: Neighbourhoods are created by clustering properties based on their z-scores and employing a bounding geometry (the circular polygons) to envelope hotspot and coldspot properties in the smallest geometry possible. This figure illustrates an example of the bounding process using inputs from Figure 3.1.

3.3.4 Summary Statistics

While this cross-section of property sales is sufficient to identify every efficient cluster, it does not provide the most appropriate setting to estimate a price premium for a number of reasons: i) the inability to control at a granular level for location characteristics, ii) the computational inefficiency associated with trying to do so, and iii) issues related to the correlation between the location of the cluster and the energy efficiency of residing properties. For this reason I adopt a repeat-sales approach that observes the same properties in at least two different time periods (two different sales). From the dataset I obtain a panel of 190,000 unique observations. The full sample of repeated sales are retained, including those properties that are purchased for speculative reasons or for quick renovations, as such buyers are expected to respond to market information and peer influences. A benefit of this approach is that it avoids bias in estimation by using a different sample from that used to define the clusters. I can be confident in this way that correlations

between similarities in energy efficiency and prices are minimised. Summary statistics for the sample used within the econometric analysis are presented in Table 3.1. These are presented at the border for both the treatment and control group, where the former identifies properties within an efficient cluster and the latter captures those on the periphery.

Table 3.1: Balancing Test at the Border of Efficient Clusters


	Treatment		Control		Balancing	
	Mean	Sd	Mean	Sd	t	p> t
Price	435,381	272,871	481,044	332,648	-7.41	0.000
Adj. price	542,018	129,456	519,241	137,901	8.38	0.000
Dist. to border	45	31.9	47.1	37.8	-3.02	0.003
Current efficiency	73.8	14.6	64.3	15.3	31.15	0.000
Potential efficiency	81.4	10.9	76.8	12.3	19.38	0.000
Unrealised Efficiency	7.6	10	12.5	12.04	21.9	0.000
Energy cost (.sqm)	7	3.9	9.3	4.6	-26.72	0.000
Energy consumption	169.4	122.7	245.4	128	-29.80	0.000
Window glazing %	84.4	31.8	81.8	33.7	2.99	0.003
Floor Area	101.6	42.1	99.6	43.3	2.37	0.018
No. rooms	4.6	1.4	4.7	1.4	-1.62	0.105
Observations	5,146		4,533			

NOTES: This table presents the unconditional means for a set of property characteristics. Data presented are for properties within 112m of the border, the bandwidth selected for the baseline RD later. Treated neighbourhoods are the energy-efficient neighbourhoods denoted by the Getis-Ord clusters. Prices here are stripped of their time and location drivers by including property and time fixed effects. Energy cost measures the estimated spend on utilities (heat, lighting and hot water) in GBP for a house of given structural characteristics and size.

The first thing to notice is that regular prices are significantly higher at the border for properties in the control group – £481k versus £435k. One can be confident that this discrepancy does not reflect differences in the structural characteristics of properties as the number of bedrooms – an important determinant of prices – is balanced at the border. A more plausible explanation lies in the significant differences between current energy efficiency scores from which the clusters are defined. This likely drives a vintage effect whereby properties on the periphery of clusters are a little older. The average current energy efficiency score for treated properties (74) exceeds that of control properties (64) by roughly ten points. The difference is large and statistically significant, potentially explaining the price differentials. Such a difference confirms that the approach taken

to define clusters is robust and delivers the expected tipping points. If buyers truly value higher energy efficiency, then this ought be reflected in higher prices in treatment neighbourhoods. Adjusting prices to remove the effects from energy efficiency characteristics reflect exactly this. The average price for a property within a treated cluster is now around £542k compared with just £519k for a property in the control group. Higher prices provide descriptive evidence that buyers can distinguish between treatment and control neighbourhoods through their attentiveness to energy efficiency scores. Reassuringly, differences between the potential energy efficiency scores are not economically significant – though they are statistically significant. With scores of 81 and 77 in treatment and control neighbourhoods respectively, it is reasonable to assume that this is not large enough to warrant changes in buyer preferences. A slight concern is that a score of 81 is the cut-off for energy rating band B, leading the average property in each group to fall into different rating bands. Controls for unrealised efficiency – the difference between the potential and current energy efficiency score – are included in the RD which control for a property’s potential. This measure captures the energy efficiency not yet realised by the household. Balancing potential efficiency scores between groups is less important given that I control for it in this way. Such differences are important for this paper and instead explain why properties in some areas make energy efficiency improvements more so than others. Only minor differences in structural characteristics exist at the border – mainly in terms of floor area which is controlled for in the empirical strategy. Price discontinuities at the border should therefore reflect differences in current energy efficiency scores. It is worth noting that treatment groups have a notably higher proportion of homes constructed post-1976 than in control groups (see Appendix B.16). I take the stance that this is not considerable enough to, first, be easily noticeable, and second, to alter buyers preferences for a neighbourhood outside of that explained by energy efficiency. To support this view, I show in a similar comparison in Appendix B.14 that there are no noticeable differences in the built form of these properties (i.e. detached, semi-detached, terraced etc.). The repeat sales approach in any case should capture any differences in construction age beyond what is already controlled for.

Figure 3.3: The Distribution of Energy Ratings at the Border



NOTES: This figure illustrates the proportion of properties in the sample that occupy each combination of current-potential energy rating. The total in the right-hand column represents the proportion of properties falling in each current energy rating band.

Having established significant differences in energy efficiency scores at the border, Figure 3.3 presents a graphical representation of the distribution by treatment status. The vertical and horizontal spread of both graphs are similar – illustrating a similar distribution of potential energy efficiency scores. Higher energy efficiency scores are more prevalent in the treatment group where 71% of properties rated B-C compared with only 46% in the control group. Conversely, lower scores are more prevalent in the control group where nearly 50% are rated D-E compared with just 26% in treatment neighbourhoods. The differences are stark – evidencing the identifying assumption that buyers are able to identify neighbourhoods with clusters of energy efficient properties. An important feature of the approach to defining neighbourhood treatment is that both groups have some degree of distribution across the rating spectrum. In other words, both treatment and control groups have a common support from each rating band. This way, I am not identifying the returns to energy-efficient properties, but to neighbourhoods with potential to make returns from improvements captured by the signalling of energy efficiency through property listings. Finally, the presence of inefficient properties within treatment neighbourhoods provides incentives to make improvements, and variation enough to identify neighbourhood effects.

3.3.5 Identification of Neighbourhood Effects

Any price discontinuity at the border should reflect a premium associated with locating within an efficient property cluster. In using property clusters to identify this premium, a number of assumptions should be made to ensure their exogeneity. Firstly, with regards to the cluster's timing, which acts as the shock. Standard RD approaches require an exogenous event that modifies behaviour, but often this is difficult to find in spatial RD.²⁴ Exogeneity in this study requires that clusters are formed independently of other characteristics, meaning they cannot be manipulated by households. I exploit the random timing of household listing behaviour across the neighbourhood, leveraging the fact that nearby properties must i) list in large-enough numbers, ii) have similar energy efficiency scores, and iii) must be located within spatial proximity. Owing to these three constraints placed on the algorithm, the ability of households to manipulate their inclusion within a property cluster should be limited. Secondly, cluster borders should not be correlated with other factors that could confound property prices. For example, observing discontinuous levels of education between households at the border could affect property prices through unobserved mechanisms related to income or property quality. The identification assumption requires that households on both sides of the border are indistinguishable in terms of their characteristics and surroundings – differing only in their energy efficiency levels. This would make buyer attentiveness to EPCs across the neighbourhood a necessary condition for identifying the price premium (otherwise, one would observe no price differences). A downside of the data is that I cannot observe individual household attributes. To circumvent this, I use proxies to disprove the fact that the formation of clusters is driven by inherent differences in these locations. Efficient property clusters that contain different property compositions to bordering locations by design will house households of differing characteristics. I present a set of RD estimates for a handful of neighbourhood characteristics, covering the ownership rate, qualifications, occupation and age demographics, in Appendix B.18. Estimates show that there are no discontinuities in any of the measures at the border of efficient clusters, suggesting that the results do not capture differences in demographics. Instead, balancing tests at the border (see Table 3.1)

²⁴Mainly because spatial policies would induce residential selection that confounds estimates.

suggest that this is driven by energy efficiency alone.

Thirdly, beyond neighbourhood characteristics, we should also not observe structural differences between properties. If this were the case, one may attribute the premium to better quality properties within efficient clusters and not from differences in residential energy efficiency. Summary statistics presented in Table 3.1 show that this is not the case, as properties are very similar in their size, number of rooms, and window glazing. We also observe similarities in build type across the border (see Appendix B.14). In defining clusters, I am confident then that they accurately capture differences in energy efficiency scores and not differences in demographics or structural characteristics. This removes the threat of factors that may confound property prices, meaning that the clusters should be exogenous in their size and location across space. Finally, given that differences at the border arise from energy efficiency scores alone, I control for the property-level current efficiency scores so as to avoid issues related to omitted variable bias. This allows me to capture how households value location within a cluster alone.


I show that these exogeneity conditions hold and deliver an exogenous treatment through which neighbourhoods can be separated based on their energy efficiency. I use this treatment to identify two different types of neighbourhood effects: firstly, neighbourhood effects in the market for property purchases. A necessary condition to identify this effect is that prospective buyers are attentive to EPCs across the neighbourhood, such that they can distinguish the efficient property clusters from the regular neighbourhood. Exogeneity conditions satisfied, there would be no other reason for buyers to choose the property cluster. Evidence of a price premium for locating within a property cluster would then suggest a neighbourhood effect arising from the mass-listing of efficient properties. Identification therefore requires that buyers use EPCs to source information regarding the energy efficiency scores of nearby properties and use this to select into efficient property clusters. Secondly, I identify neighbourhood effects in the market for property improvements. I exploit the selection of buyers into these efficient property clusters as an exogenous shock to the peer group of incumbent households (those households already established within

the area). Each year of exposure to this change in peer group ought to increase their likelihood of making property improvements due to exposure to selecting households with preferences for energy efficiency. Identification of this effect rests on three assumptions: i) the random location of property clusters ensures that these people ought not differ in terms of their characteristics (exogeneity), ii) that price effects from property clusters at the border do not persist over time (excludable), and iii) that exposure to the improvements made by selecting households is important in explaining those made by incumbent households (relevance).

3.3.6 Discontinuities in Price

Figure 3.4 shows the discontinuity in prices at the border of efficient clusters. The border is defined at 0 on the x-axis, with properties to the right representing those falling within an efficient property cluster.

Figure 3.4: Price Discontinuity



NOTES: This figure illustrates the regression discontinuity plot of property prices at the border of energy efficient clusters (neighbourhoods). The x-axis measures distance in metres to this border, with negative (positive) values representing residence within the control (treated) group. The plot is conditional on a set of property characteristics, namely the sale year, floor area, and current energy efficiency score.

One can see a declining price gradient within control neighbourhoods, but one that becomes flatter as it approaches the border of the efficient neighbourhood. At the border

we see a small up-tick in prices, highlighting the discontinuous jump in energy efficient scores. This soon dies out further into the cluster. The small spike in prices suggests that we could observe a difference in prices. One would expect this to be larger once we control for property location, where we would expect to see heterogeneous price effects across the city.

3.4 Empirical Analysis

In this section I present the empirical strategy used to estimate the price premium paid for residence within an efficient property cluster. Neighbourhood selection decisions are discrete in nature; buyers either choose to locate within an efficient or non-efficient property cluster and adjust their price offers accordingly. Given that these clusters are randomly distributed across space, I am able to exploit their spatial boundaries to estimate a discontinuity in property prices. Below I set out the empirical strategy to estimate this in a way that exploits the panel data structure and the entry and exit of households to treatment (through residence in an observed efficient property cluster). Any discontinuity estimated reflects the premium buyers are willing to pay to reside within an energy efficient property cluster over-and-above that of a non-efficient one. Estimates from this approach represent the extensive margin and the main result of this section. Estimates of this result at the intensive margin are presented in Appendix [B.3](#).²⁵

3.4.1 Identifying a Premium

The discrete choice buyers make between property clusters allows this chapter to estimate a discontinuity at the precise point when the cluster becomes treated. Section [3.3](#) outlined the approach to defining property clusters in each period, creating a set of locations bounded in space. Properties listed on the market in a given year are then assigned to a property cluster based on energy efficiency scores in their surrounding area – allowing this chapter to exploit this as a treatment. I hypothesise that buyers undertake a similar exercise to locate property clusters using the energy efficiency scores reported in listing

²⁵The intensive margin measures the strength of spatial relationships in energy efficiency scores between properties. These scores are used to create property clusters and act as a proxy to measure neighbourhood effects and the degree to which this is reflected in prices.

adverts. Here I test whether we can accept this hypothesis that buyers are attentive to energy efficiency scores across their chosen neighbourhood and whether this generates a discontinuity in property prices. The identification strategy therefore relies on buyers gathering information from the market for their chosen neighbourhood. I provide evidence later in the section that buyers do in fact undertake this procedure. Gathering such information would allow buyers to distinguish with relative precision the spatial extent of property clusters. Using the bounded cluster defined in the previous section, I employ a spatial Regression Discontinuity (RD) framework that exploits turning points in energy efficiency scores, but that keeps constant all other covariates. I estimate this model on a sample of repeat sales to better control for property composition over time.²⁶ This is represented in equation in Eq (12).

$$\log(P_{it}) = \beta_0 + \beta_1 T_{it} + \beta_2 \delta_{it} + \beta_3 X_{it} + \psi_i + \phi_t + \epsilon_{it} \quad (12)$$

The RD framework differs from the canonical model in that the running variable captures spatial distance from the treatment border δ allowing the regression to place more weight on properties closer to the border of the property cluster. Distance to the border is centred at zero, meaning that the border of the property clusters act as the treatment. Treatment captures the sale of a property i at time t by cluster type such that T takes the value 1 if a property sells in an efficient property cluster – when $\mathbb{1}[\delta_{it} > c]$ – and 0 otherwise. The coefficient β_1 captures the price premium paid to reside within an efficient property cluster over-and-above residing outside on the periphery. Exploiting the border as a treatment is possible given randomness in their location, as characteristics on both sides of the border should be balanced. To ensure this is the case, I include a vector of property-level controls X . Included in this is a control for property floor space, owing to a marginal difference in property size outlined in Table 3.1. I further control for current energy efficiency scores at the property-level given that property clusters are defined on the back of this and the inherent differences at the border. Marginal differences exist in the potential of properties at the border identified in the balancing tests. I augment

²⁶This approach limits over-representing certain property types in different years (i.e. detached, semi-detached, terraced).

the regression with a control for unrealised efficiency to capture differences in investment potential. This signals how buyers value a property's investment potential. Including these covariates captures most property-level variation in characteristics. Included in the regression are a series of dummy variables controlling for the year of sale ϕ which capture variation in the property market. Using a repeat sales approach further allows me to capture unobservable changes made to a property outside of those related to energy efficiency whilst further allowing the inclusion of a property fixed effect ψ to control for time-invariant characteristics between properties. Capturing within-property variation excludes properties whose treatment status does not change. Estimates of the premium thereby capture entry and exit from treatment and estimates the premium upon entering. The battery of controls used should capture all alterations made to a property over time, thus allowing a more accurate estimate of the price premium whereby the only variation arises from the change in treatment status. Property clusters are defined using data from property listing adverts, therefore one would expect the premium to coincide with changes in information from listings. A number of robustness checks are undertaken to prove this later in this section. Finally, popular property search portals began advertising EPCs within listing adverts during 2010. Properties treated before this point are removed from the analysis to ensure continuity in the information source available to buyers.

Leveraging variation in listing adverts to define property clusters, I am confident that their size and location are as good as random and cannot be manipulated by buyers. A multi-step process is undertaken to ensure that the borders defined under the cluster creation process are exogenous. Firstly, property clusters are formed based on the characteristics of proximate properties. The clustering algorithm generates a measure of spatial correlation in energy efficiency scores for each property based on the scores of their neighbours. This ensures that a household cannot manipulate their own measure. Properties that lie within 500m of each other are considered neighbours, with those of increasing distance receiving less weight in the spatial correlation measure. It is therefore the characteristics of one's most proximate neighbours that determine inclusion within a property cluster. Restricting the algorithm in this way means that households cannot

free-ride off of the characteristics of the wider neighbourhood. Only spatial correlations at a fine, granular level should be captured in this case. Secondly, to circumvent issues of spatial correlation in listing behaviour, I set a threshold requiring that each property have at least 10 unique neighbours from which to identify spatial clustering, with all selling in the same year. This feature negates the possibility that a group of properties choose to list their property in the same period for mutual benefit. Households cannot game the process under this approach by falsely signalling the quality of a neighbourhood. Besides, selling a property is both time and resource intensive. Recent evidence from Zoopla suggests that it takes the average property 185 days to complete a sale (from time of listing to selling).²⁷ It is unlikely that this approach identifies such properties, and should only capture properties who are similar in energy efficiency score. Finally, in the unlikely event that the first two criteria are not satisfied, it must be the case that properties are of similar energy efficiency score and that this be significantly different from that of the expected score. Manipulation would require that all properties be of similar score to the wider area – requiring that significant home improvements be undertaken. This also seems unlikely, given that all homes within the area would have to be on board. Border manipulation could also be the present if the identified properties were all newer-builds (i.e. properties built with the same specification and with lesser time for energy-efficient installations to depreciate). I show in Appendix B.16 that these properties are not over-represented in the sample and provide reasonable evidence that borders are not manipulated by sellers. It is therefore highly unlikely that sellers could alter their inclusion/exclusion from the treatment.


A series of robustness checks are undertaken to validate the mechanism behind any discontinuity. Firstly, I show in Appendix B.2 that for the majority of efficient property clusters defined over the sample period, the borders do not overlap with administrative boundaries – specifically, those of London boroughs (districts). While this is generally true, there are few instances where this does occur. The empirical strategy should capture this through the use of location fixed effects (see later). As a further robustness check, and to show that such an overlap does not confound the main result, I re-estimate the al-

²⁷<https://www.zoopla.co.uk/discover/selling/how-long-does-it-take-to-sell-a-house/>

gorithm using property listings *within-borough*. Doing so avoids defining property cluster that cross borough boundaries and contains the borders within the respective borough. I detail this approach in Appendix B.2 and show that this chapter’s main result does not change.

I show in the previous iterations that, even after changing the approach to defining clusters, the results from the RDD are non-biased. I later test the robustness of the information mechanism that operates through information reported in listing adverts. I achieve this by delaying the listing channel and testing again for the presence of a premium. Delaying the listing channel involves exploiting the gap in time between requesting a new EPC and listing on the market. Doing so produces a sample of properties that have an EPC score that is not visible to the market at the time of the request (as the property is not listed). I use such properties to re-estimate the clustering algorithm and redefine property clusters, thereby ensuring no information is available to buyers. If the previous results are driven by the information channel, removing this ought to remove any price premium for efficient property clusters. I re-run the clustering algorithm using properties that request a new EPC but do not sell within a year. Running this for the year of the EPC request ensures that information for such properties has not reached the market at the time property clusters are defined. Properties are requesting new EPCs under both property cluster definitions – the only difference being that these properties are not listing on the market and therefore buyers cannot distinguish their locations. This should mean that buyers have no information about the properties and potential property clusters of efficient properties.

Figure 3.5: Efficient Clusters Across All Years: Sale vs No Sale



NOTES: Clusters are created in the same way by merging properties based on their z-scores and employing a bounding geometry to envelope hotspot properties. No sale clusters differ only in that the properties used to define them do not sell within a year of obtaining an EPC. This ensures properties are not listed on the market at the time of cluster creation.

If no significant discontinuities exist at the border, one can speak to two points: i) buyers are attentive to the energy efficiency scores obtained through property listings at a neighbourhood level, and ii) properties within efficient clusters do not differ in their amenity level relative to their peripheral neighbours.²⁸ The latter is an important validation test in proving the signalling mechanism at play. Figure 3.5 presents the efficient property clusters for both specifications across all years in the sample. Property clusters defined under this approach display a good deal of overlap with those defined under the original approach, illustrating the robustness of the approach in capturing the true efficient property clusters. I can be confident that the robustness check is capturing properties in the same neighbourhoods with similar characteristics to the main specification. Any differences in discontinuity can be directly attributed to the energy-efficiency signalling occurring within listings, not the general amenity level.

²⁸If they differed in amenities, we would still observe a price premium that was not the result of energy efficiency differentials.

3.4.2 Premium Estimates

Estimates of the premium are presented in Table 3.2 and are captured by the RD term $\mathbb{1}.\text{Cluster}$. I calculate the bandwidth for the RD in line with Calonico et al., (2014) who develop a non-parametric, data-driven method to estimate the optimal bandwidth. This procedure returns a bandwidth of 112 metres, implying that the effects are hyperlocal. This may raise concerns that specific local characteristics are driving changes in market activity at the border. One may hypothesise that property clusters are more attractive for reasons unknown that change over time. A number of papers find evidence that buyer location decisions are driven by the attractiveness of an area (Koster et al., 2014; van Duijn and Rouwendal, 2021; Guerrieri et al., 2013; Autor et al., 2014). To circumvent other confounding factors, I augment the spatial RD model with $\text{Location} \times \text{Year}$ fixed effects in Panel B and present the results alongside those for the baseline model in Panel A. Location is captured at the Middle Layer Super Output Area (MSOA), a Census-defined geography that partitions districts into smaller areas comprising between 5,000-15,000 individuals. This acts as a dynamic control for neighbourhood characteristics that capture the area's general attractiveness. The temporal nature of the control covers changes in both the provision of amenities and changes in property quality at the neighbourhood level, as one could imagine that gentrification effects arise from neighbourhoods that have undertaken home improvements en masse. Reverse causality between amenities and property quality in an area are corrected for using the aforementioned fixed effects approach. Re-estimating this model, capturing instances where buyers may select into these neighbourhoods for reasons outside of preferences for energy efficiency, ensures the resulting estimates are not biased.

Panel A captures the baseline estimates for the RD including property and year fixed effects. Column 1 estimates this premium without controlling for unrealised efficiency – a proxy for a property's investment potential. Including other property controls captures the property-level determinants of prices, meaning the premium reflects a willingness to pay for residence within a cluster as opposed to for the property itself. Estimates show that buyers pay a premium 2.5% premium to reside within an efficient property clus-

ter over-and-above that of an inefficient cluster. This is significant at the 1% level and translates to a monetary value of around £9,500 for the average property residing in an efficient property cluster. This is economically significant and reflects the willingness-to-pay to reside in these locations. One would expect such a large premium to reflect an expected benefit to household. Returns to an increase in the current energy efficiency are significant and positive, highlighting that buyers increase their price offers in return for more efficient properties. This signals buyer's preference for properties that have made energy efficiency improvements. I augment the regression with the property-level measure of unrealised efficiency in Column 2 to capture household's preferences for making improvements and to explain how buyers perceive their expected investment returns. Results from my previous chapter suggest that buyers decrease their bids in anticipation of undertaking and financing improvements. Running this regression on the population sample finds the same reduction in bids (see Appendix B.6). Around the border, an increase in unrealised efficiency between sales actually leads buyers to *increase* their bids for housing.

Paying more for properties in poorer condition, but that have potential, would suggest that buyers anticipate higher returns on their investments within efficient property clusters. This explains a small part of the premium paid for the neighbourhood – reducing the coefficient marginally to 2.4% which is significant at the 5% level. Evidence from the RD suggests that buyers are selecting into efficient property clusters in order to chase higher returns on investment. I come back to this later in the chapter. Alternative tests controlling for the impact of amenities are presented in Panel B. The coefficients are of the same sign to those in panel A, but are smaller in magnitude. Column 1 suggests that the premium is positive and significant, around 1.7% above that of contiguous inefficient property clusters. Including unrealised efficiency in Column 2 slightly lowers the coefficient to 1.6%. Both premium estimates are significant at the 5% level and reflects a willingness-to-pay for residence within an efficient property cluster of roughly £6,100. Expected returns to investment remain positive but become insignificant in this specification, perhaps owing to the smaller sample. Put together, controlling for the location's attractiveness and access to amenities is important, otherwise it introduces positive bias.

That being said, the estimates are robust to controlling for these factors and remain significant.

I undertake placebo tests in Columns 3-4 to determine the accuracy with which buyers can identify an efficient property cluster. I re-estimate the spatial RD after swelling and shrinking the radii of each property cluster by 50 metres. This alters the size of the clusters and the resulting location of the borders. Any discontinuity should disappear if buyers can accurately determine the true location of efficient property clusters based on property listings. In both panel A and panel B, this is the case. The discontinuity coefficient is insignificant in all specifications and either turns negative or falls close to zero. This suggests that the border is well-defined, as the price premium dissipates quickly with distance from the true border. It further suggests that buyers are attentive to energy efficiency scores of properties listed around these neighbourhoods. Estimates for the effect of both current and unrealised efficiency remain generally unchanged, pointing to similarities in preferences at either side of the border. This provides suggestive evidence that the premium arises as a result of the energy efficiency reporting in listing adverts and the impact this has on buyer perceptions of an area.

Table 3.2: Estimates of the Price Premium for Efficient Clusters

	(1)	(2)	(3)	(4)
	Discontinuity	Estimates	Discontinuity	Robustness
Panel A: Baseline RD				
1.Cluster	0.025** (0.009)	0.024* (0.009)	-0.021 (0.020)	-0.001 (0.010)
Unrealised		0.003*** (0.001)	0.002* (0.001)	0.003* (0.001)
Current	0.007*** (0.001)	0.008*** (0.001)	0.007*** (0.001)	0.010*** (0.001)
Property FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Adj. R^2	0.784	0.786	0.755	0.792
Observations	8,871	8,871	9,346	8,193
Panel B: Baseline RD w/ Location×Time Fixed Effect				
1.Cluster	0.017* (0.008)	0.016* (0.008)	-0.015 (0.024)	-0.007 (0.009)
Unrealised		0.001 (0.002)	0.001 (0.002)	0.001 (0.002)
Current	0.007*** (0.001)	0.008*** (0.001)	0.007*** (0.001)	0.008*** (0.002)
Property FE	✓	✓	✓	✓
Location×Year FE	✓	✓	✓	✓
Bandwidth	112	112	112	112
Adj. R^2	0.958	0.958	0.945	0.961
Observations	2,576	2,576	2,004	2,921

NOTES: Standard errors clustered at the LSOA level in parentheses. Treated clusters are those defined as energy-efficient under the Getis-Ord calculation. The regression includes controls for the year of sale and follows a second-order polynomial. Flat's are excluded from the analysis. I swell and shrink the borders by 50m and re-estimate as a robustness checks in columns 3 and 4, respectively. Bandwidth measured in metres from the border of the treated clusters. + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

3.4.3 Robustness Checks

I undertake a series of checks to validate the robustness of the estimates. The previous results speak to the discontinuity in prices induced by information reported in property listing adverts. To show that it is the information contained within these adverts that drives the effect, I undertake a set of robustness checks that estimate the discontinuity after removing the listing channel. I remove this channel in two different ways. First, I run

the RD using properties that have been reassigned to efficient property clusters outside of the treatment window. In other words, I extract the properties located within the same property clusters, only before/after the cluster's treatment window. The entry-exit style treatment means that these clusters are only treated when energy-efficient properties are listed on the market. When buyers do not observe such properties through listings, the information channel should be removed and the premium along with it. If it truly is listings adverts that propagate the premium, then a discontinuity should only arise in the year of reporting. Assuming that characteristics stay the same over time, we ought not observe a premium. A potential issue with this approach is that popular property portals such as Rightmove and Zoopla offer services that allow customers to view the listings for historic sales. Such a feature provides a threat to this identification strategy as one may suspect that buyers form neighbourhood expectations based on previous sales, threatening the importance of information reported through property listing adverts in the present year. While it seems implausible that buyers would study previous sales in such a detailed way, I propose a second approach to remove the information channel. In this approach, I tweak the property cluster definitions applied in Section 3.3 in two ways: by creating property clusters i) based on the year of a property's EPC inspection, and ii) removing properties that request a new EPC within a year of selling the property. The former captures each property's energy efficiency scores at the time at which they make home improvements, while the latter ensures that, when capturing this, the properties are as-good-as guaranteed to not be listed for sale on the market. Buyers have no way of sourcing information on energy efficiency scores under this approach while simultaneously removing the influence of previous listings on the estimates. In fact, the property clusters created here are the true energy-efficient clusters as they do not rely on households selling to release information to the market. Proof of no discontinuity under this approach are more robust. I overlay these clusters with those of the original definition in Figure 3.5.

Table 3.3: Placebo Test Estimates of the Price Premium for Efficient Clusters

	(1)	(2)	(3)	(4)
	Out of Treatment Window		Clusters using Non-Listers	
1.Cluster	-0.029 (0.024)	-0.005 (0.028)	-0.025 (0.022)	-0.058 (0.037)
Unrealised	0.001** (0.000)	-0.000 (0.000)	0.002 (0.003)	0.007 (0.005)
Current	0.006*** (0.000)	0.005*** (0.000)	0.011*** (0.002)	0.011*** (0.002)
Property FE	✓	✓	✓	✓
Year FE	✓		✓	
Location×Year FE		✓		✓
Bandwidth	112	112	112	112
Adj. R^2	0.836	0.967	0.602	0.890
Observations	16,446	13,425	4,182	285

NOTES: Standard errors clustered at the LSOA level in parentheses. Regression follow a second-order polynomial. Flat's are excluded from the analysis. Bandwidth measured in metres from the border of the treated clusters. The RDD is re-estimated on two different samples: Columns 1-2 capture properties within efficient clusters that sell in years outside of the treatment window. Columns 3-4 redefine clusters according to Eq.(11) using properties that do not sell within a year of the EPC inspection – removing the information channel. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Both approaches produce remarkably similar definitions with regards to the locations and size of the property clusters, providing a strong counterfactual for the exercise. The tests outlined in this section should provide adequate evidence that property listings create a premium. Together, these robustness checks shed light on the drivers of a price premium and motivate the idea that buyers select into neighbourhoods based on energy-efficiency-related characteristics. Results from two sets of placebo tests are presented in Table 3.3. Column 1-2 present estimates for properties located in efficient property clusters but outside of the treatment window.²⁹ Odd-numbered columns contain the baseline results while even-numbered columns include Location×Time fixed effects. The results show that the premium is no longer significant, nor is it positive in either specification or under either approach. The estimates are noisy and so their magnitudes are interpreted with caution. This suggests that the premium wanes over time and that reporting plays a big role in the attraction of buyers to a cluster. Positive coefficients on unrealised efficiency shows that

²⁹For example, if the cluster is identified as efficient through listings in year $t = 2015$, I only consider properties that sell in periods outside-of-treatment in years $t \neq 2015$.

buyers still expect investment returns, but not to the same degree. This would explain the lack of premium. Columns 3-4 present estimates for clusters defined using properties not listed on the market. Estimates are similar to that of the first placebo test. The premium becomes negative and insignificant in both Columns 3 and 4. Point estimates are close to zero in both cases and the signs remain the same on both the current and unrealised efficiency coefficients. Both placebo differ from the main results only in that they exclude the reporting mechanism. One can be confident therefore that listings create the premium within energy efficient property clusters as buyers learn from the information reported within them.

Estimates from Table 3.2 return positive estimates for unrealised efficiency suggesting that buyers pay more for properties with greater investment potential. Apartments are excluded from these estimates on account of the difficulties posed in undertaking energy-efficiency improvements. This would bias estimates of the premium if it truly is the expectation of greater returns to investment that buyers are paying for. I present estimates of the price premium using a sample of apartments only in Table 3.4. The results confirm my hypothesis. Point estimates of the premium are either negative and/or close to zero and statistically insignificant in all instances. The estimates are noisy, possibly owing to limited variation in characteristics within blocks of flats. Significant coefficients for current energy efficiency show that buyers still value energy efficiency improvements that deliver benefits in the present (i.e. through reduced bills and improved thermal comfort). However, we do not see any evidence that unrealised efficiency is valued by buyers purchasing flats at the border. While buyers value energy efficiency, they do not anticipate returns to undertaking investments and therefore reduce their bids for properties that are further from their energy efficiency potential. This provides suggestive evidence that these households are constrained in their scope for making improvements in apartments and gaining adequate price returns. To summarise, estimates in this section provide evidence of two interesting features of the market: that i) that selection into efficient neighbourhoods occurs, as illustrated by the associated premium, and ii) those selecting into these neighbourhoods do so under the expectation of higher returns to in-

vestment, illustrated by the premium paid for poorly-efficient properties.

Table 3.4: Apartments: Estimates of the Price Premium for Efficient Clusters

	(1)	(2)	(3)
	Discontinuity Estimates	Discontinuity	Robustness
Panel A: Baseline RD			
1.Cluster	0.007 (0.009)	-0.012 (0.019)	0.006 (0.012)
Unrealised	-0.002 (0.003)	-0.001 (0.003)	-0.005 (0.003)
Current	0.003* (0.001)	0.004*** (0.001)	0.003* (0.001)
Property FE	✓	✓	✓
Year FE	✓	✓	✓
Adj. R^2	0.760	0.755	0.765
Observations	7,737	7,998	7,487
Panel B: Baseline RD w/ Location×Time Fixed Effect			
1.Cluster	-0.001 (0.008)	0.008 (0.019)	-0.006 (0.010)
Unrealised	-0.003 (0.004)	-0.002 (0.004)	-0.008* (0.004)
Current	0.002 (0.002)	0.003* (0.001)	0.001 (0.001)
Property FE	✓	✓	✓
Location×Year FE	✓	✓	✓
Bandwidth	169	169	169
Adj. R^2	0.957	0.954	0.959
Observations	3,052	2,893	3,269

NOTES: Standard errors clustered at the LSOA level in parentheses. Treated clusters are those defined as energy-efficient under the Getis-Ord calculation. The regression includes controls for the year of sale and follows a second-order polynomial. Flat's are excluded from the analysis. I swell and shrink the borders by 50m and re-estimate as a robustness checks in columns 3 and 4, respectively. Bandwidth measured in metres from the border of the treated clusters. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

3.5 Mechanisms

The previous section provided evidence to support the presence of a premium for efficient property clusters in the period of treatment. This is shown to dissipates in periods after the treatment. In this section, I provide further evidence that energy efficiency reporting

in listings generates this premium and that this creates neighbourhood effects among incumbent households. The former is achieved by illustrating the exact timing of when the premium arises. The latter I show to work through an exposure mechanism whereby incumbent households (never movers) are exposed over time to new peers with preferences for energy efficiency investments. The premium is found to reflect buyer expectations of higher investment returns, resulting in a higher probability of making investments in these neighbourhoods and higher returns for previously invested-in properties.

3.5.1 Identifying Premium Timing

Section 3.4 provides evidence of a discontinuity in prices owing to a premium for residence within efficient property clusters. Removing the information mechanism removes the premium. While this sheds light on the hypothesised mechanism, it does not illustrate the exact time a premium arises and whether it is contemporaneous with energy efficiency reporting within listings. Without this, one could speculate about the mechanism or even favour other prognoses. To tease out the timing of the premium is to uncover whether buyers are attentive to neighbourhood-level energy efficiency scores, whether they truly select into specific neighbourhoods on the back of them, and whether this is reflected in prices. Property clusters are defined annually meaning that properties could be treated in each year. Standard event study or Difference-in-Difference approaches to measuring price changes across time are therefore subject to bias. I instead employ an event study approach with multiple treatment periods developed by Callaway and Sant’Anna, (2021) to determine the timing at which the premium arises. Time is normalised (centred) to the period of treatment – time zero when the property cluster is first treated – while time for the other observations is measured as the number of years between sale and treatment.³⁰ Property prices are then observed before, during and after treatment, and can be visualised on a continuous time scale. The treatment group comprises properties that fall into a future- or previously-observed efficient property cluster – with properties treated at time zero being those that fall into an observed efficient cluster. Treatment occurs in every year, and properties remain treated after their first sale in an efficient

³⁰For example, a property in an efficient cluster that becomes treated in year 2015 but sells only in 2013 would be assigned a time-period $\tau = -2$.

property cluster. This is justified by the methodology used to value properties on the market, whereby RICS surveyors take as a property's starting valuation the value of recent nearby sales of similar quality (RICS, 2022). Estimates produced under this approach are notoriously sensitive to how the control group is defined. The inability to observe a panel of sales, given the infrequency of transactions, creates a further complication. Results from the previous section evidence buyer's ability to locate efficient property clusters. I exploit this result here and use as the control group properties that lie outside but on the periphery of an efficient property cluster. I find no previous evidence to suggest buyers pay a premium for these properties given that they possess no intrinsic added value to selecting households. This group should provide a strong counterfactual given that covariates are balanced at the border and the only other difference between the two groups arises at time of treatment when property listing adverts report energy efficiency information. If buyers are truly chasing returns to their property investment and listing adverts provide the source of information, then prices at the border should not differ in years prior to the cluster being treated (observed). I test this by constructing a repeated cross-section of property sales that occur within 50 metres of a property cluster border (for all years in the sample). Various bandwidths are tested and the results were found to be hyperlocal. Dissipating treatment effects after this threshold means that I choose a bandwidth of 50 metres. Given that RICS assessments produce property valuations using anchor properties within these neighbourhoods, one would expect to see persistence in the premium at the border as sales in subsequent years adopt the new baseline price at valuation. If this is not the case and prices between the two groups do not diverge, this would provide evidence of price spillovers at the border onto control properties. Appendix B.11 provides a graphical illustration of the treatment and control groups while Appendix B.10 provides intuition behind the construction of the time variable. Property prices are regressed on a series of time dummies to compare differences between the treatment and control group over time both before and after the period of treatment. The estimating equation is presented in Eq (13).

$$\log(P_{it}) = \alpha + \sum_{\tau=-5}^{\tau=3} \beta \cdot \mathbb{1}(E_{it} = \tau) + X'_i \beta + \phi_i + \epsilon_{it}, \quad (13)$$

Where the event E_{it} occurs when the time of sale t is equal to the time of treatment T such that $\tau = t - T_i$. Captured in τ is a property's year of sale relative to its time of treatment.³¹ In other words, treatment occurs at $\tau = 0$ when buyers observe an efficient property cluster from properties listed on the market. I regress the log of prices on this treatment indicator which captures properties selling within an efficient property cluster both within and outside of the treatment year for each period within the sample window. To increase the sample size within-year, I restrict the sample to include properties that sell five years prior and three years after a treatment window. This provides a large enough sample window to observe pre-trends and the post-listing effect, respectively. Included in the regression are property controls X_i that capture the current and unrealised energy efficiency score, the total floor area, the year of sale, as well as distance from the efficient property cluster and location controls. I further control for the property type and its year of construction.³² There should be little variation in built form and year of construction at the border, but inclusion in the regression should remove any bias arising from changes in the composition of properties listed over time. The bandwidth around the borders are very small geographical areas and should capture any further differences in property composition and location characteristics. The results should therefore identify price differences between the treated and control groups arising from listings information alone – not property characteristics. If buyers respond to energy-efficiency reporting within listings, price differences between the two groups should be significantly different and positive at $\tau = 0$ when an efficient property cluster is observed and insignificant in all other years ($\tau \neq 0$). A similar Difference-in-Differences approach is undertaken by Ghosh et al., (2024) who find that making it mandatory to supply buyers with an energy performance certificate at time of sale generated a premium. Possession of an up-to-date EPC became a legal requirement to sell from 2008. Reporting of these certificates within property search portals only started taking place in 2010, however. For this reason, I choose to drop properties that sell nearby property clusters created prior to 2010. I later re-include these clusters and re-estimate the event study to see whether their inclusion affects the observed Av-

³¹For example, $\tau = -2$ for a property residing within a treated cluster in 2015 but selling in 2013; $\tau = 2$ if treated in year 2015 but selling in 2017.

³²Flats are removed from the analysis like in previous estimates. Property type here refers to the built form, such as detached, semi-detached or terraced.


verage Treatment effect on the Treated (ATT); results are presented in the Appendix [B.13](#).

3.5.2 Premium Timing Estimates

Estimates from the event study are presented in Figure [3.6](#). Large confidence intervals towards the poles of the time axis are symptomatic of the DiD approach with multiple time periods. However, moving closer towards the treatment date at $\tau = 0$, we observe less noisy estimates. The parallel trends assumption is upheld with next-to-no differences in prices between the two groups prior to treatment, as expected. It also proves that coefficients are well-balanced, so that price differences do not arise as a result of discontinuities in other characteristics at the border. The figure shows a spike in prices at time $\tau = 0$ when listings adverts are able to observe clusters of energy efficient properties. Property prices for treated units in this time period increase above that of the control group, generating a price premium. The only difference between periods $\tau < 0$ and $\tau = 0$ is the reporting of energy efficiency characteristics within listings, meaning the premium can be directly attributed to the information spillover. While noisy and insignificant, the figure demonstrates a slight persistence in the premium in the subsequent period before prices in the two groups return back towards parity. One could hypothesise this to be the result of spillovers onto control properties. This is not the case, however. A simple plot of property prices for treated units in Appendix [B.12](#) shows that, relative to the baseline year in which clusters are treated, prices begin to decline after the treatment period $\tau = 0$. If the price of treated properties begin to fall after treatment and price differences between the treated and control groups converge towards zero, then properties in the control group must experience no change in prices. I can be confident therefore that price premium does not spillover across the border and that buyers are simply purchasing property at higher prices based on signalling coming from the market. Robustness checks are implemented to check whether the results are sensitive to changes in the property clusters included in the sample. Property portals reported EPCs in listings starting in 2010. Replicating the event study removing properties treated a year before and a year after this period changes the results very little. Results from this exercise are presented

in Appendix B.13. The only notable change belongs to the post-period results where the ATT is marginally higher. Lower uptake of EPC in the early years of the EPC initiative likely dampens the results slightly and explain why the results are larger when incorporating properties treated a year later. Because of sample size issues I cannot estimate a placebo test using clusters prior to 2010, but this approach provides suggestive evidence that there were limited effects of energy efficiency reporting on prices prior to 2010.

Figure 3.6: Event Study Estimates of the Price Premium by Period



NOTES: This figure presents the results from the event study. Properties selling within an efficient cluster at any time represent the treated group, while the control group consists properties selling on the edge of these clusters. A bandwidth of 50m has been chosen, capturing properties either side of the efficient cluster (treatment) border. Property controls include the current efficiency score, total floor area, property age, property type, and build type. Tests of the pre-trend return a p-value = 0.977 in favour of accepting the null hypothesis H_0 : that all pre-treatment ATT's are equal to zero.

3.5.3 Identifying Returns to Improvements

Having shown that energy efficiency reporting within listings is generating a price premium – and with it, selection of households with preferences for energy efficiency into efficient property clusters – I now explore the impacts this has on the uptake of improvements. Two types of households exist: those that move property and select into efficient property clusters, and those that do not move property. The latter are referred to as incumbent

households and are defined as those that were already residing within an efficient property cluster at the time it became treated. I explore the impact of undertaking improvements on property prices for both groups. If we take the view that incumbents are more likely to learn from selecting households and upgrade their properties, the capitalisation literature would argue that this results in higher prices (Fuerst and Warren-Myers, 2018; Aydin et al., 2020; Sejas-Portillo et al., 2020). An alternative strand of research would suggest that the mere presence of selecting households generates neighbourhood effects that increase prices (Guerrieri et al., 2013; Koster et al., 2014; González-Pampillón, 2022). The latter can be easily estimated by including a variable in the regression denoting whether or not the property resides within an efficient property cluster at time of sale, but the decision to make improvements is endogenous. This could simply reflect the fact that wealthier households reside within these neighbourhoods and are more likely to make upgrades in the first place (correlated effects, a la Manski, 1993). To separate the two explanations, I employ an instrumental variable approach that exploits spatial variation in the location of efficient property clusters to explain household improvement decisions. Results from Section 3.4 illustrate that efficient clusters are exogenous and are not subject to manipulation. I use these to create an Instrumental Variable that measures the number of years a property is exposed to an observed cluster. A mathematical representation of this can be seen in Eq.(14)

$$expos = \sum_{\substack{t < \tau_2 \\ t > \tau_1}} T_{it} \quad \forall t \in [\tau_1, \tau_2] \quad (14)$$

To estimate the 2SLS regression I exploit the same panel of properties used in the border discontinuity design in Section 3.4. Here I collapse the panel to a first differenced dataset, given that exposure to efficient property clusters cannot be measured before the first observed sale. Exposure captures the number of years each property i is exposed to a newly observed cluster (defined in each period t) between its first τ_1 and second sale τ_2 . This approach takes the extensive margin and develops it to measure an intensive margin (years of exposure) of treatment. In other words, the IV captures the number of years a household is exposed to selecting households with preferences for energy efficiency.

Exposure to selecting households overtime should impact households only through their likelihood of making improvements, not through prices directly. I use this to instrument for the endogenous decision of the household to undertake improvements, proxied for by the presence of an EPC ($\widehat{EPC}_i$). Proxying for improvement decisions rests on the assumption that rational occupiers request a new certificate after improving their property's energy efficiency. New certificates are more likely to be requested at time of sale, but this does not affect the estimation, given that properties are only observed at time of sale. If incumbents are conducting upgrades, then prices ought reflect this capitalisation. I set out the IV specification in Eq (15).


$$\log(\Delta P_i) = \beta_0 + \beta_1 \widehat{EPC}_i + \beta_2 \mathbb{1}T_{i\tau_2} + \beta_3 \delta_i + X_i' \beta + Z_i' \beta + \phi_i + \epsilon_i \quad (15)$$

The variable EPC is binary and takes the value 1 if household i request a new EPC between sales. This is regressed on the exposure instrument in the first stage such that $\widehat{EPC}_i$ isolates only the exogenous components that drive improvements. This essentially returns the likelihood a property makes an energy efficiency improvement given its level of exposure to an efficient property cluster. The IV specification means that the coefficient β_1 captures the change in prices between sales for those making improvements over-and-above those that do not. Estimates are conditional on other variables, such as whether or not the household sells within an active cluster $T_{i\tau_2}$. The treatment in this case is the same as in the RD framework. I further control for the same property-level characteristics X_i as in the RD: current energy efficiency and total floor space at the time of first sales, as well as the change in each between the two sales. A control for the distance from the efficient cluster border is also included to ensure that buyer perceptions about the location/size does not influence the estimated coefficient. I finally include controls for the year of sale ϕ_i and location Z_i to remove any trends in local property markets. Location is captured in the same way as in the fixed effects model at the MSOA level.

Given that efficient clusters are defined by property listings, the exposure measure by design captures property churn within the locality. Based on the previously presented empirical strategies, I show that buyers pay more to reside within these neighbourhoods

based on energy-efficiency reporting – suggesting that churn is driven by buyers with preferences for energy efficiency.³³ In measuring exposure, I therefore capture the exogenous changes in the composition of one’s peer group and the exposure to households with greater preferences for energy efficiency. Following the social norms argument around peer effects, or simply that households seek to imitate their peers, I would expect the instrument to be relevant in explaining investment decisions. I exploit the same borders produced using the approach in Eq (11) with the same exogenous characteristics. Creation of the clusters exploits random variation in the distribution of listings across space, as well as with regard to the quantity of local listings and their quality (with respect to their energy efficiency). Incumbent households should therefore be unable to influence their inclusion within an efficient property cluster or the characteristics of those purchasing within – meaning that changes in peer composition are as good as random.³⁴ I highlight the instrument’s relevance by swapping the dependent variable in the border discontinuity in Eq (12) with $\mathbb{1}.EPC_i$ – which captures whether a household requests a new EPC between sales – and re-estimating to produce estimates for the probability of upgrading at the border. Figure 3.7 presents the results.

Figure 3.7: Probability of Upgrading at the Border By Household Type



NOTES: This figure illustrates the probability of upgrading at the border of the efficient neighbourhood. The x-axis captures distance to the efficient cluster border measured in metres, with observations to the left of the border representing the control group. Property controls include the current efficiency score and the property’s total floor area. Non-selection captures properties sold in a cluster in years outside of the treatment window. Selection should not take place given the absence of the information channel.

I estimate two versions of the RDD are estimated: the first includes all properties within the baseline sample (the left-hand figure), and the second includes households residing

³³This assumes that buyers without such preferences would not pay more for properties in such neighbourhoods and would instead locate on the periphery.

³⁴Incumbent households are defined as those that do not sell in a given year

within, but not choosing to purchase within, an efficient property cluster at time of treatment (the right-hand figure). In other words, the sample used to estimate the right-hand figure captures households who choose an efficient property cluster after it has already been signalled as such. This captures households that do not select into it based on its energy efficiency. The difference between the two samples is that the left-hand figure includes those households selecting into efficient property clusters in the first sale. We can use the two figures to examine whether these individuals are more likely to then request a new EPC. The border to the efficient cluster is represented by 0 on the x-axis, with households to the right forming the treatment group. Selecting households have a higher likelihood of requesting a new EPC, with the probability jumping discontinuously at the border and increasing further. On the contrary, households who did not actively select into an efficient cluster exhibit no discontinuous effects at the border – instead, households display a lower likelihood of upgrading. This suggests that households with preferences for energy efficiency are selecting into efficient clusters and should exert some degree of influence over incumbent households in the improvements they make.

Regarding the exogeneity of the instrument, it is also unlikely to believe that years of exposure to an efficient property clusters should influence any property-level characteristics other than the decision to make energy-efficiency investments. This should be true given the selection of households with energy efficiency preferences. To guard against the indirect influence that energy efficiency investments have on energy efficiency scores, I control for current energy efficiency. Further to this, I show earlier in the event study analysis that the premium for residing in a cluster occurs contemporaneously with the reporting in listings. Effects in the preceding and succeeding periods are close to zero. This would suggest two things: i) that the premium accrues only to those selling within an observed cluster, and ii) that there are no long-lasting price effects, meaning that exposure does not affect prices for incumbent households indirectly. Because there are no indirect price effects spilling over from the treated to the non-treated sellers, the instrument should not violate the exclusion restriction that exposure is correlated with prices.

To ensure the results are robust, I re-estimate the 2SLS model inclusive of properties for whom I am sure have conducted improvements. A new EPC can affect prices through i) changes in energy efficiency scores, or ii) through the undertaking of home improvements. Changes in the former are not necessarily a product of the latter as scores may change as a result of changes in the energy assessment criteria or through unobservable factors affecting the engineering estimates (i.e. different engineers working under different assumptions). To isolate the properties with observable improvements, I extract the set of improvements recommended by the engineer at time of inspection and compare them between certificates. Improvements that appear in the first but not in the second EPC are assumed to have been implemented, allowing me to identify those properties that have made improvements. I include a dummy variable that takes the value 1 if properties have made visible improvements and 0 otherwise. Controlling for these factors ought to return a more robust estimate for the price effects of undertaking improvements.

3.5.4 Returns to Improvements Estimates

I now present the results from the 2SLS to ascertain whether buyers value the efficient property clusters themselves or the improvements undertaken within them. Estimates from the previous section showed that buyers select into these locations based on access to the efficient cluster – but is this driven by the improvements undertaken by previous owners? Table 3.5 present the results from the IV specification. Columns 1-2 present the results for the whole sample. The decision to invest in energy efficiency improvements is proxied by the request of a new EPC ($1.EPC$) and is instrumented for using the number of years exposed to an active cluster. The first stage returns the expected sign and suggests that every extra year of exposure to an efficient cluster increases the likelihood of requesting a new EPC by 5%. This matches the a-priori expectation that exposure to households in efficient clusters over time leads to a higher probability of making home improvements. The instrument is highly relevant with an F-statistic around 120 in both specifications. The instrument captures exposure to newly observed clusters (i.e. those identifiable through property listing adverts) – suggesting that it is the churn of households into the area that creates this effect. Estimates here are taken from the whole sample

of properties in the original panel, however, and includes those that do and do not move into an efficient cluster at time of treatment.

Table 3.5: 2SLS Estimates of Exposure to Clusters on Property Prices

	(1) Full Sample	(2)	(3) EPCs with Upgrades	(4)
First Stage				
Years in cluster	0.052*** (0.005)	0.050*** (0.005)	0.060*** (0.006)	0.060*** (0.006)
1.Cluster	-0.002 (0.006)	-0.001 (0.006)	-0.011 (0.007)	-0.011 (0.007)
Second Stage				
1.EPC	2.334*** (0.226)	0.764*** (0.146)	2.027*** (0.222)	0.745*** (0.139)
1.Cluster	-0.094*** (0.027)	-0.006 (0.023)	-0.077** (0.028)	0.004 (0.024)
Current	-0.009*** (0.000)	-0.001** (0.000)	-0.009*** (0.000)	-0.001** (0.000)
Δ Current	0.016*** (0.001)	0.010*** (0.001)	0.015*** (0.001)	0.010*** (0.001)
Unrealised	-0.023*** (0.001)	0.005*** (0.001)	-0.025*** (0.001)	0.005*** (0.001)
Δ Unrealised	0.009*** (0.001)	-0.004*** (0.000)	0.011*** (0.000)	-0.003*** (0.000)
Location Controls	✓	✓	✓	✓
Year Controls	✓	✓	✓	✓
Time Between Sales Control		✓		✓
F Stat	120.7	115.6	100.4	99.7
Observations	49,234	49,234	48,731	48,731

NOTES: Standard errors clustered at the LSOA level in parentheses. Years in cluster represents the instrumental variable, while the endogenous variable is binary and takes the value 1 if the property requests a new EPC and 0 otherwise. Second stage dependent variable is the change in log prices. Controls include distance from the treatment border, as well as the total floor area of the property measured in m². Columns 3-4 remove properties that have requested new EPCs but that have undertaken no home improvements, so as to improve the accuracy of measured improvements. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Estimates of the probability of making improvements therefore includes the endogenous effect (the effect of one's peers) as well as the correlated effect (the effect of being more likely to make improvements regardless), in the terminology of Manski, (1993). Though this chapter is concerned with neighbourhood effects and not peer effects, the premise

still holds. I eliminate the source of the neighbourhood effect (selecting households) and just observe the effect on the recipients (incumbent households). The purpose of the exercise at hand is not yet to estimate the neighbourhood effect, but to use this framework to disentangle the price impacts arising from neighbourhood effects with those from making energy efficiency improvements. In Columns 3-4, I restrict the sample by removing properties that have made no visible home improvement but who have requested an EPC. The probability of requesting a new EPC rises slightly to 6%, illustrating that those making improvements are influenced more by exposure to observed efficient clusters. In Columns 2 and 4 I control for the time (in days) between the two sales – but as expected, this has little to no effect on the probability of making upgrades. Both columns are taken as the preferred specifications. Estimates from the second stage do not differ greatly between samples. The dependent variable is the log of the price change between the two sales. The returns to requesting a new EPC after making improvements are incredibly large in Column 3 – around 900% larger for those requesting a new EPC. Upon controlling for the time between sales in Column 4, this falls to around 100%, meaning that price changes between sales are twice as large for households that request a new EPC following improvements, suggesting that these are valued by buyers. However, we now see that price returns to locating within an efficient cluster are insignificant and close to zero upon controlling for home improvements. While I show in Table 3.2 that buyers are able to identify the location of efficient property clusters and pay a premium for residence within them, I show here that this reflects the gains in energy efficiency delivered in these areas. This is confirmed by the positive coefficient on the change in current energy efficiency scores between sales. Buyers still value the investment potential of these neighbourhoods. Increases in the starting unrealised efficiency lead to higher prices as buyers chase investment returns from improvements. Estimates from the 2SLS model show that value accrues from the home improvements undertaken and the potential to undertake one's own. The first stage confirms that buyers selecting into the neighbourhoods – who have preferences for energy efficiency – create spillovers onto incumbent households, as evidenced in the previous section. I now turn my attention to isolating the endogenous effects to estimate the degree to which neighbourhood effects occur in propagating home

improvements.

3.5.5 Identifying neighbourhood Effects

The first stage of the above 2SLS regressions estimates the probability of making improvements for each extra year of exposure to an efficient property cluster. This gets us part way to thinking about neighbourhood effects, though using the whole repeat-sale sample in estimation raises issues regarding the reflection problem coined in Manski, (1993). The sample includes selecting households – those actively choosing residence within efficient clusters – who I have shown to be more likely to upgrade, regardless of neighbourhood effects. This makes any estimates prone to such correlated effects. To overcome this, I re-estimate Eq (15) separately for a sample of households that move into the same highly-efficient cluster, but outside of a treatment year. Households in this sample should be unable to observe the neighbourhood’s efficiency prior to relocating and should not possess any notable preferences for energy efficiency. The right-hand panel of Figure 3.7 proves that these households should not exhibit ex-ante preferences for energy efficiency and should be less likely to make improvements. Removing the correlated effect helps provide a more accurate estimate of the neighbourhood effect resulting from exposure to new peers within efficient clusters. I take the same approach to estimating the regression – calculating the years of exposure to efficient property clusters between sales for incumbent households. One would expect that living nearby households with preferences for energy efficiency would induce higher levels of energy-efficiency investment locally. I now turn to presenting the results of the different specifications

3.5.6 Neighbourhood Effects Estimates

I adopt the preferred 2SLS specification presented in Column 4 of Table 3.5 to estimate the neighbourhood effects associated with proximity to upgrading neighbours. This limits the sample of new EPCs to those that have implemented home improvements. Estimates are presented in Table 3.6 and are analogous to the first stage estimates presented in the previous 2SLS models.

Table 3.6: First Stage IV Estimates of Neighbourhood Effects in Improvements

	(1) Neighbourhood Effect Estimates	(2) Effect Estimates	(3) Baseline Estimates
Years in cluster		0.011*** (0.003)	0.060*** (0.006)
1.Cluster	0.005 (0.010)	0.001 (0.010)	-0.0070 (0.0091)
Current	0.007*** (0.001)	0.007*** (0.001)	0.001*** (0.000)
Unrealised	-0.008*** (0.002)	-0.008*** (0.002)	-0.000 (0.000)
Location Controls	✓	✓	✓
Year Controls	✓	✓	✓
R ²	0.430	0.431	0.505
Observations	7,206	7,206	48,731

NOTES: Standard errors clustered at the LSOA level in parentheses. The dependent variable is binary denoted by 1 if the property requests a new EPC and 0 otherwise. The empirical strategy mimics that of Table 3.5; please refer to its notes. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Column 3 presents the previous estimates of the probability of making investments using the sample of selecting households, taken from the first stage in Table 3.5. These figures are used as a baseline upon which to compare with the neighbourhood effect estimates. Column 1 presents baseline characteristics and their impacts on requesting a new EPC following improvements. Estimates show that presence within an efficient cluster alone has no bearing on this decision (though the estimates are noisy) – same as in the previous estimates. Crucially for the estimate of neighbourhood effects, this makes me confident that this sample does not capture correlated effects. Regardless of their precession, the point estimate is close to zero. The likelihood of requesting a new EPC also increases when a property is more energy efficient, reflected in a positive coefficient on current efficiency and a negative coefficient on unrealised efficiency. The sample considers properties making improvements, meaning those undertaking home improvements are doing so on properties that are already more energy efficient. This is interesting and suggests that improvements are concentrated within efficient clusters. Including the measure of exposure to clusters, we see that the coefficient on location within an efficient cluster effectively becomes zero. This is the same as the previous set of estimates, suggesting again that it is exposure to

the preferences of selecting households that drives the higher incidence of improvements among incumbents. Every extra year of exposure to an efficient cluster increases the probability of making improvements by 1.1%. This is statistically significant at all levels. Compared to the baseline estimates from the whole sample, neighbourhood effects account for roughly one-fifth of the effect of exposure. Correlated effects drive a lot of improvement behaviour accounted for based upon selection of individuals into efficient clusters, but neighbourhood effects still make up a non-negligible share of the overall effect. Coefficients on the control variables also do not change upon controlling for exposure, highlighting that they are not correlated. One can see that exposure is exogenous to the model, making it a strong predictor of neighbourhood effects. The churn of properties nearby therefore induces changes in local preferences and encourages incumbents to make home improvements. The presence of neighbourhood effects is important in this context, given that earlier findings suggest investments to be a leading cause of the price premium. This requires further study to understand the implications this has on sorting behaviour and any subsequent impacts on individuals being priced out of markets.

3.6 Conclusion

The energy efficiency of one's peers and the subsequent impacts on sorting, and therefore prices, is an under-researched field. This paper fills this gap by exploring how reporting of energy efficiency scores causes buyers to differentiate between different neighbourhoods based on their distribution of energy efficient properties. By exploiting random variation in property listings, I find that buyers are able to identify energy-efficient clusters through their attentiveness to efficiency scores and respond to this reporting by bidding higher for properties in efficient clusters. Buyers pay a 1.6% premium to reside within such clusters, translating to roughly £6,100. While small, this premium is directly attributable to energy efficiency reporting and represents a significant premium for properties whose structural characteristics differ very minimally from the control group on the periphery of clusters. Premiums do not exist in years when the efficient clusters are not observed, nor do they exist when one defines efficient clusters using properties not-yet-listed on the

market.³⁵ Prices are shown to vary across neighbourhoods depending on the characteristics of properties listed on the market. A series of robustness checks show that prices respond to information shock that make efficient clusters relatively more attractive for investment. This is reflected in buyer's expectations of higher returns to investments in energy efficiency. A premium in this case reflects a selection of individuals into efficient property clusters and creates an exogenous change to the peer group composition of incumbent households. Selecting households display a higher probability of making investments over-and-above that of contiguous neighbourhoods, and along with it, introduce interactions within the neighbourhoods that influence others to follow suit. Part of this can be attributed to a neighbourhood effect whereby households who do not purchase in an efficient cluster, but are exposed to households that do, observe a 1.1% higher probability of making upgrades. Conditional on such upgrades, there is no significant impact of residence within these neighbourhood themselves, meaning that increased uptake of improvements are driven by the exposure to selecting households. This invariably creates more efficient households and more efficient neighbourhoods. Prices are shown to respond to the increased uptake of improvements – induced by the actions of one's peers – as buyers seek out more energy efficient homes. But like with upgrades, prices are not influenced by residence within an efficient cluster. Though in the presence of random variation in listings and sorting based on such listings, the path of development in energy efficiency remains unknown and unpredictable.

This paper's findings raise some interesting questions related to the policy debate and areas for further research. Firstly, around Net Zero. The push to de-carbonise housing via retrofitting continues to plod on. Selection is found to play a big role in uptake based on households with preferences for energy efficiency and the spillover to nearby properties. This presents an area from which we can explore more given the UK's wealth of granular housing data and previous green initiatives, such as the Green Deal, the Green Homes Grant, and subsidies surrounding solar panel installation. Furthermore, the impact of improving uptake has consequences for local renters. This study focuses mainly

³⁵i.e. properties that have new Energy Performance Certificates (EPCs) but that do not provide information spillovers through listing.

on owner-occupied housing or rentals that have converted to owner-occupied housing. Persistent price premiums within efficient neighbourhoods threaten to price renters out of the market. While premiums remain modest in magnitude, this presents a real risk given issues of affordability in big cities. Further risks exist from potential general equilibrium effects arising from gentrification and the access to amenities. Though beyond the scope of this paper, one could envisage a shock to the spatial equilibrium of local economies with wider potential knock-on effects. As the need for new technologies expands – such as air-conditioning and heat pumps – so too will the demand for homes that can support such features. This paper therefore provides early evidence as to the demand implications that threaten to become greater in the face of an ever-changing climate.

B Appendix

B Further Technical Information

B.1 Merging Energy Efficiency and Property Price Data

This chapter merges property-level data on energy-efficiency characteristics with temporal data on transaction prices (outlined in Section 3.3). The lack of a common unique property identifier makes merging the data difficult. Without this unique property reference number (UPRN) common in most geospatial data in the UK, I instead develop a multi-rule matching process similar to that of Chi et al., (2021) which merges data based on address characteristics. These address fields are first cleaned of punctuation or errand characters to eliminate discrepancies between the two datasets. They are then split into separate fields based on their unique components – such as property name/number, street name and postcode. Properties were first matched on their postcodes to ascertain the right locality. Under this approach, I avoid matching properties with others that just so happen to have the same street name. From these matches, records were then matched based on their street name, and then finally, on their property number to remove the risk of duplicated (or even misspelled) street names or property numbers. This procedure returned a match-rate of over 83% on properties within the London Travel to Work Area (TTWA). Remaining properties are likely those whose information contains spelling errors in either the address or house name fields, resulting in failure to match. Given the high match rate, the decision was taken not to further complicate the algorithm and collect missing properties.

B.2 Redefining Clusters Within Borough

In Section 3.3.3 I set out the approach to defining efficient clusters and the outputs of the method. Figure B.2 shows the resulting efficient clusters overlapped with borough boundaries across the London TTWA. While borders of these clusters do not overlap with the borough boundaries in most instances, there are the odd exceptions. I make the case that location fixed effects applied in Section 3.4 should deal with such instances by controlling for the impact of borough-specific characteristics on prices. To provide further robust-


ness to the estimates produced in this chapter, I outline here an alternative approach to defining clusters that deals with this feature, alongside the RD accompanying results.

I undertake the same approach outlined in the chapter. The only difference is the sample upon which the algorithm in Eq.(11) is re-estimated. To avoid any possibility that cluster borders overlap with administrative boundaries, I run the algorithm separately for properties i within each borough b for each year t using properties that list in the given year. Mathematically, one can rewrite the algorithm as

$$G_{ibt}^* = \frac{\sum_{j=1}^{n_t} w_{i_b, j_b} x_{ibt} - \bar{X} \sum_{j=1}^{n_t} w_{i_b, j_b}}{S \sqrt{\frac{n_t \sum_{j=1}^{n_t} w_{i_b, j_b}^2 - (\sum_{j=1}^{n_t} w_{i_b, j_b})^2}{n_t - 1}}} \quad (16)$$

where the weights are now calculated with respect to neighbours exclusively residing in the same neighbourhood. Households make location decisions based on a number of factors: by income Couture et al., (2023); amenities (Koster et al., (2014)); van Duijn and Rouwendal, (2021)); and by school quality Bayer et al., (2007). Households choose their location based on these factors and prioritise certain locations as a result. Property prices surge in areas with good schools, high income households, and areas with high amenities. This further re-enforces spatial sorting as each household – constrained by their budget – has a given number of viable locations for which they can purchase. Therefore clusters located in one borough, with proximate neighbours in the borough next door, may not actually reflect the true market search procedure for many households. The alternate approach undertaken here seeks to capture this by clustering within-borough in order to capture household preferences for certain boroughs. Following the same method applied earlier, I then group properties based on their similarity with respect to energy efficiency and their proximity into efficient cluster. Using a snapshot from one borough, Islington, Figure B.1 maps the resulting output.

Figure B.1: Efficient Clusters in Islington, London: All Years



NOTES: The figure illustrates the local authority district (borough) of Islington overlapped with efficient clusters as defined over the whole sample period 2009-2019. The figure replicates that of Figure 3.1, but restricts the computation of the Getis-Ord equation in Eq.(11) to operate within the borough – thereby not seeking spatial relationships with properties in the next borough.

Borders from the clusters now no longer overlap borough boundaries and instead end at this point. To ensure that boundary does not interfere with the estimation, distance from the cluster border is calculated for each property within the borough. For example, clusters whose western border intersects the borough boundary will be assigned properties residing to the north, east and south, for which distance to the border will be calculated. Borders now no longer overlap administrative boundaries. Results presented below show that conclusions regarding the price premium for these clusters remain unchanged.

Extensive margin estimates: I now present the results from the RD using the new cluster definitions in Table B.1. Estimates show that price premiums remain and remain significant upon redefining the clusters.

Table B.1: Estimates of the Price Premium for the Redefined Efficient Clusters

	(1)	(2)	(3)
	At Border	Border Swell	Border Shrink
Panel A: Baseline RDD			
1.Cluster	0.013*	0.000	-0.004
	(0.006)	(0.021)	(0.007)
Unrealised efficiency	0.005**	0.004**	0.004**
	(0.002)	(0.001)	(0.002)
Current efficiency	0.008***	0.008***	0.007***
	(0.001)	(0.001)	(0.001)
Property/Year FE	✓	✓	✓
Bandwidth	165	165	165
Adj. R^2	0.828	0.815	0.850
Observations	7,117	8,316	6,025
Panel B: RDD w/ Amenities			
1.Cluster	0.018**	-0.012	0.004
	(0.006)	(0.020)	(0.006)
Unrealised efficiency	0.004**	0.004**	0.003
	(0.001)	(0.001)	(0.001)
Current efficiency	0.008***	0.008***	0.006***
	(0.001)	(0.001)	(0.001)
Property/Year FE	✓	✓	✓
Bandwidth	165	165	165
Adj. R^2	0.103	0.104	0.029
Observations	3,256	3,418	3,007

NOTES: Standard errors clustered at the LSOA level in parentheses. Treated clusters are those defined as energy-efficient under the Getis-Ord calculation. The regression includes controls for the year of sale and follows a second-order polynomial. Flat's are excluded from the analysis. I swell and shrink the borders by 50m and re-estimate as a robustness checks in columns 2-3. Bandwidth measured in metres from the border of the treated clusters. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Buyers pay a premium of roughly 1.3% to reside within energy efficient clusters over and above those on the periphery. This estimate is robust to the inclusion of LocationxYear fixed effects in Panel B to control for changes in amenities over time in nearby locations. Panel B Column 1 forms the main specification; the estimate of the premium of 1.8% is very similar to that of the main result (1.6%) and highlights that the empirical strategy is robust to different methods of defining clusters. Robustness checks in columns 2-3 of both

panels show that the premium tends to zero or even becomes negative and insignificant when we alter the size of the clusters and hence the locations of the borders. The clusters estimated then are an accurate reflection of energy-efficient clusters and buyers are able to identify the size and location of them from property listings. Consistent with the main results in the chapter, we also observe in nearly all specifications the positive returns to energy efficiency improvements, as well as to unrealised efficiency. Highlighted in these results is that the main driving mechanism behind the premiums – the higher expected returns to investments in these neighbourhoods – remains regardless of the approach applied to defining clusters.

B.3 Intensive Margin Estimates

Intensive margin approach: Results in Section 3.4 highlight the premium residents pay to reside within efficient clusters – under the extensive margin. This approach represents a discrete choice of residing in either a efficient or non-efficient neighbourhood. This choice is motivated by the strength of spatial relationships locally with respect to energy efficiency scores, a proxy for peer effects. Strength in this case can be measured on a continuous scale. I use this proxy to estimate the intensive margin in order to motivate the discontinuity found in earlier estimates; specifically, how buyers price peer effects into their property bids. Estimates from the intensive bear light on the mechanism behind the price premiums observed in that it aims to prove that buyers draw inferences from property listings. It therefore serves as a way to elicit buyer’s perceptions of an area. The strength peer effects can be quantified using z-scores reported in the clustering algorithm that denote the strength of spatial correlation in energy efficiency scores. Higher correlation suggests similarity between scores, achieved only when households conduct similar home improvements. A continuous measure, these scores are applied in the creation of clusters through their statistical significance. Only significant scores were used to form clusters, thereby omitting a good deal of variation. I exploit this variation to tease out whether buyer’s choice of neighbourhood – and expected returns to investment – depends on the strength of these spatial relationships. I do this by estimating the spatial RD model separately for both the treated and control group at the border. The model differs only

in that I replace the treatment indicator with the z-score. Buyers possess the ability to identify clusters, thereby allowing me to adopt the approach of estimating separately for both groups. The spatial RD approach is formalised in Eq.(17).

$$\log(P_{it}) = \beta_0 + \beta_1 Zscore_{it} + \beta_2 \delta_{it} + \beta_3 X_{it} + \psi_i + \phi_t + \epsilon_{it} \quad \forall T_{it} \in \{T, C\} \quad (17)$$

The algorithm calculates a property's locally weighted energy efficiency score against the expected score and takes the difference. One would interpret the z-scores as the number of standard deviations this difference lies away from the mean. The coefficient β_1 therefore represents the change in price following a 1 s.d increase in the z-score – or as the energy efficiency of one's neighbours increases relative to that that is expected. The RD regression includes the same controls X as those in Eq.(12), only I now exclude current energy efficiency scores given their potential collinearity with the z-scores. Property fixed effects ψ are included as a granular location control to observe changes in valuation of local energy efficiency scores within property. Year of sale dummies are included to capture any variation in property markets between years.

Intensive margin estimates: Table B.2 presents the estimates of the spatial RD separately for properties in efficient (treatment) and inefficient (controls) clusters. Columns 1 and 4 start with naive specifications to identify the correlation between z-scores and property prices. Both estimates are negative and significant, suggesting that buyers bid lower for properties with higher energy efficiency scores locally. This is counter-intuitive as one would expect higher returns properties in areas similarly efficient properties. A low R^2 of 0.07 suggests that the Z-score alone provides weak explanatory power for prices and may instead be affected by other factors. Including property fixed effects in Columns 2 and 5 show this to be the case as the coefficients become positive and significant, meaning better match in areas with higher energy efficiency translates into higher prices in both specifications. Prices respond at a lower rate in the treated clusters, suggesting a lower price response to increased efficiency. Controlling for property-level characteristics has opposing effects on the magnitude of the z-score coefficients between neighbourhoods. Estimates in Column 2 were negatively bias while estimates in Column 4 were positively

bias. Much of the bias in the control group is explained by floor space, not the functional form. Functional form does explain a lot of the bias in the case of treatment clusters, however. The relationship appears to be heavily non-linear, with the coefficient increasing by around 0.04 s.d and with 0.037 of that explained by the squared z-score. A negative and significant coefficient on the squared z-score points to an effect of local efficiency scores on prices that increases at a decreasing rate. Buyers are willing to increase their bids for strong spatial relationships but only up to a certain point. The extensive margin shows that buyers identify a group of peers and pay premiums to locate alongside them. These neighbourhoods contain a non-negligible number of peers, while the z-score in comparison is confined to estimation based on a small number of neighbours. So while one's immediate peers are important, the wider neighbourhood is also important. This likely explains why unrealised efficiency becomes insignificant – given that buyers are unable to identify under this approach the full extent of the efficient cluster. The coefficient on the z-scores are relatively similar between groups owing to the fact that buyers respond to changes in the local efficiency scores. For treated clusters, a 1 s.d increase lead to a 5.4% increase in prices.


Table B.2: Estimates of the Strength of Spatial Relationships on Prices

	(1)	(2)	(3)	(4)	(5)	(6)
	Treatment			Control		
Z-score	-0.032*** (0.004)	0.013** (0.004)	0.053*** (0.011)	-0.024*** (0.006)	0.073*** (0.015)	0.051*** (0.014)
Z-score ²			-0.006*** (0.002)			-0.006 (0.007)
Unrealised			0.001 (0.001)			-0.000 (0.001)
Floor Area			0.000 (0.000)			0.003** (0.001)
Controls			✓			✓
Property FE		✓	✓		✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Bandwidth	170	170	170	170		
Adjusted R^2	0.077	0.817	0.819	0.057	0.728	0.736
Observations	6,298	6,298	6,298	7,027	7,027	7,027

NOTES: Standard errors clustered at the LSOA level in parentheses. Treated neighbourhoods are those defined as energy-efficient under the Getis-Ord calculation. Sale clusters are those defined using sellers in each year. No sale clusters use those properties that do not sell within a year of a new EPC in each inspection year. The regression includes property fixed effects alongside a series of dummy's controlling for year of sale. Though not reported here, the regression follows a second-order polynomial. Flat's are excluded from the analysis. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$


I show then that buyers do responding with positive bids for properties are surrounded by more energy-efficient peers and that this an important determinant of prices. Findings suggest that this is only important up to a point, and that other factors may well be explaining the premiums observed in estimates of the extensive margin – such as the return to investment. In what follows, I present results that speak to the impact of buyers selecting into these clusters to chase higher returns. Namely, the impact on incumbent households and their willingness to also make improvements.

Figure B.2: Efficient Clusters and Administrative Boundaries




NOTES The figure illustrates the extent to which local authority districts overlap with efficient clusters defined in Figure 3.1. The figure includes all efficient clusters defined over the sample period 2009-2019. We observe some overlaps at the border, where properties attached to the same cluster belong to different boroughs.

Figure B.3: The Distribution of Energy Ratings at the Border: No Sale Clusters



NOTES: This figure illustrates the proportion of properties in the sample that occupy each combination of current-potential energy rating, using properties that are not-yet listed on the market. The total in the right-hand column represents the proportion of properties falling in each current energy rating band.

Figure B.4: Rightmove Neighbourhood Definition



NOTES: This figure provides an illustration of a property search using the Rightmove portal and the associated radii with which buyers can use to search nearby properties.

© Crown copyright 2009

Energy Performance Certificate (EPC)

17 Any Street, District, Any Town, B5 5XX

Dwelling type: Detached house
Date of assessment: 15 August 2011
Date of certificate: 13 March 2012

Reference number: 0919-9628-8430-2785-5996
Type of assessment: RdSAP, existing dwelling
Total floor area: 165 m²

Use this document to:

- Compare current ratings of properties to see which properties are more energy efficient
- Find out how you can save energy and money by installing improvement measures

Estimated energy costs of dwelling for 3 years	£5,367
Over 3 years you could save	£2,865

Estimated energy costs of this home

	Current costs	Potential costs	Potential future savings
Lighting	£375 over 3 years	£207 over 3 years	 You could save £2,865 over 3 years
Heating	£4,443 over 3 years	£2,073 over 3 years	
Hot water	£549 over 3 years	£222 over 3 years	
Totals:	£5,367	£2,502	

These figures show how much the average household would spend in this property for heating, lighting and hot water. This excludes energy use for running appliances like TVs, computers and cookers, and any electricity generated by microgeneration.

Energy Efficiency Rating

Very energy efficient - lower running costs

(92 plus) **A**
(81-91) **B**
(69-80) **C**
(55-68) **D**
(39-54) **E**
(21-38) **F**
(1-20) **G**

Not energy efficient - higher running costs

Current	Potential
49	76

The graph shows the current energy efficiency of your home.

The higher the rating the lower your fuel bills are likely to be.

The potential rating shows the effect of undertaking the recommendations on page 3.

The average energy efficiency rating for a dwelling in England and Wales is band D (rating 60).

Top actions you can take to save money and make your home more efficient


Recommended measures	Indicative cost	Typical savings over 3 years	Available with Green Deal
1 Increase loft insulation to 270 mm	£100 - £350	£141	✓
2 Cavity wall insulation	£500 - £1,500	£537	✓
3 Draught proofing	£80 - £120	£78	✓

See page 3 for a full list of recommendations for this property.

To find out more about the recommended measures and other actions you could take today to save money, visit www.direct.gov.uk/savingenergy or call 0300 123 1234 (standard national rate). When the Green Deal launches, it may allow you to make your home warmer and cheaper to run at no up-front cost.

145

Figure B.6: Typical Energy Efficiency Rating Distribution



NOTES: The figure presents an example of an EPC rating distribution passed over to the buyer at the time of the property transaction.

Figure B.7: Example Property Listing Advert from Rightmove

PRICE

Monthly mortgage payments

PROPERTY TYPE

BEDROOMS

BATHROOMS

TENURE ⓘ



1

PROPERTY IMAGES

Key features

Property description

Read more

Tenure: Leasehold ⓘ

GROUND RENT ⓘ

ANNUAL SERVICE CHARGE ⓘ

LENGTH OF LEASE ⓘ

Energy Performance Certificates ⓘ


Council Tax ⓘ

epc

Ask agent

NOTES: This figure provides an illustration of a property listing presented on the Rightmove portal and the associated information contained within – namely, the property’s EPC certificate.

Figure B.8: Average Property Prices: London vs UK




NOTES: This figure illustrates the ratio of London-to-UK house prices. Data are sourced from HM Land Registry House Price Index and are generated from a hedonic price model accounting for differences in characteristics. The figure presents the order by which prices in London exceed the UK average.

Figure B.9: Defining The Spatial Extent of Efficient Clusters




NOTES: This figure illustrates the extent of the neighbourhoods defined as energy efficient. This map is the precursor to Figure 3.1, which uses the extents defined here to bound the efficient clusters. Here, properties are aggregated in each sale year based on their treatment status (whether they are hotspot or coldspot) which comes directly from the z-scores obtained from the G_i^* algorithm. The output here represents these points connected in a dot-to-dot fashion.

Figure B.10: Time Periods Defined in the Event Study




NOTES: This figure illustrates how the treatment is defined within the event study. Treatment properties are those that sell within a cluster at the time it is defined as being efficient by the market (at time $\tau = 0$).

Figure B.11: Construction of the Event Study Control Group




NOTES: This figure illustrates how the control group is formed in the event study. The control group comprises properties on the periphery (yellow) of the treatment groups (orange). These properties do not receive the treatment, and therefore should not experience premiums, as they are outside of the efficient cluster.

Figure B.12: Property Prices in Efficient Clusters Relative to the Treatment Period




NOTES: This figure illustrates how property prices for treated properties change over time relative to the baseline year $\tau = 0$ (when properties sell within an active efficient cluster). Controls include the property's current efficiency score and total floor area, as well as controls for property age, type, and build. Location controls are also included for the MSOA in which the property sits.

Figure B.13: Event Study Robustness Checks




NOTES: This figure presents the results from the event study for all years in the sample (left), as well as after removing properties selling prior to 2012. Property portals began formally including EPCs within property listings in 2012, prior to this, it was a choice. Evidence shows that they began being listed as a choice in 2010, hence these figures provide a robustness check of the main result in the paper. Properties selling within an efficient cluster at any time represent the treated group, while the control group consists properties selling on the edge of these clusters. A bandwidth of 50m has been chosen, capturing properties either side of the efficient cluster (treatment) border. Property controls include the current efficiency score, total floor area, property age, property type, and build type.

Figure B.14: Properties by Built-Form by Treatment Status




NOTES: This figure illustrates the distribution of properties by built form within both the treatment and control group. The bars represent the proportion of the sample corresponding to each property type.

Figure B.15: Properties by Built-Form by Treatment Status – No Sale




NOTES: This figure illustrates the distribution of properties by built form within both the treatment and control group for properties that have not-yet listed. The bars represent the proportion of the sample corresponding to each property type.

Figure B.16: Properties by Age of Construction by Treatment Status




NOTES: This figure illustrates the distribution of properties by year of construction within both the treatment and control group. The bars represent the proportion of the sample corresponding to each year of construction.

Figure B.17: Properties by Age of Construction by Treatment Status – No Sale



NOTES: This figure illustrates the distribution of properties by year of construction within both the treatment and control group for properties that have not-yet listed. The bars represent the proportion of the sample corresponding to each year of construction.

Figure B.18: Balancing of Neighbourhood Covariates at Cluster Border



NOTES: This figure presents results from a series of RDDs that explore whether there exists discontinuities in neighbourhood characteristics at the border of clusters. Data are collected at the output area (OA) level, with distance from the border measured in miles. The plots are unconditional, representing the pure correlations between the intended outcome and the efficient cluster location.

Table B.3: Balancing Test at the Border of Efficient Clusters: No Sale

	Treatment		Control		Balancing	
	Mean	Sd	Mean	Sd	t	p> t
Price	378,853	223,300	424,393	299,327	-5.7	0.000
Adj. price	478,071	97,261	422,449	105,636	17.8	0.000
Dist. to border	41.6	30	45.4	37.3	-3.6	0.000
Current efficiency	77.4	12.8	67.2	16.3	22.6	0.000
Potential efficiency	82.4	10.9	76	14	16.8	0.000
Unrealised Efficiency	5	6.8	8.7	9.6	-14.6	0.000
Energy cost (.sqm)	6.1	3.2	8.4	4.2	-20.2	0.000
Energy consumption	142	109.3	226.1	139.3	-22	0.000
Window glazing %	81.2	34.5	80.9	33.8	0.2	0.816
Floor Area	99.8	36.1	101.2	40.5	-1.1	0.262
No. rooms	4.6	1.4	4.7	1.4	-0.7	0.498
Observations	2,211		2,013			

NOTES: This table presents the unconditional means for a set of property characteristics for properties not-yet listed. Data presented are for properties within 112m of the border, the bandwidth selected for the baseline RD later. Treated neighbourhoods are the energy-efficient neighbourhoods denoted by the Getis-Ord clusters. Prices here are stripped of their time and location drivers by including property and time fixed effects. Energy cost measures the estimated spend on utilities (heat, lighting and hot water) in GBP for a house of given structural characteristics and size.

Table B.4: Estimates of the Strength of Spatial Relationships on Prices

	(1)	(1)	(1)
Zscore	-0.035*** (0.001)	0.078*** (0.001)	-0.002 (0.003)
Zscore ²			0.005*** (0.001)
Current			0.007*** (0.000)
Floor Area			0.001** (0.000)
Controls			✓
Property FE		✓	✓
Year FE	✓	✓	✓
Adjusted R^2	0.122	0.813	0.821
Observations	141,583	141,583	141,583

NOTES: Standard errors clustered at the LSOA level in parentheses. The regression explores the intensive margin of efficient clusters to estimate whether the strength of spatial relationships matter. The regression includes property and year of sale fixed effects alongside a series of dummy's controlling for year of sale. I further control for the property's total floor area. Flat's are excluded from the analysis, like in all other estimates. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.5: Estimates of the Price Premium for Efficient Clusters with Unrealised Efficiency

	(2)	(2)	(3)	(4)	(5)	(6)
	Baseline	RDD	No Sale	Clusters	Sales Not in	Clusters
1.Cluster	0.0160*	0.0157	0.000320	0.00504	0.0106	0.0126
	(0.00760)	(0.0101)	(0.0189)	(0.0196)	(0.0195)	(0.0195)
1.ClusterxUnrealised		-0.000191		-0.00130		-0.000240
		(0.00124)		(0.00226)		(0.000449)
Unrealised		0.00239**		0.0000829		0.000810*
		(0.000866)		(0.00244)		(0.000373)
Distance	-0.000114	-0.000108	-0.0000460	-0.0000417	0.000235	0.000223
	(0.00006)	(0.00006)	(0.00014)	(0.00014)	(0.00016)	(0.00016)
Current	0.00668***	0.00735***	0.00797***	0.00791***	0.00683***	0.00712***
	(0.0008)	(0.0009)	(0.0012)	(0.0013)	(0.0002)	(0.0003)
Floor Area	0.000877	0.000888	-0.000145	-0.000120	0.000283	0.000280
	(0.0007)	(0.0007)	(0.0014)	(0.0014)	(0.0003)	(0.0003)
Property FE	✓	✓	✓	✓	✓	
Year FE	✓	✓	✓	✓	✓	✓
Bandwidth	150	150	150	150	150	150
Adjusted R^2	0.782	0.783	0.644	0.644	0.831	0.831
Observations	12,100	12,100	5,059	5,059	19,115	19,115

NOTES: Standard errors in parentheses. Unrealised efficiency is defined as the difference between a property's current and potential energy efficiency. Treated neighbourhoods are those defined as energy-efficient under the Getis-Ord calculation. The regression includes property fixed effects alongside a series of dummy's controlling for year of sale. Though not reported here, the regression follows a second-order polynomial. Flat's are excluded from the analysis. Column 3-4 presents estimates when clusters defined using never-sellers. Column 5-6 instead uses the original clusters, but sales which do not fall into a live cluster. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.6: Changes in Unrealised Efficiency on Prices – Whole Sample

	(1) Baseline RDD	(2) No Sale Clusters	(3) Sales Not in Clusters
1.ClusterxUnrealised	-0.000962* (0.000375)	-0.00431***	-0.000734* (0.000342)
Unrealised	-0.000328** (0.000113)	-0.000394*** (0.000112)	0.000335 (0.000329)
Property FE	✓	✓	✓
Year FE	✓	✓	✓
Bandwidth	None	None	None
Adjusted R^2	0.820	0.821	0.820
Observations	141,753	142,178	117,436

NOTES: Standard errors clustered at the LSOA level in parentheses. Treated neighbourhoods are those defined as energy-efficient under the Getis-Ord calculation. The regression includes property fixed effects alongside a series of dummy's controlling for year of sale. Though not reported here, the regression follows a second-order polynomial. Flat's are excluded from the analysis. Column 3-4 presents estimates when clusters defined using never-sellers. Column 5-6 instead uses the original clusters, but sales which do not fall into a live cluster.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Chapter 4: Household Location and Investment Decisions Under Neighbourhood Effects: A Theoretical Approach

4.1 Introduction

In chapter two I show that energy efficiency investments are capitalised into property prices, in line with findings from recent studies (Fuerst et al., 2015; Fowlie et al., 2015; Sejas-Portillo et al., 2020; Aydin et al., 2020). Though returns to these investments are evidenced in the literature, their uptake remains low, pointing to a wider problem in the incentives facing households. This represents a shortcoming of this literature that likely finds its roots in the lack of spatial evidence on both the magnitude and drivers of such returns. I address this shortcoming in chapter three by showing that returns vary across space owing to neighbourhood effects that determine household location decisions and their uptake of home investments. Buyers learn from the characteristics of their peers when moving property and choose the most energy efficient neighbourhoods that deliver the greatest returns. These returns, coupled with the global push to achieve ambitious Net Zero targets, makes investments into energy efficiency improvements a particularly viable option for households looking to build wealth. Currently, this is an under-researched topic. Studies in the urban economics literature fail to rationalise how household location decisions depend on such factors, instead focusing on other explanations such as access to good schools (Bayer et al., 2007), access to cultural amenities (Koster et al., 2014; van Duijn and Rouwendal, 2021; Autor et al., 2014), and the general affluence of a locality (Couture et al., 2023; Guerrieri et al., 2013). While these papers provide the mechanisms underpinning the sorting of households across space, they are not directly applicable to the literature on residential energy efficiency. With that in mind, I develop in this chapter a model that rationalise the findings from the previous two chapters. The model explores how household's demand for investment affects their location decisions and how this subsequently affects property prices in the same location. I pin down the mechanisms to a neighbourhood effect; when strong, the neighbourhood effect delivers higher levels of

investment and higher property price, versus the opposite effect when the neighbourhood effect is weak.

I develop the model by drawing on the urban economics literature to explain how investment demand and location decisions are influenced by one's peers and how this translates into property bids. Two mutually exclusive groups of households are considered – those that move house and those that do not move house (referred to hereafter as movers and non-movers). Movers take decisions on how much to consume of a numeraire good, investment services (for housing), and housing – obtaining utility which I assume to follow Stone-Geary preferences. The Stone-Geary functional form imposes non-homothetic preferences and requires that households obtain a minimum subsistence level of a good, and only after this do they receive utility from consuming it. Papers in the spatial sorting literature require non-homothetic such that households with differing incomes consume differently – making the Stone-Geary functional form an attractive option (Schmidheiny, 2006; Eeckhout et al., 2014; Gaigné et al., 2022; Blouri et al., 2023; Grossmann et al., 2021). Movers also draw utility from the investment potential provided by their chosen neighbourhood. Neighbourhood effects, which following on from the previous chapter takes the form of local neighbour's energy efficiency scores, enter the model through this incentive with a parameter defining the buyer's degree of attentiveness to scores across the neighbourhood. I model this in such a way that choosing to reside within areas of higher energy efficiency crowds out one's own investment through the decision to free-ride on the back of their neighbour's investments. This is consistent with the effect of an increase in spending on regeneration projects modelled in Rossi-Hansberg et al., (2010). A benefit of developing the model in this way is that it returns a set of demand functions for each good as a function of the energy efficiency levels of the surrounding neighbourhood, meaning I am able to capture how choosing a neighbourhood with higher levels of energy efficiency impacts all forms of consumption that ultimately changes the budget constraint for housing. This is particularly important when deriving the demand for housing and investment services. A second component of the model explores the subsequent effects this has on investment spending within the wider neighbourhood among non-movers. Neigh-

bourhood effects for non-movers come in the form of exposure to the investment services consumed by movers and increase in intensity with the volume of services observed nearby. I show in the empirical analysis in chapter three that movers select into the most efficient neighbourhoods on the back of preferences for energy efficiency, and that exposure to this group increases the likelihood that non-movers make home improvements of their own. The model requires that each household faces a degree of exposure to local investment – measured as the spatially-weighted sum of local investment expenditure undertaken by movers – and that buyers learn about this through the neighbourhood effect which induces them to increase their own investment.

Findings from the model align with those of the third chapter, shedding further light on the mechanisms at play. Firstly, both movers and non-movers increase their demand for investment in neighbourhoods with higher energy efficiency. Each household demands a baseline level of investment in line with their preferences and are subject to spillovers in the form of neighbourhood effects that increase their demand; investment among non-movers is increasing in the investments made by movers – the latter of which is increasing in the energy efficiency of the chosen neighbourhood. This occurs through exposure, as movers have preferences for energy efficiency that translate into investments observable by non-movers. The degree to which neighbourhood effects increase investments among movers is determined by the degree of attentiveness they have for energy efficiency scores across the neighbourhood at the time of purchase. Greater attentiveness means the buyer places more importance on higher levels of efficiency – translating into greater consumption of investment services and greater spillovers to non-movers. Intuitively, this makes sense given that households would only increase investment beyond the baseline if they perceived a benefit to doing so. This aligns with Baldauf et al., (2020) who find that the subject (in their case, house prices) only responds to a phenomenon (flood risk) if the agent believes in its existence. By incorporating the degree of attentiveness parameter is to show that households would in fact undertake different levels of investment across different neighbourhoods, that vary in their efficiency. Secondly, because property prices in equilibrium are a function of the investment potential, the price equation reflects the

resulting changes in investment demand. Doing so matches the empirical estimates and confirms that highly-efficient neighbourhoods attract a price premium over-and-above those that are less efficient. I also show that higher levels of investment demand leads buyers to bid higher for properties. By characterising the equilibrium in the economy, I pin this down to the fact that buyers expect higher returns on their investments into these neighbourhoods at the time of sale. This return is expected on the investments made beyond those required to match the investment levels of the neighbourhood; which I show to be undertaken in the most efficient neighbourhoods by those with preferences for energy efficiency. Finally, higher levels of investment undertaken by movers induces higher levels of investments by non-movers through greater exposure, but this is only true in neighbourhoods that are among the most energy efficient. Though exposure drives higher investments, this exposure is greatest in the most efficient neighbourhoods where greater investments take place among movers. This acts in a similar way whereby neighbourhood effects are present when non-movers observe home improvements undertaken locally. These findings match those of the third chapter and shed light on the mechanisms driving the increased investment – namely the attentiveness to energy efficiency scores and the increased incidence of home improvements across the neighbourhood.

I contribute to the literature in three ways. Firstly, there is a dearth of theoretical evidence exploring the drivers of residential energy efficiency investments. Most models from the literature focus on the issue of underinvestment. Allcott and Greenstone, (2012) provide a conceptual framework that prescribes the low uptake of residential energy efficiency investments as resulting from poor returns on investment overshadowed by high capital cost. Others attribute this to high discount rates and produce models that optimise human behaviour (Hausman, 1979; Dubin and D. L. McFadden, 1984; Ansar and Sparks, 2009). While a separate strand of literature explores policy-based remedies to underinvestment and the resulting welfare effects after correcting for behavioural bias (Allcott et al., 2012; Allcott and Taubinsky, 2013). These papers favour explanations for the poor uptake of improvements at the household-level. While this is important, it is a narrow view of the overall picture that neglects the important influence one’s peers have on be-

haviour. Findings from this chapter actually contradict the findings from these papers, providing evidence of increasing returns to investment under market forces as opposed to a policy response.

Secondly, I provide a unique, spatial approach to exploring the drivers of investment by modelling uptake as a function of neighbourhood effects. Common approaches in the literature adopt a more restrictive definition of neighbourhood effects as deriving from specific individuals with specific attributes – as in the case of models that exploit networks. These applications find their roots in education and public economics. I deviate from this and provide a more simple approach by modelling neighbourhood effects as arising from a neighbourhood’s energy efficiency. Such an approach aligns with chapter three, where movers identify peers through property listing adverts based on their energy efficiency characteristics and form expectations around the returns on investment. Neighbourhood effects in this way enter the utility function and affect the household’s investment potential in different neighbourhoods – thereby changing the demand for other consumption bundles as a result. Using this approach is more preferable for two reasons, as it: i) avoids the need to model explicitly the nature of household interactions across the neighbourhood, and ii) allows me to capture the shifts in investment demand external to the household. Studies from the housing and energy literature provide evidence that peer effects change household behaviour. Patacchini and Venanzoni, (2014) use network analysis to show that peer effects are present in home maintenance decisions as households conform to social norms; Bollinger and Gillingham, (2012) demonstrate similar conformism in the adoption of solar PPV; while Allcott, (2011b) and Bollinger et al., (2020) highlight that conformism is also present in reducing the demand and consumption of utilities and water resources, respectively. Evidence that peer effects modify behaviour make this model an important one to understand the mechanisms at play in the market for household investment and location decisions under neighbourhood effects.

By focusing solely on the gains realised through energy savings or the wider uptake of improvements, previous papers neglect a key feature: the effect that exposure to peer ef-

fects has on property prices. My third and final contribution to the literature is to model the effect of peer-induced investment demand on property prices. Conflicting evidence of the effects on prices from the previous two chapters fails to establish a clear theoretical grounding. Chapter two would suggest that buyers discount prices in order to fund investment, while chapter three suggests that buyers pay a premium in expectation of receiving it back at the time of sale. Findings from the latter suggest that neighbourhood effects drive this result. I fill this gap in the literature by producing a model that explores the effect on property prices of greater demand for investment when neighbourhood effects are present. Any estimate of the property price equation is by design a function of a neighbourhood's energy efficiency and allows me to represent prices as a function of the peer effect. I can therefore explore the price implications of varying degrees of the neighbourhood effect. Overlooking the impact on prices is to miss an important piece of the puzzle; how the resulting premium in these neighbourhoods affects the location decisions of buyers. Both Guerrieri et al., (2013) and Couture et al., 2023 show that gentrification causes richer households to price out poorer households in Downtown locations. While the chapter in questions is not concerned with this, it lays the foundations for thinking about such an issue and the effects on the spatial equilibrium of a city as buyers become priced out of markets. I therefore contribute to the debate about residential sorting and household location decisions by viewing this through the lens of wealth accumulation arising from residential investments.

I present some stylised facts about investment demand in Section 4.2 and the model for both movers and non-movers in Section 4.3. Conclusions are drawn in Section 4.4, and derivation for the model can be found in Section 4.5.

4.2 Stylised Facts

The previous chapters show that investment potential plays a key role in determining household location decisions and property prices in a neighbourhood. Chapter two finds that buyers value the potential energy efficiency of properties only if undertaking improvements makes this is achievable. Chapter three on the other hand shows that the same

potential is valued at the neighbourhood-level in locations that means it delivers returns on investment. The model presented in this chapter looks to reconcile these two findings and show that neighbourhood effects generate a price premium in neighbourhoods with investment potential. Before presenting the model, I present some stylised facts that motivate the relationships between investment, energy efficiency and property prices. I align this with chapter three and present descriptive statistics using data gathered for the London Travel To Work Area (TTWA).


4.2.1 #1: Moderately Efficient Neighbourhoods Have Investment Potential

A stereotypically rational households wanting to undertake home improvements would only do so if they saw a return on their investments, either financially or through the added comfort of their living environment. One can capture both returns through measures of energy efficiency. Current energy efficiency scores proxy for the level of thermal comfort provided by a property, while unrealised efficiency – the difference between a property’s current and potential score – captures its investment potential and hence the prospect of future returns. The data for both measures come from the universe of EPCs recorded in the London TTWA – for which I take the scores from each property’s latest certificate. I take the median value for both scores at the output area level and plot them in Figure 4.1. Output areas are a Census-define classification which splits local authorities into smaller geographical areas with a usual resident population ranging between 100-625 people. I present the median to avoid capturing neighbourhoods with polarised efficiency scores and to avoid the over-representation of more efficient properties.

From the plot, one can see that the density of neighbourhoods is highest around the middle of the distribution. Darker areas represent a larger cluster, with a large proportion of output areas having a median current efficiency score between 50-60. This is only a low D-rating which represents poor efficiency. Included in this cluster are areas with higher efficiency scores around 70. Properties nearer the centre of the distribution exhibit a high degree of investment potential – with the prospect of increasing their efficiency by around one-third. Such neighbourhoods should be attractive to buyers with preferences for energy

efficiency and those looking to invest in their property. The same is not true for areas with greater efficiency. A negative correlation between current and unrealised efficiency suggests that highly-rated neighbourhoods have lower potential. One should expect this given that there are limits to the improvements that can be made. The gradient is steep, however, illustrating that investment potential is highest towards the centre of the distribution and lowest even for properties just beyond the centre. A tight and relatively linear gradient like this should make it easier for buyers to identify areas with potential, thereby ensuring that buyers can accurately respond to information surrounding investment potential.


Figure 4.1: Investment Potential by Neighbourhood



NOTES: This figure takes the median current and unrealised efficiency for each neighbourhood and plots the result. I take the median to avoid over-representing neighbourhoods with polarised levels of energy efficiency. Analysis is undertaken at the level of output area.

Greater potential at the centre of the distribution tells an interesting story, but fails to provide insight into the spatial patterns behind this. In Figure 4.2 I take both axes from the scatter plot and impose them on separate maps to illustrate this relationship spatially. The left-hand map presents current energy efficiency scores and the right-hand map unrealised efficiency (both still representing the median property in each output area).

Figure 4.2: Energy Efficiency by Neighbourhood: London TTWA



NOTES: This figure maps median current and unrealised efficiency scores by output area for each property's latest EPC. Scores are broken down into bands.


Taken individually, the maps illustrate the expected pattern that Central London has higher levels of energy efficiency and lower investment potential owing to the modern housing stock. This is also evident in the gentrified west of the city and regenerated regions along the river Thames in the east. The effects of gentrification can be seen by poorer areas in the east that have more modest energy efficiency but more potential. Energy efficiency dissipates as one leaves the centre and more so as we reach the outskirts, illustrated by neighbourhoods with more modest scores. The previous pattern of modest efficiency and high potential seems more apparent in these suburban neighbourhoods. While neighbourhoods closer to the city centre offer better access to amenities, more distant neighbourhoods offer better investment potential – highlighting a trade-off in residential sorting decisions and motivating the need for a model to understand the impacts of investment demand.

4.2.2 #2: Investment is Driven by Preferences and Potential

Investments into improving residential energy efficiency are often expensive and the returns uncertain, making for variation spatially. In chapter three of this thesis I show this to be the case. Contained within EPCs are a set of recommendations put forward by the assessing engineer that enable the household to obtain their potential efficiency. By accessing the costs associated with these recommendations, I am able to identify the investments undertaken by households between certificates and aggregate them by output area to derive a measure of total investment expenditure. I map the spatial distribution

of this expenditure in Figure 4.3.

Figure 4.3: Investment Expenditure by Neighbourhood: London TTWA




NOTES: Investment expenditure is created as a proxy for the amount households spend on energy efficiency improvements. I calculate this using a repeat sales approach, whereby an improvement is considered undertaken if it was listed as a recommendation in the property's first EPC but not in subsequent EPCs. I observe the cost of each recommendation and sum these across the neighbourhood, weighting by the number of properties in that location.

Bearing in mind that population is pretty much fixed across all output area, a lot of the investment occurs in areas outside of central London. While this is likely driven by the age of the housing stock, as noted before, it is also driven by build type – as a higher share of central properties are apartments that are harder to improve. Pockets of East London have high levels of investment; so do too more peripheral regions in the north and south where investment potential is highest. Gentrified neighbourhoods in the more affluent west of the city demonstrate high levels of investment, but less than the aforementioned locations. One would also imagine that this is a product of demographics, owing to a larger share of family dwellings. That being said, we do observe a high degree of spatial variation which seems to follow a pattern of investment in suburban locations. I break this down further by extracting the neighbourhoods with the highest levels of investment and

superimposing them on the map of the median investment potential (unrealised efficiency). The resulting map is presented in Figure 4.4.

Figure 4.4: Investment Potential and Expenditure by Neighbourhood: London TTWA



NOTES: The left-hand figure maps investment potential – defined as the median unrealised efficiency within each output area – and superimposes the output areas with the highest levels of investment expenditure (those areas that fall into the top decile in Figure 4.3). This is taken for each property’s latest EPC. The right-hand figure plots the median current and unrealised expenditure (like in Figure 4.1), where the size of the bubbles corresponds to the size of the investment spend within the output area.


One can see from the left-hand map that investment is highest in neighbourhoods with the highest potential – evident by the smattering of green. High levels of unrealised efficiency means that higher levels of investment can be undertaken, so one would expect this results. Investment is heavily concentrated in suburban areas, but pockets are evident in more central locations as well with lower levels of unrealised efficiency – those most constrained in terms of the possible investment expenditure. One can draw from this that preferences also play a role, as investment is high in neighbourhoods that have little to gain. I explore this in the right-hand visual of Figure 4.4 by plotting current against unrealised efficiency. The size of the bubbles vary by the investment spend in the neighbourhood from which two clusters are identified. The first shows that investment spend is high in areas with modest efficiency and high unrealised efficiency; supporting earlier conclusions that investment occurs in areas with most to gain. The second cluster shows that high levels of investment also occur in the most efficient neighbourhoods with less to gain and points the idea that this is in fact preference-driven. The descriptive statistics presented here align with the findings from chapter three, which demonstrate that i) households with preferences for residential energy efficiency purchase in the most efficient neighbourhoods, and that ii) investment in these neighbourhoods occurs as households aim to gain from improving

poorer properties. I rationalise these findings with the model developed in this chapter.

4.2.3 #3: Energy Efficiency and Investment Potential Draw Higher Prices

While traditional explanations of property price variation point to access to employment, amenities and public services, the previous descriptive analyses would suggest that property quality is also an important explanatory factor. I delve into this as an explanation by exploring the correlation between property prices and energy efficiency. Using property-level data on transaction prices supplied by HM Land Registry, I calculate the change in property prices between 2009-2019 by Lower-Layer Super Output Area (LSOA). These are similar to output areas, only they are slightly larger and contain between 1,000-3,000 individuals. Figure 4.5 takes the LSOAs with the highest growth over the period and superimposes them on the previous maps depicting energy efficiency.

Figure 4.5: Energy Efficiency Among Priciest Neighbourhoods: London TTWA



NOTES: In this figure, I replicate the maps presenting the median current and unrealised efficiency scores by output area in Figure 4.2 and superimpose the LSOAs with the highest levels of property price growth over the period 2009-2019 – calculated using data from HM Land Registry. Property price figures are produced by the Office for National Statistics (ONS) using hedonic regressions to perform mix-adjustment in property types.


Property price growth is highest in neighbourhoods with more energy-efficient properties – reflected in the large pockets of yellow and green on the left-hand map. Growth clustered around central London is no surprise, but this also branches out across the whole TTWA in neighbourhoods in both the North- and South-West. These areas still have energy efficiency levels greater than the average which provides descriptive evidence of a relationship between property prices and energy efficiency. At a glance, one would not see the same relationships for unrealised efficiency. Generally speaking, neighbourhoods with

the largest growth in property prices are those with lower potential (those shaded in red). Central locations – which previously show to have lower potential – are the main drivers of this trend. Looking further out towards suburban locations, areas with potential do exhibit price growth; northern and eastern locations in-particular. This provides suggestive evidence that price growth in this case is more likely when investments deliver more tangible benefits; i.e. when properties can be taken from energy-inefficient to efficient.

4.2.4 #4: Property Prices Reflect Higher Levels of Investment

While the previous subsection provides descriptive evidence that property prices respond positively to higher levels of energy efficiency (which is well-noted in the literature review in chapter two) – a lesser understood question is whether households recoup their investments at the time of sale. The purpose of this chapter is to model how buyers undertake investment and the impact that this demand has on the prices buyers offer for property. As buyer's perceptions of investment returns are important for their ultimate offers, the model requires that there be evidence of returns. I show this to be the case in Figure [4.6](#).

Figure 4.6: Investment Expenditure and Property Prices (2009-2019): London TTWA



NOTES: In this figure, I plot the LSOAs with the highest property price growth over the period 2009-2019. I then take the plot of investment expenditure by output area from Figure 4.3 and superimpose those locations with the highest level of expenditure (locations in the highest decile).

Mapped here is property-price growth by neighbourhood between 2009-2019, upon which I superimpose the neighbourhoods with the highest levels of investment expenditure. Yellow and green neighbourhoods are those with positive price growth while red and orange are those with price contraction. The majority of neighbourhoods with the highest levels of investment experience positive price growth, with most growing at rates upwards of 50%. Neighbourhoods are evenly distributed across the TTWA showing that there are no location biases in the relationship between property prices and investment. Central locations are predictably of higher growth, but suburban neighbourhoods with high potential also experience relatively high levels.

The previous figures highlight a significant degree of correlation between investment, energy efficiency, and property prices. Given that these are explored in isolation, one may think that these correlations are driven by unobserved variables. I therefore run two re-

gressions gathering data at the output-area level to show that this is not the case.

Column 1 presents the results of a regression of energy efficiency characteristics on the level of investment expenditure in a log-log specification. Neighbourhoods with higher levels of energy efficiency and investment potential experience higher levels of investment. A unit-percent increase in both leads to a 10.6% and 0.2% increase in investment expenditure, respectively. One can see the differing investment incentives when observing the coefficients on the squared terms, however. The negative coefficient for current efficiency suggests that investment expenditure increases at a decreasing rate beyond a certain point. We can see this in the coefficients on the dummy variables, that expenditure is lower and statistically insignificant in areas with higher levels of current efficiency (1.High). Conversely, the positive coefficient on the squared unrealised efficiency term suggests that investment expenditure may be increasing with investment potential across the distribution – although this is statistically insignificant. One can see this is the case by looking at the coefficients on the dummy variables for the unrealised distribution; the coefficients grow in magnitude with the level of unrealised efficiency, suggesting higher investment expenditure in neighbourhoods with greater potential. Findings from Column 1 align with Fact #2 – that investment is driven by preferences as households invest in neighbourhoods with moderately-high efficiency, and by investment potential as households invest in neighbourhoods with moderate-efficiency and high unrealised efficiency.

Table 4.1: Drivers of Investment Expenditure and Property Prices

	(1)	(2)
Inv. Expenditure		0.840** (0.259)
Current	10.58*** (1.26)	312.3*** (41.70)
Current ²	-0.0014*** (0.00017)	-0.03*** (0.0053)
Unrealised	0.208*** (0.0509)	0.857 (1.646)
Unrealised ²	0.0002 (0.00017)	0.0327*** (0.006)
Current Distribution		
1.Mid	0.143*** (0.0294)	1.982 (1.125)
1.High	0.122 (0.0555)	3.662 (1.931)
Unrealised Distribution		
1.Mid	0.236*** (0.0334)	4.049** (1.268)
1.High	0.270*** (0.0508)	5.104** (1.836)
Adj. R ²	0.108	0.107
Observations	22370	20890

NOTES: Standard errors clustered at the output area in parentheses. All continuous variables are measured in logs. The dependent variable in Column 1 is the log of investment expenditure; for Column 2 it is the log of property prices. Data are collected at the output area level, using the median of current and unrealised efficiency. I control for Gross Value Added (GVA), distance to CBD, population, and density to capture effects arising from the location, size, and income of an area. 1.Low represents the hold-out category for both the Current and Unrealised Distribution. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Columns 2 on the other hand presents the results of a regression of the same characteristics on the change in property prices over the sample period using a linear-log model.³⁶ One can see that the energy efficiency of the neighbourhood generates an increase in property prices. A 1% increase in the median current energy efficiency score of a neighbourhood

³⁶The dependent variable is a percentage change and so contains positive and negative values, precluding the use of a log-log specification.

generated a 3 percentage-point increases in property price growth over the period. This represents significant growth – especially when one compares this with a smaller and statistically insignificant coefficient for unrealised efficiency. However, turning again to the squared energy efficiency variables tells a better story. A negative and statistically significant coefficient on the current efficiency term suggests that prices only increase up to a certain level of efficiency, whereas a positive and statistically significant coefficient on the unrealised efficiency terms suggests that prices are growing faster in neighbourhoods with greater investment potential. The latter is very-much driven by areas with higher levels of investment potential, as can be seen by magnitude of the coefficients across the unrealised distribution. Findings here align with Fact #3 that higher prices are found in neighbourhoods with greater energy efficiency and investment potential. Column 2 also presents estimates of the relationship between investment expenditure and property prices. The coefficient shows that property prices grew on average by 0.8 percentage points in response to a £1,000 increase in investment expenditure – highlighting the fact that buyers capitalise such expenditure. This corroborates the descriptive analysis presented in support of Fact #4 that neighbourhoods experiencing price growth are the ones previously having received the highest levels of investment.

Using both descriptive statistics and empirical analysis, I provide evidence to support four key facts: that investment potential is highest in moderately-efficient neighbourhoods; investment expenditure is highest in energy-efficient neighbourhoods with investment potential (driven by preferences); both are reflected in higher property prices in these neighbourhoods; and that property prices further reflect the investment expenditure undertaken. Previous research within the empirical literature has failed to recognise the effects of investment potential and expenditure on property prices. Both impact the location decisions taken by movers, adding further complication. The stylised facts presented here show that the effects are non-negligible and therefore motivate the need for better understanding. The model presented in this chapter seeks to shed light on the mechanisms driving these features and their effects on property prices. In the next section I introduce and solve the model.

4.3 Model

Consider a closed city composed of $n = \{1, \dots, N\}$ neighbourhoods with a fixed number of households i that belong to one of two mutually exclusive groups – those choosing to move property (movers) and those that do not move property (non-movers).³⁷ Both make decisions in the same period: movers make location decisions based on the prevailing price of property and their own demand for housing investment within a neighbourhood, while non-movers make investment decisions based on exposure to those taken by movers. The latter occurs as non-movers observe the investments undertaken in their neighbourhood. The purpose of this model is to capture the drivers of demand for investment and to explore the implications this has for property prices and household location decisions: specifically, the model pins down the impact in both the presence and absence of neighbourhood effects. Both movers and non-movers are exposed to a neighbourhood effect originating from different sources. I develop a model of an economy composed of two sectors, one producing final goods and the other providing investment services. In solving the model, I show that property prices increase in response to household investment and the neighbourhood's energy efficiency. I then characterise the equilibrium of the economy before studying the decisions of non-movers. I start by building the structure of the economy and introducing the utility framework for movers.

4.3.1 Movers

Moving households obtain utility from the consumption of a freely-traded final (numeraire) good C_i , housing H_i and investment services I_i according to the utility specification in Eq.(18). Preferences for each are captured by the parameters α_i , β_i and γ_i , respectively, which I allow to vary across households and which sum to unity. The share of income assigned to each good or service is captured by these parameters. Income is captured through the wage w_i which the household earns through employment in either of two

³⁷I remove inter-city migration so that labour supply remains fixed in the model.

sectors.

$$\max_{\{C_i, H_i, I_i\}} U_i^M = C_i^{\alpha_i} (H_i - \bar{h})^{\beta_i} \left(I_i - \frac{1 - \theta_n^{\mu_i}}{\gamma_n} \right)^{\gamma_i} \quad (18)$$

$$\text{subject to } w_i = C_i + r_n H_i + p_n I_i$$

Utility is modelled on the assumption that households possess Stone-Geary preferences for housing such that utility is obtained only after a minimum level of consumption has been satisfied. Households require this minimum level for survival according to the term $\bar{h}$ which is homogenous across all households.³⁸ I model housing consumption in this manner for two reasons: as i) this matches the true approach taken by households in a property search (Eeckhout et al., 2014; Grossmann et al., 2021), and ii) it allows households to trade off expenditure on housing for home improvements attained through the consumption of investment services. I assume the constraint to be binding such that households facing this trade-off may always reserve income for consumption, such that $C_i > 0$. Those that opt for lower levels of housing demand may direct resources to the consumption of investment services to upgrade their housing stock. Household investment is determined by one's own demand for investment services I_i and by the innate *investment potential* provided by a neighbourhood $\frac{1 - \theta_n^{\mu_i}}{\gamma_n}$. Drawing on the public economics literature Calabrese et al., (2006), investment potential is composed of two terms. Firstly, it is a negative function of θ_n which captures the average energy efficiency of the neighbourhood E_n relative to the average across the wider city $\bar{E}$, such that³⁹

$$\theta_n = \left(\frac{E_n}{\bar{E}} \right) \quad (19)$$

This is measured on a continuum $\theta_n \in (0, \infty]$ where neighbourhoods more energy efficient than the city-wide average achieve $\theta_n > 1$. In this event, investment potential becomes negative and reflects the fact that unrealised efficiency is lower and so is the scope for investment. A second explanation – modelled in line with Rossi-Hansberg et al., (2010) –

³⁸For example, a roofed dwelling with access to an energy source for heating and food preparation.

³⁹Where E_n captures the similarity in energy efficiency scores between each reference property i and its neighbours j with distance weight $\omega_{i,j}$, such that $\theta_n = \sum_i \left[\sum_{j=1}^J \omega_{i,j} E_i - \bar{E} \sum_{j=1}^J \omega_{i,j} \right]$

is that neighbourhoods with historically high levels of investment crowd out future investment by allowing buyers to free-ride on the back of other's efforts. The extent to which buyers are aware of the neighbourhood-city differential is captured by the term μ_i which models buyer's attentiveness to energy efficiency scores across the neighbourhood and which allows investment potential to vary across households. In other words, the more a household values residence within an energy-efficient neighbourhood, the higher the perceived investment potential. Energy efficiency at the neighbourhood-level is defined as the neighbourhood effect for movers. Like in the empirical exercise where movers make location decisions based on information reported in property listing adverts, I assume that this group can observe the characteristics of prospective peers which then influences the decisions they take with regards to housing and demand for investment services. Secondly, investment potential is a negative function of the investment preferences of the wider neighbourhood γ_n . This acts in a similar way in that it captures the inclination of buyers to free-ride. Linking back to chapter three, I distinguish between the two in that one can observe the former through information reported in property listings, while buyers are left to infer the latter. Both terms are shown to be exogenous in the empirical analysis when reported in property listings; random variation in the timing and spatial dispersion of property listings sees to this. Investment potential is then a function of the two inputs which directly enter the household's utility function, allowing them to choose neighbourhoods based on their preferences for residential energy efficiency. I show later in the chapter that buyers pay more to reside in the most energy efficient neighbourhoods and that this leads to an increase in investments undertaken by both movers and non-movers.

I develop the model in such a way that movers receive dis-utility when their investment demand is lower than the potential – or what would be optimal for their chosen market. Investment potential can take negative values (as we will see later) which corresponds to individuals forgoing investments in favour of consuming more of housing or the freely-traded good in the most energy efficient neighbourhoods. Such a scenario ensures that households only choose locations in which their own preferences match that of the neigh-

bourhood, inferring that $\gamma_i \equiv \gamma_n$. I only consider in this model households that choose neighbourhoods that represent a suitable match for their investment preferences. This a reasonable assumption and one that matches the results of chapter three, finding that buyers choose locations that align with their own preference. Including investment potential within the model allows me to determine how property prices are affected by the energy efficiency level and investment demand within a neighbourhood, given that this term enters into the demand function of all three goods. Maximising utility subject to the prices and income highlights this in the demand functions presented in Eq.(20)

$$\begin{aligned}
C_i &= \alpha_i(w_i - r_n \bar{h} - p_n \frac{1 - \theta_n^{\mu_i}}{\gamma_n}) \\
H_i - \bar{h} &= \frac{\beta_i(w_i - r_n \bar{h} - p_n \frac{1 - \theta_n^{\mu_i}}{\gamma_n})}{r_n} \\
I_i - \frac{1 - \theta_n^{\mu_i}}{\gamma_n} &= \frac{\gamma_i(w_i - r_n \bar{h} - p_n \frac{1 - \theta_n^{\mu_i}}{\gamma_n})}{p_n}
\end{aligned} \tag{20}$$

Demand in all cases increases linearly with the income disposable after consuming the remaining goods and is weighted by the preference for the respective good. Allowing housing demand to be a function of the investment potential that varies by household highlights the clear trade-off between the two. One can write the demand for investment services as

$$I_i = \frac{\gamma_i \left(w_i - r_n \bar{h} - p_n \left(\frac{1 - \theta_n^{\mu_i}}{\gamma_n} \right) \right)}{p_n} + \left(\frac{1 - \theta_n^{\mu_i}}{\gamma_n} \right) \tag{21}$$

where demand is equal to the amount demanded by the households (as dictated by their preferences) plus that arising from investment potential. One can see that investment is crowded out in the most efficient neighbourhoods ($\theta_n^{\mu_i} > 1$). Taking the earlier fact that the model constrains movers to purchasing in neighbourhoods with similar preferences to their own ($\gamma = \gamma_n$), one can express this purely in terms of household preferences

$$I_i = \frac{\gamma_i (w_i - r_n \bar{h})}{p_n} + (1 - \theta_n^{\mu_i}) \left(\frac{1 - \gamma_i}{\gamma_i} \right) \tag{22}$$

Demand for investment services increases with the disposable income attributed to investment services and decrease with the price. Captured by the first term is the household's baseline level of demand weighted by preferences for investment. The second term instead captures how the average energy efficiency of the neighbourhood compared with the city-wide average (the neighbourhood effect) impacts one's demand. Taking the first order derivative with respect to θ_n shows that the demand for investment services is a decreasing function of the neighbourhood's energy efficiency. This is larger for higher levels of μ_i – in other words, for households who are more attentive to energy efficiency scores in the surrounding area and who are more likely to be aware of the surrounding historic investment stock. Taking the second order derivative with respect to θ_n illustrates the importance of this attentiveness; highlighting the presence of a U-shaped relationship between investment and a neighbourhood's energy efficiency. The parameter therefore dictates the extent to which buyers are knowledgeable about their potential and the investments they are willing to undertake.

$$\frac{\partial I_i}{\partial \theta_n} = -\mu_i \theta_n^{\mu_i-1} \left(\frac{1 - \gamma_i}{\gamma_i} \right) \quad \frac{\partial^2 I_i}{\partial \theta_n^2} = (1 - \mu_i) \mu_i \theta_n^{\mu_i-2} \left(\frac{1 - \gamma_i}{\gamma_i} \right)$$

The degree to which investment varies with attentiveness depends also on one's own preferences γ_i which weight the impact of one's potential on investment. In this case, the neighbourhood effect dictates that households choosing more efficient locations should match the investments of their neighbours but that the incentive to further it is minimal when viewing it in isolation. Highlighted in this section then is that investment preferences could be a key driver in one's location decision, given that one's demand for investment services is shown to vary across neighbourhoods depending on its energy efficiency. Location decisions are then a product of the household's remaining income. I now explore the extent to which the costs of obtaining such services influence this demand and, hence, location.

4.3.2 Investment Services

Demand functions derived in the previous section highlight the interconnectedness of housing and investment demand. Equilibrium property prices – where supply meets demand

– would then be a function of buyer’s investment demand, but household consumption is constrained by the price of investment services. These prices are determined in equilibrium by suppliers. Previous research has evidenced that investment has limits, most notably in the context of price ceilings that limit returns on investment and the subsequent home improvements undertaken (Dixit, 1991). This chapter recognises that housing investment is not limitless and makes the assumption that households only invest what they are confident they will make back by way of price returns at sale – captured by the investment potential term. Investment undertaken beyond this point is deemed risky and the returns uncertain.

Without information on ceiling price, I instead augment the approach of Severen, (2023) and develop a formula for the price of investment services determined by suppliers. Firms supply services in a competitive market using two inputs: materials M_n and labour L_n^I , expressed in Eq.(23). Materials can be thought of as the inputs to make home improvements, such as wall insulation, roofing tiles, or windows. Demand for labour and materials varies by neighbourhood, and so do the respective prices. The price of labour varies owing to the fact that individuals can either work in this sector or in the production of the final good; the price of materials varies owing to the supply of materials and the endowment in each neighbourhood (see later). I denote the wage paid in neighbourhood n as w_n . Like in the case of materials, this speaks to the empirical analysis whereby labour at one end of the London TTWA would be unlikely to offer services at the opposite end, creating regional price-based competition. Included in the model is a productivity term λ_i which captures conditions that make delivering investment services in certain locations more productive (i.e. owing to differences in topology, planning regulations locally, or local knowledge). One could also think of productivity as function of the number of similar investments undertaken locally, such that the search and monetary costs of procuring labour and materials is reduced.

$$I_n = \lambda_n L_n^{I^\Omega} M_n^{1-\Omega} \quad (23)$$

Firms provide investment services to maximise profits subject to the cost of inputs and

their budget C_n ; one can then write the firm's production function in the form of a zero profit condition in Eq. (24)

$$\pi = p_n \lambda_n L_n^{I\Omega} M_n^{1-\Omega} - P_n^m M_n - w_n L_n^I = 0 \quad (24)$$

Where p_n represents the price set for services and $p_n \lambda_n L_n^{I\Omega} M_n^{1-\Omega}$ the revenue obtained.

Taking first-order conditions returns

$$\begin{aligned} \frac{\partial \pi}{\partial L_n^I} &= \Omega p_n \lambda_n L_n^{I\Omega-1} M_n^{1-\Omega} - w_n = 0 \\ \frac{\partial \pi}{\partial M_n} &= (1 - \Omega) p_n \lambda_n L_n^{I\Omega} M_n^{-\Omega} - p_n^m = 0 \end{aligned}$$

Combining the two conditions provides the following identity showing that the demand for both inputs is inversely related to its respective cost and positively related to its productivity.

$$\frac{M_n}{L_n^I} = \frac{1 - \Omega}{\Omega} \frac{w_n}{p_n^m} \quad (25)$$

Rearranging this equation for M_n and substituting into the budget constrain returns the demand function for both inputs, such that

$$L_n^{*I} = \frac{\Omega C_n}{w_n} \quad M_n^* = \frac{(1 - \Omega) C_n}{p_n^m} \quad (26)$$

Demand for inputs is also positively related to the firm's budget, allowing firms to deploy more resources as their budget increases. Taking the first order condition for labour and substituting in the demand equations in Eq.(26), we can back out an equation describing the price of investment services in each neighbourhood.

$$p_n = \frac{w_n}{\Omega \lambda_n} \left(\frac{L_n^I}{M_n} \right)^{1-\Omega}$$

$$p_n = \frac{w_n^\Omega}{\lambda_n \Omega^\Omega (1 - \Omega)^{1-\Omega}} (p_n^m)^{1-\Omega} \quad (27)$$

Eq.(27) represents the price of investment services supplied. Prices increase with the cost of both labour and material inputs and decrease when these inputs become more productive. These costs vary across neighbourhoods which subsequently alters the distribution of investment across the city as the costs influences consumer demand. The interconnectedness of investment and housing demand means that changes in the supply of investment services alters the demand for housing and affects the spatial sorting of households across the city (though, this chapter is not directly concerned with this). The same is true when the price of supplying investment falls, as is the case in neighbourhoods that have location-specific factors that make the installation of said services more productive. Investment ought to be higher in locations that are cheaper to invest in or in locations that have lower housing costs.

Currently I cannot characterise the equilibrium price of investment services because the price of both materials and labour are endogenous to the model. For material inputs, I obtain the equilibrium price by introducing a sector that supplies materials to the firms providing investment services. I model supply using a flexible function described in Eq.(28).⁴⁰

$$M_N^S = \bar{M}_n (p_n^m)^\eta \quad (28)$$

Each neighbourhood has a fixed endowment of materials $\bar{M}_n$ from which supply then adjusts according to the elasticity of supply of materials η . A benefit of this approach is that one can remain agnostic as to the budget constraint faced by suppliers. One can also assume that developers are based out-of-town and abstract from issues around the repatriation of profits into the city.⁴¹ The structure of the equation can be better seen if one takes logs.

$$\ln(M_n^S) = \ln(\bar{M}_n) + \eta \ln(p_n^m) + \epsilon_n \quad (29)$$

⁴⁰A number of papers have adopted this functional form to model the supply of housing (i.e. Saiz, (2010); Diamond, (2016); Monte et al., (2018); Quentel, (2024)). Materials are essentially a component of housing construction and can be appropriately modelled in this way.

⁴¹If this were not the case, suppliers would invest in the city, creating general equilibrium effects.

From the regression we see that the endowment of materials is a constant and η captures the elasticity of materials supply with respect to material prices in neighbourhood n . This approach then captures how responsive developers are to prices in a locality and whether this encourages buffers in supply. Neighbourhoods with higher elasticities experience greater increases in materials that add to the existing stock. For the basis of the model, η_n is assumed inelastic. Setting both the supply and demand for materials equal to each other, one can recover an equilibrium price for materials

$$p_n^m = \left[\frac{(1 - \Omega)C_n}{\bar{M}_n} \right]^{\frac{1}{1+\eta}} \quad (30)$$

The resulting price depends positively on the productivity of using materials in providing investment services and on the budget allocated to employing inputs. A greater endowment of materials acts as an increase in supply within the neighbourhood and lowers the final price. Now, the price of investment services in Eq.(27) only depends endogenously on the wages paid to labour w_n .

4.3.3 Goods Production

To ensure that the equilibrium can be characterised, I model the production of the final good (that serves as the numeraire) such that wages can be determined for labour in both sectors. Wages determined in this sector clear the labour market by allocating workers to employment in either the production of the final good or the provision of investment services. The production sector is of lesser importance for the conclusions of the model and has no application to the reduced form estimates. For that reason, I abstract from complications and assume that production within a neighbourhood y_n is the product of the labour employed in that sector L_n^T and some productivity term that captures the location-specific factors z_n that make some neighbourhoods better at producing the good, such that

$$y_n = z_n L_n^T \quad (31)$$

As the final good acts as the numeraire, the price can be normalised to unity. Production of the final good under perfect competition implies that revenue must be equal to the cost

$$z_n L_n^T = w_n L_n^T \quad (32)$$

Expressing this as a profit function by subtracting costs from revenue and taking the derivate with respect to labour returns

$$\frac{\partial}{\partial L_n^T} z_n L_n^T - w_n L_n^T = z_n - w_n = 0$$

Which means that the wages paid in the production sector are equal to each location's productivity.

$$w_n = z_n \quad (33)$$

This would imply that production concentrates in the most productive neighbourhoods. It further implies that, as wages must equalise between sectors, wages paid to labour providing investment services is also equal to the location-specific productivity term, such that $w_n = w_n^I = w_n^T = z_n$. In this respect, wages are equal across sectors but vary between neighbourhoods in the city. This creates competition between the two sectors in the land market as firms chase the most productive locations. I later characterise the equilibrium of the economy and show that the labour market clearing condition requires employment in both sectors be equal to the total resident population of the city.

4.3.4 Housing and Investment Demand

Now we have the ingredients to compute the equilibrium price of investment services and to characterise the equilibrium of the economy. Before doing so, I turn my attention to the housing market by exploring the effects on housing demand of an increase in i) the demand for investment services, and ii) the neighbourhood's investment potential (neighbourhood effects). Motivating this model is the need to explore how mover's demand for investment services is affected by neighbourhood effects, and how this subsequently affects household's property bids and location decisions. Earlier computations of the buyer demand

functions show that housing consumption is a function of the investment potential and the prevailing market price – providing a clear channel through which prices change. A limitation of this approach is that the Stone-Geary structure imposes linearity in demand with respect to the price of each good. Property prices enter both directly and indirectly (through their effects on disposable income). In other words, they enter the housing demand equation multiple times – in the denominator and in the numerator – making it difficult to equate prices at the equilibrium where demand meets supply. To overcome this, in the following section I impose a market clearing condition that allows me to back out equilibrium property prices. A limitation is that housing supply is exogenous to the model. Before I characterise the equilibrium in the property market, I provide intuition here on the expected effects on both housing demand and on prices.

I make the assumptions that housing supply is exogenous and fixed. While the latter is a strong assumption, London is notorious for being a land-constrained city. Surrounded by greenbelt land, the city has only increased the supply of new housing by 1% per year on average between 2009-2019.⁴² I make this assumption in line with Glaeser and Gyourko, (2005) who show that housing stock is slow to adjust in US cities owing to its slow depreciation over time. The authors imply that supply as a result can be treated as fixed in the short-run. In their seminal paper, and with focus on US cities, Saiz, (2010) shows that the supply elasticity of housing is also relatively low and inelastic in land-constrained cities and in those that are larger. London fits this billing. Furthermore, Ball et al., (2010) show that supply elasticity is actually larger in US cities owing to less restrictive regulation – suggesting that housing supply in London may in fact be even more inelastic than previously determined. London sits within a highly restrictive regulatory system dictated by strict planning laws – given its size and constrained land. While the assumption of exogeneity simplifies the model and its solution, one would expect (based on evidence) the housing supply curve to be inelastic and to shift very little in actuality. A close-to-fixed housing supply curve means that any price increases induced by shifts in demand are less likely to be offset by an increase in supply.

⁴²Figure calculated using data on new completions from Planning London and live tables on dwelling stock sourced from MHCLG.

If this is the case, I can link property price increases to changes in investment potential and the demand for investment services by showing that: i) the housing demand curve takes the normal, downward-sloping shape, and ii) the housing demand curve shifts outward (inward) in response to an increase in neighbourhood energy efficiency (demand for investment services). The former shows that the demand curve behaves in the expected way, while the latter would imply the price change. To show that the demand curve is well-behaved, I take the derivative of the housing demand function in Eq.(20) with respect to the property price

$$\frac{\partial H_i}{\partial r_n} = -\beta_i \bar{h} r^{-1} - \beta_i \left(z_n - r_n \bar{h} - p_n \left(\frac{1 - \theta_n^{\mu_i}}{\gamma_n} \right) \right) r_n^{-2}$$

From the derivative we can see that the demand curve is downward sloping as increases in property prices across a neighbourhood reduce the demand for property. Any increase in demand would induce a rightward shift in the curve and an increase in prices. To make this a function of both the demand for investment services and the neighbourhood effect, I rearrange the demand equation for the former in Eq.(20) to become

$$p_n \frac{(1 - \theta_n^{\mu_i})}{\gamma_i} = p_n I_i - \gamma_i (w_i - r_n \bar{h}) + p_n (1 - \theta_n^{\mu_i})$$

We can see then that investment potential is smaller in the most efficient neighbourhoods (as homes are already closer to achieving their potential) and is an increasing function of the spending on investment services by households. Substituting this back into the housing demand equation means that we can express this as a function of the two

$$H_i = \frac{\beta_i (z_n + (\gamma_n z_n - p_n I_i) - p_n (1 - \theta_n^{\mu_i}))}{r_n} + \bar{h} (1 - \beta_i (1 + \gamma_n)) \quad (34)$$

$$\frac{\partial H_i}{\partial I_i} < 0 \quad \text{and} \quad \frac{\partial H_i}{\partial \theta_n^{\mu_i}} > 0 \quad (35)$$

A fixed supply of housing would imply that changes in housing demand are analogous to changes in property prices. Prices are a positive function of the wage z_n paid to labour

employed in the production of the final good and investment services. Prices are also increasing in the share of income attributed to the consumption of investment services $\gamma_i z_n$. However, the actual spending on investment services the household is willing to undertake leads to a reduction in prices $p_n I_i$. The two are essentially equivalent, but for the fact that spending on investment services is driven by the respective price. That the two have opposing signs could be attributed to the effect of the price of investment services and the fact that households discount their property bids when faced with financing costlier improvements. Such a situation corroborates the findings from chapter two that suggests buyers discount their bids for properties further from their potential energy efficiency in order to finance further home improvements.

The investment potential term has a negative relationship with price. Like before, this suggests that areas with higher investment potential ($\theta_n^{\mu_i} \in (0, 1]$) observe price discounts as buyers look to fund improvements. However, the same buyers actually pay more for neighbourhoods with higher levels of energy efficiency that have lower investment potential $\theta_n^{\mu_i} > 1$. As previously stated, I only model the scenario for which buyers choose neighbourhoods that represent a good match for investment preferences; this result would therefore suggest that buyers choosing more efficient neighbourhoods do so on the basis of the historic investments undertaken that align with their tastes. A second explanation would be that properties in these locations are more likely to be of better quality and reflect the fact that household's perceive tangible benefits to locating within them. In chapter three I find that this benefit arises from higher expected returns from investment. One can see this in the housing demand equation if we observe the case where $z_n \gamma_i > p_n I_i$. When the income set aside for consuming investment services is greater than that required to align with the neighbourhood (captured by $p_n I_i$), households are expected to increase their housing bids. This creates a setting in which households complete these necessary energy efficiency improvements, but expect higher returns from those improvements that go beyond those of the wider neighbourhood – thereby bidding higher in order to secure this type of housing.

Each neighbourhood comes with an investment requirement that the mover must match. I show in the previous paragraph that exceeding this requirement leads movers to increase their property bids. By once again manipulating the demand equation for investment services, one can model the demand for such services made above-and-beyond those driven by a neighbourhood's investment potential.

$$I_i - \frac{(1 - \theta_n^{\mu_i})}{\gamma_i} = \frac{\gamma_i (w_i - r_n \bar{h})}{p_n} - (1 - \theta_n^{\mu_i}) \quad (36)$$

This is depicted by the right-hand side of the equation. We can see that these additional investments are increasing in energy efficient neighbourhoods, reflecting that these neighbourhoods do in fact offer greater returns to investment; returns that can be seen in Eq.(34) for greater levels of energy efficiency.

I show in this section that buyers increase their property bids for more energy-efficient neighbourhoods and decrease their bids in response to an increase in the required investment spend. The former aligns with findings from chapter three while the latter aligns with findings from chapter two. The neighbourhood effect explains the difference in the two findings in that it alters movers' preferences. Movers anticipate returns to investment on the improvements that exceed those undertaken across the neighbourhood, which are financed by the household's remaining investment budget. These investments are undertaken in the most efficient neighbourhoods owing to the expectation of greater returns in areas where property prices are already higher (and reflect energy efficiency). Movers increase their property bids in order to secure housing in this location. As such, this will affect the locations in which movers choose to locate as they trade-off investment returns with property prices in different neighbourhoods.

4.3.5 Characterising Equilibrium

I now characterise the equilibrium in the economy. Equilibrium requires that the markets for labour, housing, and investment services clear. I first present the market clearing conditions for the labour and housing markets before characterising the equilibrium for the consumption of investment services.

Labour market clearing requires that the labour employed in the production of the final good L_n^{*T} and in providing investment services L_n^{*I} is equal to the population of the economy $\bar{L}$, such that

$$\bar{L} = L_n^{*T} + L_n^{*I} \quad (37)$$

Substituting in the equilibrium employment for investment services in Eq.(26) and rearrange for L_n^{*T} returns the equilibrium employment in final good sector that clears the market

$$L_n^{*T} = \bar{L} - \frac{\Omega C_n}{z_n} \quad (38)$$

An increase in the local productivity (wages) of delivering investment services increases the wage bill for these firms and encourages more labour to transition into the sector. Conversely, a decrease in wages would steer more workers into production of the final good. A similar result occurs as labour becomes less productive relative to materials in the provision of investment services (Ω); in this case, fewer labour units are required to turn around the same output and more workers migrate to the final good sector.

In modelling the provision of investment services, I derived the demand equations for labour and materials within the sector and used these to construct an equation detailing the price of investment services as function of the price of the respective inputs. The price of materials derived in Eq.(30) and the prevailing wage for labour derived in Eq.(33) can now be incorporated to characterise the equilibrium price. Substituting both equations into the price of investment services in Eq.(27) returns

$$p_n = \frac{z_n^\Omega C_n^{\frac{1-\Omega}{1+\eta}}}{\lambda_n \Omega (1-\Omega)^{\frac{\eta(1-\Omega)}{1+\eta}} \bar{M}_n^{\frac{1-\Omega}{1+\eta}}} \quad (39)$$

The price of investment services increases with the wage paid to labour and with the budget of the providing firm. Both represent an increase in either the cost, or the use,

of inputs. An increase in the productivity of providing these services decreases the equilibrium price for consumers. Incorporating the price of materials now means that the neighbourhood's endowment of materials $\bar{M}_n$ enters the pricing equation. Neighbourhoods with a greater supply offer lower prices to consumers, which depends on both the elasticity of supply of materials η and on the productivity with which firms transform them into investment services Ω – for which prices are decreasing in both.

The consumption of investment services in equilibrium depends on property prices which is a component of the household's disposable income. Before defining the equilibrium, we must clear the housing market to determine the equilibrium property price. The market clears when the housing consumed is equal to that supplied by the neighbourhood $\bar{H}_n$.

$$\bar{H}_n = \sum_{i \in n} H_i \quad (40)$$

Substituting the housing demand equation into the right-hand side of this clearing condition and summing elements at the household-level means I can express prices in terms of housing supply.⁴³

$$\bar{H}_n = \frac{\beta_i(z_n - r_n N_n \bar{h} - p_n \sum_{i \in n} \frac{1 - \theta_n^{\mu_i}}{\gamma_n})}{r_n} + N_n \bar{h}$$

By equating the supply and demand for housing, rearranging for property prices r_n returns the price.

$$r_n = \frac{z_n \sum_{i \in n} \beta_i - \frac{p_n}{\gamma_n} \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i})}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i}$$

Substituting in the equilibrium price of investment services p_n gives us the equilibrium price

$$r_n = \frac{z_n \sum_{i \in n} \beta_i - \frac{z_n^\Omega C_n^{\frac{1-\Omega}{1+\eta}}}{\gamma_n \lambda_n \Omega^\Omega (1-\Omega)^{\frac{\eta(1-\Omega)}{1+\eta}} \bar{M}_n^{\frac{1-\Omega}{1+\eta}}} \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i})}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \quad (41)$$

⁴³Constants, such as $\bar{h}$, are summed by multiplying by the number of households in n .

The equilibrium price for property in Eq.(41) captures the effect of income and investment potential in the numerator and the net housing supply of the neighbourhood in the denominator. An important prediction comes from this condition: that an increase in the price of investment services causes buyers to increase their price offers for property exclusively in energy-efficient neighbourhoods. For energy efficient neighbourhoods where $\theta_n^{\mu_i} > 1$, the sign on numerator flips and turns positive, meaning that prices increase. Bidding higher for a good as the price increases is not economically rational – unless movers expect greater investment returns in the most efficient neighbourhoods to recover the higher property price paid. This matches the predictions from earlier in the chapter, as well as from the empirical setting in chapter three.

Now I have both the equilibrium price of property and investment services, I can substitute them into the equation for investment demand in Eq.(22). The equation shows that investment demand contains a baseline and spillover component arising from the neighbourhood effect that drives the neighbourhood's investment potential. For the former, households invest in line with their preferences; for the latter, they undertake less investment in neighbourhoods with lesser investment potential. Incorporating the equilibrium prices for property and investment services allows me to explore their effects on the consumption of said services. By first incorporating the price of investment, we get

$$I_i = \frac{\gamma_i (z_n - r_n \bar{h}) \lambda_n \Omega^\Omega (1 - \Omega)^{\frac{\eta(1-\Omega)}{1+\eta}} \bar{M}^{\frac{1-\Omega}{1+\eta}}}{z_n^\Omega C_n^{\frac{1-\Omega}{1+\eta}}} + (1 - \theta_n^{\mu_i}) \left(\frac{1 - \gamma_i}{\gamma_i} \right)$$

The baseline level of investment is now determined by buyer's preferences, their income, and the prevailing market price for housing and investment services. Compared to the equation for the price of investment services, the numerator and denominator have flipped. Intuitively, factors that lower prices for the providing firm (on the denominator) increase the investment demand for the consumer and appear on the numerator – such as the productivity of providing the services and the endowment of material inputs. The opposite is true for factors that increase firm prices. Interestingly, prices do not affect the impact of the neighbourhood effect on investment; this is driven purely by personal preferences and the investment potential of the neighbourhood. Given that market forces do not

determine its magnitude, the equation for the neighbourhood effect would suggest that social norms play a role in propagating investments. I show this by expressing the equation in terms of the wider neighbourhood preference for investment services. By incorporating the equilibrium property price, we get the final equilibrium consumption of investment services

$$I_i^* = \gamma_n \frac{z_n^{1-\Omega} \lambda_n \Omega (1-\Omega)^{\frac{\eta(1-\Omega)}{1+\eta}} \bar{M}_n^{\frac{1-\Omega}{1+\eta}}}{C_n^{\frac{1-\Omega}{1+\eta}}} \left(1 - \frac{\bar{h} \sum_{i \in n} \beta_i}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \right) \quad (42)$$

$$+ \frac{\bar{h} \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i})}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} + (1 - \theta_n^{\mu_i}) \left(\frac{1 - \gamma_n}{\gamma_n} \right)$$

As preferences for investment services within the neighbourhood increase (γ_n), buyers have split preferences between neighbourhoods. Investment in the least-efficient neighbourhoods increases at a decreasing rate, whereas investment in the most-efficient neighbourhoods decreases at a decreasing rate. Part of the household's investment decisions are tied to that of their neighbours, as well as prices. This makes conformity a key part of the model and one which drives the resulting uptake of home improvements.

The consumption of investment services in equilibrium – like in the previous derivations – reduces with the cost of providing them. However, we now see that an increase in wages increase the consumption at a rate that depends on the productivity of turning labour units into investment services (Ω). We also now see that the spillover term depends on preferences for investment, like before, and on the net supply of housing ($\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i$). As net supply (and the denominator) increases, the strength of the spillover on investment dampens – owing to a wider availability of housing on the market that widens the buyer's choice set and dampens the need for investment. Taking the derivative of the consumption of investment services with respect to β_i , we see that housing and investment are substitutes. Buyers divert resources away from investing in their property when spending a larger share of their income on housing. This is more evident in the most efficient neighbourhoods and underscores the fact that buyers increase their housing budget in the expectation of obtaining better quality housing in these locations. As a result, they purchase the investments made by previous owners that satisfy the investment

potential.

$$\frac{\partial I_i^*}{\partial \beta_i} = \left[\left((1 - \theta_n^{\mu_i}) - \gamma_n \frac{z_n}{p_n} \right) \left(\frac{\bar{h}(\bar{H}_n - N_n \bar{h})}{(\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i)^2} \right) \right] < 0 \quad (43)$$

One can see from the derivative that movers – in matching their chosen neighbourhood’s investment preferences – purchase better quality housing in order to avoid making improvements. However, this does not subvert their ability to make improvements as we see in Eq.(36). It is therefore the case that better quality housing is a substitute to those investments that allow the mover to attain the investment potential. One can see this from the inverse relationship between β_i and γ_n inside the derivative.

By characterising the equilibrium of the economy, I conclude three points from the model. Firstly, property prices are increasing in the energy efficiency of the neighbourhood. Secondly, this is a result of buyers undertaking levels of investment over-and-above that of the wider neighbourhood (as dictated by the investment potential). Buyers may attain this level of investment by consuming investment services or by purchasing housing of better quality. Finally, these investments are undertaken exclusively in the most efficient neighbourhoods – dictated by the fact in point two that buyers pay higher prices for property exhibiting improvements in these neighbourhoods. Higher levels of investment locally are observable across the neighbourhood. I now turn my attention towards exploring the effects of this on the investment levels of the wider neighbourhood.

4.3.6 Never Movers

Given that movers bid higher for property in more energy efficient neighbourhoods in anticipation of higher returns in such, I now turn my attention to exploring how this impacts the investment decisions of non-movers. I make three assumptions that differentiate this group from the movers: i) non-movers do not take decisions with respect to housing, meaning it only enters the household’s budget constraint; ii) they do not observe the quality of the neighbourhood (as reported through listings in the reduced form exercise), so neighbourhood effects are observed and not inferred; iii) I assume that the two groups – movers and non-movers – are mutually exclusive such that households can’t

transition between groups, allowing me to model decisions separately between the groups in the same time period. Given that non-movers do not take decisions over housing (and thereby do not explicitly chase locations delivering greater investment returns), I deviate from the Stone-Geary functional form that requires households to consume a minimum housing requirement. I instead model preferences using a Cobb-Douglas functional form as depicted by Eq.(44)

$$\max_{C,I} U_{it}^{nm} = \left(\frac{C_i^{nm}}{\psi} \right)^{\psi} \left(\frac{(I_i^{nm})^{\Gamma_i}}{1 - \psi} \right)^{1-\psi} \quad (44)$$

$$\text{subject to } z_n - r_n H_i = C^{nm} + p_n(I_i^{nm})$$

Households maximise consumption of the numeraire good and investment services (I_i^{nm}) as in the case of movers, but subject to income after accounting for housing costs. The Cobb-Douglas set up assumes that the share of household income attributed to consumption of both goods is fixed across locations. This provides a better setting to motivate how investment differs by neighbourhood as a result of neighbourhood effects. In the empirical setting I show that the price premium paid by movers for efficient neighbourhoods represent the minimum expected returns from investment, which differs by household, and which are driven by their attentiveness to energy efficiency scores. I now use the consumption of investment services by movers to model the investment decisions made by non-movers. This represents the neighbourhood effect facing non-movers that I show later transmits through exposure to movers. I define this exposure by taking the equilibrium consumption of investment services by movers across a neighbourhood in Eq.(42) and reassigning it to a new variable such that $\nu_j = I_i^*$. One can define the exposure of non-movers to local investment as

$$\Gamma_i = \sum_{j=1}^J e^{\delta \tau_{ij}} \nu_j \quad (45)$$

Captured by Γ_i is the distance-weighted sum of all investment services consumed by one's neighbours, where distance between non-mover i and neighbour j is measured by τ ($i \neq j$). The importance of distance in determining exposure is captured by δ . Defining

the neighbourhood effect in such a way captures the fact that proximity to investing households matters as this allows households to observe the investments made by movers. Maximising with respect to both goods and rearranging returns the non-mover's demand function for investment services.

$$I_i^{nm} = \frac{(1-\psi)}{\psi} \left(\Gamma_i + \frac{\psi}{(1-\psi)} \right) \frac{(z_n - r_n H_i)}{p_n} \quad (46)$$

Solving for demand returns an equation of two parts. The first denotes how both individual preferences and the preferences of movers for investment services impact the demand of non-movers. Demand for investment services increases with the share of income attributed to such activities $(1-\psi)$ and decreases when more resources are diverted to consumption of the final good. Consumption also increases with Γ which captures the distance-weighted sum of investment spending within the neighbourhood. This suggests that neighbourhood effects are stronger when investments happen close-by or in densely-population neighbourhoods (as a result of the distance weighting). The second term provides the level of investment services consumed at baseline by the households. This reflects the investment undertaken by households regardless of preferences and dependent only on their disposable income. Incorporating the equilibrium price of investment and the equilibrium property price into this equation returns

$$I_i^{nm*} = \frac{(1-\psi)}{\psi} \left(\Gamma_i + \frac{\psi}{(1-\psi)} \right) \frac{1}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \quad (47)$$

$$\left(\frac{z_n^{1-\Omega} \lambda_n \Omega^\Omega (1-\Omega)^{\frac{\eta(1-\Omega)}{1+\eta}} \bar{M}_n^{\frac{1-\Omega}{1+\eta}}}{C_n^{\frac{1-\Omega}{1+\eta}}} \left(\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i - H_i \sum_{i \in n} \beta_i \right) + \frac{H_i}{\gamma_n} \left(\sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i}) \right) \right)$$

which represent the equilibrium consumption of investment services for non-movers. Investment increases with the share of income devoted to consuming investment services $(1-\psi)$ and decreases as households divert resources towards consuming the final good. Investment also increases as household income increases. Like before, I observe the same relationship whereby reductions in the price of purchasing investment services increase their uptake; this is unrelated to energy efficiency. The equation provides the same conclusion for both movers and non-movers: that overall levels of investment are lower in the

most efficient neighbourhoods that have the lowest investment potential. For non-movers, this is expected, as it is unlikely that they are aware of energy efficiency scores across their neighbourhood; this information is more available to buyers who find it reported in property listing adverts. The neighbourhood effect does not operate through this channel. Instead, the consumption of investment services among non-movers is actually driven by their exposure to the same services consumed by movers (Γ_i), such that

$$\frac{\partial I_i^{nm*}}{\partial \Gamma_i} > 0 \quad (48)$$

The neighbourhood effect therefore operates through the exposure to movers and the extra investments they undertake in the most efficient neighbourhoods as opposed to some intrinsic characteristic of the neighbourhoods themselves. This matches the findings of chapter three in which I show that exposure to movers over time induces non-movers (with no ex-ante preferences for energy efficiency) to make their own improvements. Finally, I can show that the effect of exposure on the consumption of investment services is driven by both the price of the services and a social norm by taking its derivative with respect to the neighbourhood's preferences for investment

$$\frac{\partial \Gamma_i}{\partial \gamma_n} = \frac{z_n^{1-\Omega} \lambda_n \Omega^\Omega (1-\Omega)^{\frac{\eta(1-\Omega)}{1+\eta}} \bar{M}_n^{\frac{1-\Omega}{1+\eta}}}{C_n^{\frac{1-\Omega}{1+\eta}}} \left(\frac{\bar{H}_n - N_n \bar{h}}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \right) - \frac{1 - \theta_n^{\mu_i}}{\gamma_n^2} \quad (49)$$

This social norm is most evident in the most efficient neighbourhoods as it requires non-movers to conform with their neighbours and undertake higher levels of expenditure on investment services. The model then demonstrates that the increased incidence of home improvements evidenced in the reduced form analysis as a result of exposure are driven by households wanting to keep up with their neighbours.

In this model I show that movers with preferences for energy efficiency purchase property in the most efficient neighbourhoods and pay a premium for the privilege. This premium reflects the expectation that such neighbourhoods deliver higher returns to investment – as movers observe higher prices in these neighbourhoods prior to purchasing. Movers purchase housing in neighbourhoods for which their preferences for investment match

those of the other occupants; implying that buyers must attain a minimum investment standard dictated by the neighbourhood's potential. Derivations from the model show that returns come from investment made after this condition has been satisfied. One's choice of peers is therefore important in determining location decisions and the resulting property price. By design, this influences location decisions as movers chase higher returns. Movers choose a location and the level of investment services to consume, which are both observed by non-movers. Exposure to this investment activity induces non-movers to undertake their own investments through conformity to the social norm of the neighbourhood (alternatively referred to as "*keeping up with the Joneses*"). I finally show that this level of conformity is higher in the most energy-efficient neighbourhoods – thereby confirming that these locations foster the highest levels of investment.

4.4 Conclusion

Little is known about how the demand for investment services ex-ante impacts household property bids. Standard intuition would suggest that buyers discount their bids in expectation of financing greater investments – as chapter two would confirm. I contribute to the literature in this chapter by producing a theoretical approach to thinking about this, from which I document three findings. Firstly, movers do indeed lower their property bids in order to finance the investments that align them with their neighbours; but that they increase their bids to finance investments beyond these. Secondly, I am able to pin down the mechanism behind this to the neighbourhood effect that means movers pay higher prices for property as they chase higher returns on their investments in energy efficient neighbourhoods. Higher prices in these locations means that these further investments are exclusive to the most efficient neighbourhoods. Finally, I show that the investments undertaken by movers create a neighbourhood effect that induce non-movers to undertake their own investments. Using the market clearing conditions from the model to characterise the equilibrium of the economy, I am able to show that this neighbourhood effect works through increased exposure to investments made locally as never-movers observe that decisions taken by movers. Findings from the model confirm those of the empirical exercises in chapters two and three, suggesting that the presence of neighbourhood effects

alters household incentives with regards to housing investment and the resulting property bids. This would suggest that the model behaves well and that it captures the decisions that buyers would make in the neighbourhoods that represent a good match for their investment preferences.

The theoretical foundations of this chapter raises some interesting policy questions. Findings from the model show that neighbourhoods with higher levels of energy efficiency receive greater investment and draw higher property prices as a result of the neighbourhood effect. As we continue to understand the dynamics driving these effects, we must design effective policies that maximise the positive externalities associated with neighbourhood effects that can propagate benefits across cities and beyond the most efficient neighbourhoods. We must also recognise that such effects induce negative externalities in the form of higher prices as households face being priced out of certain markets. Future policy aiming to maximise the benefits of these investments should be mindful of the potentially negative welfare effects. Beyond the final benefits to households, future policy should explore different approaches to streamlining the delivery of investment services. The model shows that the price of such services fall with the productivity of their delivery and with decreases in the cost of inputs. Lowering the price is shown to increase investment across all neighbourhoods and would thereby allow neighbourhoods that are less efficient to catch up in terms of residential energy efficiency. Streamlining the service would also help reduce other barriers to energy efficiency investments, such as information asymmetries, home-life disruption, and concerns over returns to investment.

Finally, the model presented in this chapter provides a setup that can be applied across a range of topics in the energy literature. In the residential energy efficiency literature, for example. While I remain agnostic about the precise investments undertaken, one could indeed incorporate into the model specific investments that affect property prices by remodelling the supply of materials. A wide variety of papers explore the investments made by households into energy durables, such as air conditioning units (Dubin and D. L. McFadden, 1984; L. W. Davis et al., 2014), solar PV (Bollinger and Gillingham, 2012;

Bigler and Janzen, 2023), and home appliances (Carroll et al., 2016; Kallbekken et al., 2013), but mainly focus on the energy efficiency savings of such investments. In modifying this model, future research could explore the implications that such adoptions have on household's future adoptions and on the prices of these properties. The model itself, with simplifying assumptions, could be quantified, opening up further research possibilities. A secondary strand of literature could use this model to explore – on the back of this adoption of durables – the willingness-to-pay of specific households for properties with more energy efficient appliances/technologies. Previous work has considered the barriers to their adoption (Hausman, 1979; Metcalf and Hassett, 1999; Ansar and Sparks, 2009) – but purchasing properties with such technologies removes these barriers and presents a better case study. In tweaking the model slightly to consider ex-post investments, this could provide an interesting research question.

4.5 Derivations

4.5.1 Household Demand Functions

I start by deriving the household's demand functions for the final good, investment services, and housing. Setting the maximisation problem up as a Lagrangian using the utility function in Eq.(18) returns

$$\mathbb{L}^g = C_i^{\alpha_i} (H_i - \bar{h})^{\beta_i} \left(I_i - \frac{1 - \theta_n^{\mu_i}}{\gamma_n} \right)^{\gamma_i} + \lambda(w_i - C_i - r_n H_i - p_n I_i)$$

For simplicity of notation, let us define $\chi = \frac{1 - \theta_n^{\mu_i}}{\gamma_n}$. Taking the derivative of the Lagrangian with respect to each good returns the first order conditions

$$\begin{aligned} 1) \quad \frac{\partial \mathbb{L}^g}{\partial C_i} &= \alpha_i C_i^{\alpha_i - 1} (H_i - \bar{h})^{\beta_i} (I_i - \chi)^{\gamma_i} - \lambda = 0 \\ 2) \quad \frac{\partial \mathbb{L}^g}{\partial H_i} &= \beta_i C_i^{\alpha_i} (H_i - \bar{h})^{\beta_i - 1} (I_i - \chi)^{\gamma_i} - \lambda r_n = 0 \\ 3) \quad \frac{\partial \mathbb{L}^g}{\partial I_i} &= \gamma_i C_i^{\alpha_i} (H_i - \bar{h})^{\beta_i} (I_i - \chi)^{\gamma_i - 1} - \lambda p_n = 0 \end{aligned}$$

One obtains the demand functions by finding expressions for two out of the three variables, substituting into the constraint, and expressing the constraint in terms of the one unknown good. In this case, I first find expressions for housing and investment services by dividing 1)/2) and dividing 1)/3), such that we get $\frac{\partial U_i^M / \partial C_i}{\partial U_i^M / \partial H_i}$ and $\frac{\partial U_i^M / \partial C_i}{\partial U_i^M / \partial I_i}$, respectively. This returns

$$H_i = \frac{\beta_i C_i}{\alpha_i r_n} + \bar{h} \quad \text{and} \quad I_i = \frac{\gamma_i C_i}{\alpha_i p_n} + \chi$$

Substituting into the budget constraint and simplifying means that the constraint becomes

$$w_i = C_i \left(1 + \frac{\beta_i}{\alpha_i} + \frac{\gamma_i}{\alpha_i} \right) + r_n \bar{h} + p_n \chi$$

Which, when rearranged for C_i , becomes

$$C_i = \frac{\alpha_i (w_i - r_n \bar{h} - p_n \chi)}{\alpha_i + \beta_i + \gamma_i} = \alpha_i (w_i - r_n \bar{h} - p_n \chi)$$

I obtain this as a result of the parameters summing to unity. Substituting this into the demand equations for investment services and housing returns the final demand functions from the chapter in Eq.(20).

$$\begin{aligned} C_i &= \alpha_i(w_i - r_n \bar{h} - p_n \chi) \\ H_i - \bar{h} &= \frac{\beta_i(w_i - r_n \bar{h} - p_n \chi)}{r_n} \\ I_i - \chi &= \frac{\gamma_i(w_i - r_n \bar{h} - p_n \chi)}{p_n} \end{aligned}$$

4.5.2 Rewriting Investment Demand

Substituting χ back into the demand for investment services, one can rewrite this as

$$\begin{aligned} I_i - \left(\frac{1 - \theta_n^{\mu_i}}{\gamma_n} \right) &= \frac{\gamma_i \left(w_i - r_n \bar{h} - p_n \left(\frac{1 - \theta_n^{\mu_i}}{\gamma_n} \right) \right)}{p_n} \\ I_i &= \frac{\gamma_i (w_i - r_n \bar{h})}{p_n} + \left(\frac{1 - \theta_n^{\mu_i}}{\gamma_n} \right) - \gamma \left(\frac{1 - \theta_n^{\mu_i}}{\gamma_n} \right) \end{aligned}$$

Which, after grouping the components on the right-hand-side of the equation containing common terms, one can use to obtain the demand for investment services in Eq.(22)

$$I_i = \frac{\gamma_i (w_i - r_n \bar{h})}{p_n} + (1 - \theta_n^{\mu_i}) \left(\frac{1 - \gamma_i}{\gamma_i} \right)$$

4.5.3 Deriving the Supply of Investment Services

Firms supply investment services according to the following zero profit condition from Eq.(24)

$$\pi = p_n \lambda_n L_n^{I\Omega} M_n^{1-\Omega} - P_n^m M_n - w_n L_n^I = 0$$

Taking first order conditions of the zero profit condition with respect to both inputs

$$\frac{\partial \pi}{\partial L_n^I} = \Omega p_n \lambda_n L_n^{I\Omega-1} M_n^{1-\Omega} - w_n = 0$$

$$\frac{\partial \pi}{\partial M_n} = (1 - \Omega) p_n \lambda_n L_n^{I\Omega} M_n^{-\Omega} - p_n^m = 0$$

and dividing the conditions by each other, one can rearrange to return the following identity

$$\frac{M_n}{L_n^I} = \frac{1 - \Omega}{\Omega} \frac{w_n}{p_n^m}$$

Rearranging for the materials term and substituting into the budget constraint allows me to recover an equation describing the demand for labour

$$C_i = P_n^m \left(\frac{1 - \Omega}{\Omega} \frac{w_n}{p_n^m} L_n^I \right) + w_n L_n^I$$

$$C_i = \left(\frac{1 - \Omega}{\Omega} w_n L_n^I \right) + w_n L_n^I$$

$$L_n^{*I} = \frac{\Omega C_i}{w_n}$$

Rearranging the same identity from Eq.(25) for materials and substituting into it the demand for labour equation, I recover an equation describing the demand for materials

$$M_n^* = \frac{1 - \Omega}{\Omega} \frac{w_n}{p_n^m} L_n^{*I}$$

$$M_n^* = \frac{1 - \Omega}{\Omega} \frac{w_n}{p_n^m} \frac{\Omega C_i}{w_n}$$

$$M_n^* = \frac{(1 - \Omega) C_i}{p_n^m}$$

I use the demand equations for labour and materials to get an expression for the price of investment. Taking the first order condition for labour and rearranging

$$w_n = \Omega p_n \lambda_n \left(\frac{M_n^*}{L_n^* I} \right)^{1-\Omega}$$

$$p_n = \frac{w_n}{\Omega \lambda_n} \left(\frac{L_n^* I}{M_n^*} \right)^{1-\Omega}$$

$$p_n = \frac{w_n}{\Omega \lambda_n} \left(\frac{\Omega p_n^m}{(1-\Omega)w_n} \right)^{1-\Omega}$$

I simplify the exponents and return an equation describing the price of investment services supplied by firms from Eq.(27)

$$p_n = \frac{w_n^\Omega}{\lambda_n \Omega^\Omega (1-\Omega)^{1-\Omega}} (p_n^m)^{1-\Omega}$$

4.5.4 Deriving the Equilibrium Price of Investment Services Inputs

The price for investment services remains endogenous in both the price of labour and materials. Taking the supply of materials in Eq.(28) and setting it equal to demand returns the equilibrium price of materials

$$\bar{M}_n (p_n^m)^\eta = \frac{(1-\Omega)C_n}{p_n^m}$$

$$(p_n^m)^{1+\eta} = \frac{(1-\Omega)C_n}{\bar{M}_n}$$

$$p_n^m = \left[\frac{(1-\Omega)C_n}{\bar{M}_n} \right]^{\frac{1}{1+\eta}}$$

Deriving the Equilibrium Price of Labour

Full derivations included in the main chapter

4.5.5 Deriving the Change in Equilibrium Property Prices

I now set out the derivations to obtain the changes in property price resulting from changes in both investment and neighbourhood energy efficiency. I know from the demand

functions in Eq.(20) that housing demand takes the form

$$H_i = \frac{\beta_i(w_i - r_n \bar{h} - p_n \frac{1 - \theta_n^{\mu_i}}{\gamma_n})}{r_n} + \bar{h}$$

To show that the housing demand function behaves as normal – i.e. is downward sloping – I take the derivative of this with respect to prices r_n . This requires the use of the product rule, as the housing demand equation is a product of prices, such that

$$\underbrace{\beta_i \left(w_i - r_n \bar{h} - p_n \frac{1 - \theta_n^{\mu_i}}{\gamma_n} \right)}_{f(x)} \overbrace{r_n^{-1}}^{g(x)} + \bar{h}$$

Taking the derivative returns a negative equation, showing that demand is downward sloping in price

$$\frac{\partial H_i}{\partial r_n} = f'(x)g(x) + f(x)g'(x) = -\beta_i \bar{h} r_n^{-1} - \beta_i \left(z_n - r_n \bar{h} - p_n \left(\frac{1 - \theta_n^{\mu_i}}{\gamma_n} \right) \right) r_n^{-2}$$

To see how prices respond to investment and energy efficiency, I rearrange the investment demand above in terms of investment potential

$$I_i = \frac{\gamma_i (w_i - r_n \bar{h})}{p_n} + (1 - \theta_n^{\mu_i}) \left(\frac{1 - \gamma_i}{\gamma_i} \right) = \frac{\gamma_i (w_i - r_n \bar{h})}{p_n} + \frac{(1 - \theta_n^{\mu_i})}{\gamma_i} - \frac{\gamma_i (1 - \theta_n^{\mu_i})}{\gamma_i}$$

Simplifying this further returns

$$p_n I_i = \gamma_i (w_i - r_n \bar{h}) + p_n \frac{(1 - \theta_n^{\mu_i})}{\gamma_i} - p_n (1 - \theta_n^{\mu_i})$$

$$p_n \frac{(1 - \theta_n^{\mu_i})}{\gamma_i} = p_n I_i - \gamma_i (w_i - r_n \bar{h}) + p_n (1 - \theta_n^{\mu_i})$$

This can be substituted into the housing demand equation and simplified to obtain Eq.(34)

$$\begin{aligned}
H_i &= \frac{\beta_i(z_n - r_n \bar{h} - p_n I_i + \gamma_i(z_n - r_n \bar{h}) - p_n(1 - \theta_n^{\mu_i}))}{r_n} + \bar{h} \\
H_i &= \frac{\beta_i(z_n - p_n I_i + \gamma_i z_n - p_n(1 - \theta_n^{\mu_i}))}{r_n} + \bar{h} - \frac{\beta_i r_n \bar{h}}{r_n} - \frac{\beta_i \gamma_i r_n \bar{h}}{r_n} \\
H_i &= \frac{\beta_i(z_n + (\gamma_i z_n - p_n I_i) - p_n(1 - \theta_n^{\mu_i}))}{r_n} + \bar{h}(1 - \beta_i(1 + \gamma_i))
\end{aligned}$$

Taking the derivative with respect to both investment and the neighbourhood energy efficiency returns the conditions from Eq.(35)

$$\begin{aligned}
\frac{\partial H_i}{\partial I_i} &= -\frac{\beta p_n}{r_n} < 0 \\
\frac{\partial H_i}{\partial \theta_n} &= \mu_i \frac{\beta p_n}{r_n} \theta_n^{\mu_i-1} > 0
\end{aligned}$$

4.5.6 Characterising Equilibrium: Investment Services Price

The price of investment services in Eq.(27) is endogenous in both the prices of materials and labour. After substituting in for the price of both labour and material inputs, I get

$$\begin{aligned}
p_n &= \frac{w_n^\Omega}{\lambda_n \Omega^\Omega (1 - \Omega)^{1-\Omega}} (p_n^m)^{1-\Omega} \\
p_n &= \frac{z_n^\Omega}{\lambda_n \Omega^\Omega (1 - \Omega)^{1-\Omega}} \left(\left[\frac{(1 - \Omega) C_n}{\bar{M}_n} \right]^{\frac{1}{1+\eta}} \right)^{1-\Omega} \\
p_n &= \frac{z_n^\Omega}{\lambda_n \Omega^\Omega (1 - \Omega)^{1-\Omega}} ((1 - \Omega) C_n)^{\frac{1-\Omega}{1+\eta}} \bar{M}_n^{-\frac{1-\Omega}{1+\eta}}
\end{aligned}$$

The only common term in the equation is $(1 - \Omega)$, for which the indices become

$$\begin{aligned}
& (1 - \Omega)^{\Omega-1} \cdot (1 - \Omega)^{\frac{1-\Omega}{1+\eta}} \\
& (1 - \Omega)^{\frac{(1+\eta)(\Omega-1)}{1+\eta}} (1 - \Omega)^{\frac{1-\Omega}{1+\eta}} \\
& (1 - \Omega)^{\frac{\Omega-1+\Omega\eta-\eta}{1+\eta} + \frac{1-\Omega}{1+\eta}} \\
& (1 - \Omega)^{\frac{\Omega\eta-\eta}{1+\eta}} \\
& (1 - \Omega)^{\frac{\eta(\Omega-1)}{1+\eta}}
\end{aligned}$$

One can then solve for the equilibrium price of investment services from Eq.(39)

$$p_n = \frac{z_n^\Omega C_n^{\frac{1-\Omega}{1+\eta}}}{\lambda_n \Omega^\Omega (1 - \Omega)^{\frac{\eta(1-\Omega)}{1+\eta}} \bar{M}_n^{\frac{1-\Omega}{1+\eta}}}$$

4.5.7 Characterising Equilibrium: Property Price

The market clearing condition $\bar{H}_n = \sum_{i \in n} H_i$ requires that the supply of housing is equal to that demanded within a neighbourhood. Knowing that demand was derived in Eq.(20), I substitute this into the clearing condition and sum over all variables determined at the household-level

$$\begin{aligned}
H_i &= \frac{\beta_i(z_n - r_n \bar{h} - p_n \frac{1-\theta_n^{\mu_i}}{\gamma_n})}{r_n} + \bar{h} \\
\bar{H}_n &= \frac{z_n \sum_{i \in n} \beta_i - r_n \bar{h} \sum_{i \in n} \beta_i - \frac{p_n}{\gamma_n} \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i})}{r_n} + N_n \bar{h}
\end{aligned}$$

For constant terms, I denote by N_n the number of households in neighbourhood n . This equation can be manipulated to achieve the final property price

$$\begin{aligned}
r_n \bar{H}_n &= z_n \sum_{i \in n} \beta_i - r_n \bar{h} \sum_{i \in n} \beta_i - \frac{p_n}{\gamma_n} \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i}) + r_n N_n \bar{h} \\
r_n \bar{H}_n - r_n N_n \bar{h} + r_n \bar{h} \sum_{i \in n} \beta_i &= z_n \sum_{i \in n} \beta_i - \frac{p_n}{\gamma_n} \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i}) \\
r_n (\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i) &= z_n \sum_{i \in n} \beta_i - \frac{p_n}{\gamma_n} \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i}) \\
r_n &= \frac{z_n \sum_{i \in n} \beta_i - \frac{p_n}{\gamma_n} \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i})}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i}
\end{aligned}$$

The equilibrium price is obtained after substituting in the price of investment services – returning Eq.(41). The full derivation is available within the chapter.

4.5.8 Characterising Equilibrium: Investment Services

Now I have obtained the equilibrium price for both investment services and property, I use this to produce a formula for the services consumed in the market. I start by substituting in property prices for ease of representation.

$$\begin{aligned}
I_i &= \frac{\gamma_i (z_n - r_n \bar{h})}{p_n} + (1 - \theta_n^{\mu_i}) \left(\frac{1 - \gamma_i}{\gamma_i} \right) \\
I_i &= \frac{\gamma_i}{p_n} \left(z_n - \left(\frac{z_n \sum_{i \in n} \beta_i - \frac{p_n}{\gamma_n} \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i})}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \right) \bar{h} \right) + (1 - \theta_n^{\mu_i}) \left(\frac{1 - \gamma_i}{\gamma_i} \right) \\
I_i &= \gamma_i \left(\frac{z_n}{p_n} - \frac{z_n \bar{h} \sum_{i \in n} \beta_i}{p_n (\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i)} + \frac{p_n \bar{h} \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i})}{\gamma_i p_n (\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i)} \right) + (1 - \theta_n^{\mu_i}) \left(\frac{1 - \gamma_i}{\gamma_i} \right) \\
I_i &= \gamma_i \frac{z_n}{p_n} \left(1 - \frac{\bar{h} \sum_{i \in n} \beta_i}{(\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i)} \right) + \frac{\bar{h} \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i})}{(\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i)} + (1 - \theta_n^{\mu_i}) \left(\frac{1 - \gamma_i}{\gamma_i} \right)
\end{aligned}$$

Substituting the equilibrium price of investment services into this equation leaves me with the consumption of investment services in equilibrium from Eq.(42)

$$I_i^* = \gamma_n \frac{z_n^{1-\Omega} \lambda_n \Omega^\Omega (1-\Omega)^{\frac{\eta(1-\Omega)}{1+\eta}} \bar{M}_n^{\frac{1-\Omega}{1+\eta}}}{C_n^{\frac{1-\Omega}{1+\eta}}} \left(1 - \frac{\bar{h} \sum_{i \in n} \beta_i}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \right) + \frac{\bar{h} \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i})}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} + (1 - \theta_n^{\mu_i}) \left(\frac{1 - \gamma_n}{\gamma_n} \right)$$

One can now see that investment is dictated by a neighbourhood's energy efficiency level, weighted both in terms of preferences for investment services γ_i and housing β_i .

I then take the derivative with respect to the latter to obtain $\frac{\partial I_i}{\partial \beta_i}$ to obtain Eq.(43). This shows how the demand for investment services responds to an increase in household preferences for housing. I simplify the equation by replacing the equation with p_n and simplifying the notation of the terms containing β_i

$$I_i^* = \gamma_n \frac{z_n}{p_n} \left(1 - \frac{U}{V} \right) + \frac{U(1 - \theta_n^{\mu_i})}{V} + (1 - \theta_n^{\mu_i}) \left(\frac{1 - \gamma_n}{\gamma_n} \right)$$

Such that $U = \bar{h} \sum_{i \in n} \beta_i$ and $V = \bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i$. For the purposes of the derivation, I then only require the terms containing β_i

$$I_i^* = \gamma_n \frac{z_n}{p_n} \left(1 - \frac{U}{V} \right) + \frac{U(1 - \theta_n^{\mu_i})}{V}$$

I take the derivative using the quotient rule $\frac{U'V - UV'}{V^2}$, knowing that

$$U' = V' = \bar{h}$$

Taking the derivatives of the individual terms

$$\begin{aligned} \frac{\partial}{\partial \beta_i} \left[\gamma_n \frac{z_n}{p_n} \left(1 - \frac{U}{V} \right) \right] &= \gamma_n \frac{z_n}{p_n} \left(-\frac{U'V - UV'}{V^2} \right) \\ \frac{\partial}{\partial \beta_i} \left[\frac{U(1 - \theta_n^{\mu_i})}{V} \right] &= (1 - \theta_n^{\mu_i}) \left(\frac{U'V - UV'}{V^2} \right) \end{aligned}$$

Combining them together, I obtain

$$\begin{aligned}\frac{\partial I_i^*}{\partial \beta_i} &= \gamma_n \frac{z_n}{p_n} \left(-\frac{\bar{h}V - U\bar{h}}{V^2} \right) + (1 - \theta_n^{\mu_i}) \left(\frac{\bar{h}V - U\bar{h}}{V^2} \right) \\ \frac{\partial I_i^*}{\partial \beta_i} &= \gamma_n \frac{z_n}{p_n} \left(-\frac{\bar{h}(\bar{H}_n - N_n\bar{h} + \bar{h} \sum_{i \in n} \beta_i) - \bar{h}^2 \sum_{i \in n} \beta_i}{(\bar{H}_n - N_n\bar{h} + \bar{h} \sum_{i \in n} \beta_i)^2} \right) \\ &\quad + (1 - \theta_n^{\mu_i}) \left(\frac{\bar{h}(\bar{H}_n - N_n\bar{h} + \bar{h} \sum_{i \in n} \beta_i) - \bar{h}^2 \sum_{i \in n} \beta_i}{(\bar{H}_n - N_n\bar{h} + \bar{h} \sum_{i \in n} \beta_i)^2} \right)\end{aligned}$$

This can then be simplified to become

$$\frac{\partial I_i^*}{\partial \beta_i} = \gamma_n \frac{z_n}{p_n} \left(-\frac{\bar{h}(\bar{H}_n - N_n\bar{h})}{(\bar{H}_n - N_n\bar{h} + \bar{h} \sum_{i \in n} \beta_i)^2} \right) + (1 - \theta_n^{\mu_i}) \left(\frac{\bar{h}(\bar{H}_n - N_n\bar{h})}{(\bar{H}_n - N_n\bar{h} + \bar{h} \sum_{i \in n} \beta_i)^2} \right)$$

Such that $\frac{\partial I_i}{\partial \beta_i} < 0$. I finally simplify the equation by grouping the common term which appears in parentheses to express the derivative as

$$\frac{\partial I_i^*}{\partial \beta_i} = \left((1 - \theta_n^{\mu_i}) - \gamma_n \frac{z_n}{p_n} \right) \left(\frac{\bar{h}(\bar{H}_n - N_n\bar{h})}{(\bar{H}_n - N_n\bar{h} + \bar{h} \sum_{i \in n} \beta_i)^2} \right)$$

4.5.9 Non-Movers' Demand for Investment Services

I now turn to deriving the demand for investment services for non-movers. Non-movers take decisions over the consumption of the final good and of investment services. Given that they do not take decisions over housing, I exclude this from optimisation, but include housing spend within the budget constraint. Setting up the maximisation problem as a Lagrangian using the utility function in Eq.(44) returns

$$\mathbb{L}^g = \left(\frac{C_i^{nm}}{\psi} \right)^\psi \left(\frac{(I_i^{nm})^{\Gamma_i}}{1 - \psi} \right)^{1-\psi} + \lambda (z_n - C_i^{nm} - r_n H_i - p_n I_i^{nm})$$

Taking the derivative of this with respect to the final good and investment services returns the first order conditions

$$\begin{aligned}
1) \quad \frac{\partial \mathbb{L}^g}{\partial C_i^{nm}} &= \frac{\psi \left(\frac{C_i^{nm}}{\psi} \right) \left(\frac{(I_i^{nm})^{\Gamma_i}}{1-\psi} \right)^{1-\psi}}{C_i^{nm}} - \lambda = 0 \\
2) \quad \frac{\partial \mathbb{L}^g}{\partial I_i^{nm}} &= \frac{\left(\frac{C_i^{nm}}{\psi} \right) \Gamma_i (I_i^{nm})^{\Gamma_i-1}}{\left(\frac{(I_i^{nm})^{\Gamma_i}}{1-\psi} \right)^\psi} - \lambda p_n = 0
\end{aligned}$$

One obtains the demand functions by dividing first order condition 1) by 2) and expressing the constraint in terms either good. Doing so returns the following in terms of I_i^{nm}

$$I_i^{nm} = \frac{C_i^{nm}(1-\psi)\Gamma_i}{\psi p_n}$$

Substituting into the budget constraint and simplifying means that the constraint becomes

$$\psi(z_n - r_n H_i) = C_i^{nm} (\psi + \Gamma_i - \psi \Gamma_i)$$

Which, when rearranged for C_i , becomes

$$C_i^{nm} = (z_n - r_n H_i) + \frac{\psi(z_n - r_n H_i)}{(1-\psi)\Gamma_i}$$

Substituting this back into the equation for I_i^{nm} – which is currently in terms of C_i – I obtain the demand for investment services

$$\begin{aligned}
I_i^{nm} &= \frac{\left((z_n - r_n H_i) + \frac{\psi(z_n - r_n H_i)}{(1-\psi)\Gamma_i} \right) (1-\psi)\Gamma_i}{\psi p_n} \\
I_i^{nm} &= \frac{(z_n - r_n H_i)(1-\psi)\Gamma_i + \psi(z_n - r_n H_i)}{\psi p_n} \\
I_i^{nm} &= \frac{(z_n - r_n H_i)(1-\psi)\Gamma_i}{\psi p_n} + \frac{(z_n - r_n H_i)}{p_n} \\
I_i^{nm} &= \frac{(1-\psi)}{\psi} \frac{\Gamma_i(z_n - r_n H_i)}{p_n} + \frac{(z_n - r_n H_i)}{p_n}
\end{aligned}$$

In previously characterising the model, I determined equilibrium values for p_n and r_n . I now substitute in these values and solve for the equilibrium consumption of investment services among non-movers. For the simplification of notation, I only show the substitution for the term $\frac{\Gamma_i(z_n - r_n H_i)}{p_n}$, given that the same substitution will hold true for the rest of the equation $\left(\frac{z_n - r_n H_i}{p_n}\right)$. Like before, I first substitute in for property prices for ease of presentation

$$I_i^{nm} = \frac{\Gamma_i \left(z_n - \left(\frac{z_n \sum_{i \in n} \beta_i - \frac{p_n}{\gamma_n} \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i})}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \right) H_i \right)}{p_n}$$

Again, for the ease of notation, let us make some substitutions such that $U = \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i})$ and $V = \bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i$. We can then rewrite the equation in a format that is easier to read

$$\begin{aligned} I_i^{nm} &= \frac{\Gamma_i}{p_n} \left(z_n - \left(\frac{z_n \sum_{i \in n} \beta_i - \frac{p_n}{\gamma_n} U}{V} \right) H_i \right) \\ I_i^{nm} &= \frac{\Gamma_i}{p_n} \left(\left(\frac{V z_n - H_i z_n \sum_{i \in n} \beta_i + H_i \frac{p_n}{\gamma_n} U}{V} \right) \right) \\ I_i^{nm} &= \frac{\Gamma_i}{V p_n} \left(V z_n - H_i z_n \sum_{i \in n} \beta_i + H_i \frac{p_n}{\gamma_n} U \right) \\ I_i^{nm} &= \frac{\Gamma_i z_n}{V p_n} \left(V - H_i \sum_{i \in n} \beta_i \right) + \frac{\Gamma_i}{V \gamma_n} (H_i U) \end{aligned}$$

Substituting back into the equation the values U and V , I get

$$\begin{aligned} I_i^{nm} &= \frac{\Gamma_i z_n}{(\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i) p_n} \left(\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i - H_i \sum_{i \in n} \beta_i \right) \\ &+ \frac{\Gamma_i H_i}{(\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i) \gamma_n} \left(\sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i}) \right) \end{aligned}$$

One can simplify this and add the ratio of preferences $\frac{(1-\psi)}{\psi}$ to become

$$I_i^{nm} = \frac{(1-\psi)}{\psi} \frac{\Gamma_i}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \left(\frac{z_n}{p_n} \left(\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i - H_i \sum_{i \in n} \beta_i \right) + \frac{H_i}{\gamma_n} \left(\sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i}) \right) \right)$$

In making the same substitution of property prices into the term $\left(\frac{z_n - r_n H_i}{p_n}\right)$ and undertaking the same steps, I obtain a similar solution, but without the ratio of preferences or the degree of exposure parameter

$$I_i^{nm} = \frac{1}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \left(\frac{z_n}{p_n} \left(\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i - H_i \sum_{i \in n} \beta_i \right) + \frac{H_i}{\gamma_n} \left(\sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i}) \right) \right)$$

Combining the two expressions returns the consumption of investment services by non-movers in equilibrium

$$\begin{aligned} I_i^{nm} = & \frac{(1 - \psi)}{\psi} \frac{\Gamma_i}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \left(\frac{z_n}{p_n} \left(\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i - H_i \sum_{i \in n} \beta_i \right) + \frac{H_i}{\gamma_n} \left(\sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i}) \right) \right) \\ & + \frac{1}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \left(\frac{z_n}{p_n} \left(\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i - H_i \sum_{i \in n} \beta_i \right) + \frac{H_i}{\gamma_n} \left(\sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i}) \right) \right) \end{aligned}$$

Which one can simplify to write in the following manner

$$\begin{aligned} I_i^{nm} = & \frac{(1 - \psi)}{\psi} \left(\Gamma_i + \frac{\psi}{(1 - \psi)} \right) \\ & \frac{1}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \left(\frac{z_n}{p_n} \left(\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i - H_i \sum_{i \in n} \beta_i \right) + \frac{H_i}{\gamma_n} \left(\sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i}) \right) \right) \end{aligned}$$

Finally, by substituting in the equilibrium price of investment services, I obtain the equation in its final form

$$\begin{aligned} I_i^{nm*} = & \frac{(1 - \psi)}{\psi} \left(\Gamma_i + \frac{\psi}{(1 - \psi)} \right) \frac{1}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \\ & \left(\frac{z_n^{1-\Omega} \lambda_n \Omega^\Omega (1 - \Omega)^{\frac{\eta(1-\Omega)}{1+\eta}} \bar{M}_n^{\frac{1-\Omega}{1+\eta}}}{C_n^{\frac{1-\Omega}{1+\eta}}} \left(\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i - H_i \sum_{i \in n} \beta_i \right) + \frac{H_i}{\gamma_n} \left(\sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i}) \right) \right) \end{aligned}$$

In the chapter I show that the neighbourhood effect – captured by exposure to the investments undertaken by movers – has a positive effect on the consumption of investment services by non-movers in equilibrium. I show that this neighbourhood effect is greater in neighbourhoods with greater preferences for investment services by taking the derivative of the exposure term Γ_i with respect to investment preferences γ_i from Eq.(49). I mea-

sure non-mover's distance-weighted exposure to investment services consumed by movers in accordance with the following equation

$$\Gamma_i = \sum_{j=1}^J e^{\delta\tau_{ij}} \nu_j$$

where ν_j captures the consumption of investment services, and the only differentiable term with respect to investment preferences, such that

$$\begin{aligned} \nu_j = \Gamma_i = \gamma_n \frac{z_n^{1-\Omega} \lambda_n \Omega^\Omega (1-\Omega)^{\frac{\eta(1-\Omega)}{1+\eta}} \bar{M}_n^{\frac{1-\Omega}{1+\eta}}}{C_n^{\frac{1-\Omega}{1+\eta}}} \left(1 - \frac{\bar{h} \sum_{i \in n} \beta_i}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \right) \\ + \frac{\bar{h} \sum_{i \in n} \beta_i (1 - \theta_n^{\mu_i})}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} + (1 - \theta_n^{\mu_i}) \left(\frac{1 - \gamma_n}{\gamma_n} \right) \end{aligned}$$

Taking the derivative of this then returns

$$\frac{\partial \Gamma_i}{\partial \gamma_i} = \frac{z_n^{1-\Omega} \lambda_n \Omega^\Omega (1-\Omega)^{\frac{\eta(1-\Omega)}{1+\eta}} \bar{M}_n^{\frac{1-\Omega}{1+\eta}}}{C_n^{\frac{1-\Omega}{1+\eta}}} \left(1 - \frac{\bar{h} \sum_{i \in n} \beta_i}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \right) - \frac{1 - \theta_n^{\mu_i}}{\gamma_n^2}$$

I then simplify inside the parentheses by setting the term 1 to have a common denominator

$$\frac{\partial \Gamma_i}{\partial \gamma_i} = \frac{z_n^{1-\Omega} \lambda_n \Omega^\Omega (1-\Omega)^{\frac{\eta(1-\Omega)}{1+\eta}} \bar{M}_n^{\frac{1-\Omega}{1+\eta}}}{C_n^{\frac{1-\Omega}{1+\eta}}} \left(\frac{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} - \frac{\bar{h} \sum_{i \in n} \beta_i}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \right) - \frac{1 - \theta_n^{\mu_i}}{\gamma_n^2}$$

Which, when simplified, returns the final equation

$$\frac{\partial \Gamma_i}{\partial \gamma_i} = \frac{z_n^{1-\Omega} \lambda_n \Omega^\Omega (1-\Omega)^{\frac{\eta(1-\Omega)}{1+\eta}} \bar{M}_n^{\frac{1-\Omega}{1+\eta}}}{C_n^{\frac{1-\Omega}{1+\eta}}} \left(\frac{\bar{H}_n - N_n \bar{h}}{\bar{H}_n - N_n \bar{h} + \bar{h} \sum_{i \in n} \beta_i} \right) - \frac{1 - \theta_n^{\mu_i}}{\gamma_n^2}$$

5 Conclusion

Using advanced econometric techniques and theoretical methods applied to data on energy efficiency characteristics and property transactions, this thesis contributes to the literature by exploring the spatial implications of residential energy efficiency improvements on urban property markets. The UK provides an interesting case study for empirical study given its notoriously old housing stock. Growing urbanisation worldwide and the push to achieve Net Zero targets makes this research not only interesting, but important as we strive to lower carbon emissions.

Chapter two explores the implications that residential energy efficiency investments have on urban property prices. The literature predominantly explores how prices respond to changes in current energy efficiency scores. Ex-ante, this appears a simplistic approach that neglects buyer's valuation of the future benefits attributed to such improvements in favour of those realised in the present. This chapter proposes a new definition that I term *unrealised efficiency*, which measures the distance between a property's current and potential score. One's distance to potential determines their prospects for investment and captures how buyers value moving closer to their potential (hence, their future gains).

Chapter three investigates the role of neighbourhood effects in driving location-specific property premia and propagating energy efficiency improvements. Leveraging energy efficiency data reported in property listing adverts, I explore the implications on property prices in locations with clusters of highly-efficient properties. Specifically, whether buyers pay a premium to locate in efficient property clusters and for what benefit. Evidence of a premium infers that we observe residential selection. I exploit this to test for the presence of neighbourhood effects in improvements owing to household's exposure to other selecting households, and to test the magnitude of these effects.

Chapter four then presents a theoretical model that rationalises the findings from the previous chapter. Residential selection creates a new spatial equilibrium within a city as residential moves induce spatial shifts in preferences. I develop a spatial, one-period

model of moving and non-moving households. For the former, I explore the implications for property prices and household location decisions of ex-ante demand for investment services (used to undertake home improvements) in the presence of neighbourhood effects. For the latter, I explore how exposure to residential selection decisions of movers affect their demand for investment services. The model incorporates the neighbourhood effect to specifically rationalise how the impacts deviate from that of the expected response found in chapter two.

5.1 Summary of the Findings

Chapter two estimates how increases in unrealised efficiency are capitalised into property prices, while the literature exclusively focuses on current efficiency. The two rates differ in what they measure. Current capitalisation estimates the value buyers place on the energy efficiency improvements undertaken by previous owners. Unrealised capitalisation estimates the value buyers place on the same improvements when they have not yet been undertaken. The latter explores household myopia by estimating the value placed on future energy efficiency gains. I first estimate the current capitalisation rate (for means of comparison) by following the literature with regards to the menu of control variables and fixed effects. Doing so returns a capitalisation rate of 1.3%, which falls in line with other prominent studies. Estimates for the unrealised capitalisation rate are much lower, around -0.15%, reflecting the discount buyers apply to undertake further improvements. Comparing the two shows that buyers place a lower value on undertaking their own improvements, explaining the limited uptake of retrofitting. Estimates from a set of heterogeneity analyses show that this discount is greater for more invasive (disruptive) improvements and when improvements are required within the property. Furthermore, buyers discount property prices more when the property's potential efficiency is more attainable. Uncertainty around the financial returns to home improvements drive this result and explain why households do not invest in retrofitting.

Chapter three explores the role of neighbourhood effects in energy efficiency improvements and the impact one's peers have on property prices. To the best of my knowledge, this is

the first study to explore the hyperlocal effect on property prices arising from residential energy efficiency improvements that are driven by neighbourhood effects. Peer groups are defined based on the proximity and similarity in energy efficiency scores by leveraging data from property listings from across the London TTWA. After identifying clusters of energy-efficient properties, results from a spatial RD show that buyers pay a 1.6% (£6,100) premium to locate in these property clusters over-and-above locations on the periphery. This reflects buyers' expectations that these efficient clusters deliver higher returns to investment and suggests that the households selecting into these properties have preferences for energy efficiency. Residential sorting of this kind is exogenous and introduces incumbent households to new preferences. Using this variation to instrument for new improvements, I show that buyers are actually paying for the improvements undertaken in these efficient clusters and not location in the clusters themselves. Changes in local preferences induce households to make further improvements through neighbourhood effects as they learn from those undertaken by selecting households.

Chapter four develops a unique theoretical model to explain these results. The model comprises two mutually exclusive types of household – movers and non-movers. Both take decisions over the level of housing investment to undertake, while the latter also takes decisions over housing. Buyers have Stone-Geary preferences for a numeraire good, housing and housing investment; the latter of which is determined relative to an investment incentive that varies according to their chosen neighbourhood's energy efficiency. This acts as the neighbourhood effect which restricts buyers to neighbourhoods that provide a suitable match for their preferences for investment. Housing demand and therefore equilibrium prices are shown to be a positive function of this neighbourhood effect, suggesting that buyers pay more to locate in efficient property clusters. As a result, buyers expect higher returns to their investments and pay higher prices to access these. Non-movers have Cobb-Douglas preferences for a numeraire good and investment services, the latter being a function of the investments undertaken nearby. Results from the optimisation show that these households demand greater levels of housing investment when they are located in clusters. Households invest according to their own preferences, but are subject

to a peer effect that induces further demand.

5.2 Policy Implications

This thesis provides a set of unique results informing policymakers of the impacts of residential energy efficiency investments on urban property markets. While there is positive evidence that energy efficiency improvements are being undertaken within urban environments, there is still a long way to go. The aim of the three chapters is to provide evidence that can be used to better understand the challenges facing mass retrofitting.

Chapter two presents evidence that energy efficiency improvements are capitalised into property prices, but that the incentive to undertake improvements is constrained by uncertainty with regards to the returns on investment. Previous attempts to encourage households to make energy efficiency improvements fell short. The Green Deal – funding residential improvement works through a loan payback scheme – only financed improvements in 10,000 households. Similar programmes targetting low-income households have been more successful, like the Energy Company Obligation (ECO) scheme which provided heating and insulation services to over 1.9m households between 2013 and 2018 (The Department of Business and Strategy, [2019](#)). Like the findings from this chapter, one can draw the conclusion that households are less likely to make improvements when they have to foot a substantial bill, as evidenced by the price discount on unrealised efficiency. This points to the idea of uncertainty in the returns to such improvements. Furthermore, findings from this chapter suggest that households are less willing to make improvements *within* their properties that require changes to their daily lifestyles. Moving forward, policy should be tailored towards providing support that addresses this shortcoming and that delivers greater energy efficiency benefits. Results from this chapter could provide strong evidence to support future policy in reducing the information barrier and providing further clarity to households about the returns to undertaking energy efficiency investments.

Chapters three and four are linked and provide evidence of both the benefits and drawbacks of residential energy efficiency investments. Results from the chapters provide

evidence of neighbourhood effects in improvements benefitting households exposed to those with preferences for energy efficiency. Though this is an advantage, the scope of these effect are likely to be contained within small geographic areas. While the aforementioned policies are tailored for different income groups, policy should also be designed and tailored for different groups geographically. The ECO scheme provides a taste of this by delivering funds to be appropriated by local districts. Poor uptake of national funding schemes suggests that only those with preferences for energy efficiency take up the offer, meaning funds devolved to districts should be distributed in a way that maximises spillovers on an intra-district basis. Evidence from this chapter would suggest that policy should target households in-and-around efficient property clusters to induce such spillovers. Another concern arising from these chapters is the impact on property prices in these areas. Non-negligible premia paid for properties within efficient clusters runs the risk of pricing households out of these markets and shifting the spatial equilibrium within cities. Increasing urbanisation rates accompanied by the necessity to retrofit en-masse make this an ever-present concern that needs to be addressed.

5.3 Limitations and Future Research

While the chapters contained within this thesis provide a starting point in exploring the spatial implications of residential energy efficiency improvements, they are by no means exhaustive. Below I discuss areas in which I believe this research can be taken forward.

Chapter one introduces the concept of unrealised efficiency – a previously undiscussed and unstudied measure. Heterogeneity tests highlight that unrealised capitalisation is far from uniform and differs with regards to the type of improvements required and with one’s investment potential. Heterogeneity is expected to arise across a number of other domains not discussed within the chapter. By country, as energy potential and culture vary geographically; by tenure as owner-occupiers and renters vary in their ability to undertake improvements, the value they assign to them, and the length of occupation from which to benefit from them; and by household characteristics, as capitalisation may vary by age, marriage status, or voting behaviour, to name but a few. The latter presents

interesting avenues for researchers with access to longitudinal data at the household level. A limitation of this chapter is that I do not have access to data on the specific improvements undertaken or their cost. An interesting line of research would be to understand the actual return on these investments inclusive of the direct energy savings. Research of this type would better be able to link to research on the energy efficiency gap (see Jaffe and Stavins, 1994, Gillingham and Palmer, 2014).

Chapters three and four explore the impact of neighbourhood effects on residential energy efficiency improvements in the context of endogenous neighbourhood selection. Neighbourhood effects in this context have received very little attention in the literature. Receiving less attention is how these neighbourhood effects affect property prices. A limitation of chapter three and the accompanying model in chapter four is that I am unable to explicitly model the interactions between peers. I instead define peer groups based on proximity and the matching of energy efficiency characteristics. Using a network analysis style approach, future research in this field – with access to better data – could explore how the conclusions of this paper are affected by: the size of one’s network, the specific characteristic of one’s peers, by changes in the node of centrality, and even by the density of one’s network. Applications of this from outside the field of energy efficiency come from Patacchini and Venanzoni, (2014) in studying the drivers of regular housing investment and Patacchini and Rice, (2007) in the context of regional growth. Such applications have the ability to dig deeper into the data and deliver important findings. It would also place further structure on the importance of social relationships in delivering energy efficiency gains across neighbourhoods.

References

- Alberini, Anna, Will Gans, and Daniel Velez-Lopez (2011). “Residential consumption of gas and electricity in the U.S.: The role of prices and income”. In: *Energy Economics* 33.5, pp. 870–881. ISSN: 0140-9883. DOI: <https://doi.org/10.1016/j.eneco.2011.01.015>. URL: <https://www.sciencedirect.com/science/article/pii/S0140988311000351>.
- Allcott, Hunt (2011a). “Rethinking real-time electricity pricing”. In: *Resource and energy economics* 33.4, pp. 820–842.
- Allcott, Hunt (2011b). “Social norms and energy conservation”. In: *Journal of public Economics* 95.9-10, pp. 1082–1095.
- Allcott, Hunt and Michael Greenstone (Feb. 2012). “Is There an Energy Efficiency Gap?” In: *Journal of Economic Perspectives* 26.1, pp. 3–28. DOI: [10.1257/jep.26.1.3](https://doi.org/10.1257/jep.26.1.3). URL: <https://www.aeaweb.org/articles?id=10.1257/jep.26.1.3>.
- Allcott, Hunt and Michael Greenstone (2017). *Measuring the welfare effects of residential energy efficiency programs*. Tech. rep. National Bureau of Economic Research.
- Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky (2012). *Externalities, internalities, and the targeting of energy policy*. National Bureau of Economic Research.
- Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky (2014). “Energy policy with externalities and internalities”. In: *Journal of Public Economics* 112, pp. 72–88.
- Allcott, Hunt and Todd Rogers (2014). “The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation”. In: *American Economic Review* 104.10, pp. 3003–3037.
- Allcott, Hunt and Dmitry Taubinsky (2013). *The lightbulb paradox: Evidence from two randomized experiments*. Tech. rep. National Bureau of Economic Research.
- Anenberg, Elliot (2016). “INFORMATION FRICTIONS AND HOUSING MARKET DYNAMICS”. In: *International Economic Review* 57.4, pp. 1449–1479. DOI: <https://doi.org/10.1111/iere.12204>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/iere.12204>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/iere.12204>.

-
- Ansar, Jasmin and Roger Sparks (2009). “The experience curve, option value, and the energy paradox”. In: *Energy Policy* 37.3, pp. 1012–1020.
- Arimura, Toshi H et al. (2012). “Cost-effectiveness of electricity energy efficiency programs”. In: *The Energy Journal* 33.2, pp. 63–100.
- Attari, Shahzeen Z et al. (2010). “Public perceptions of energy consumption and savings”. In: *Proceedings of the National Academy of sciences* 107.37, pp. 16054–16059.
- Autor, David H., Christopher J. Palmer, and Parag A. Pathak (2014). “Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts”. In: *Journal of Political Economy* 122.3, pp. 661–717. DOI: [10 . 1086 / 675536](https://doi.org/10.1086/675536). eprint: <https://doi.org/10.1086/675536>. URL: <https://doi.org/10.1086/675536>.
- Aydin, Erdal, Dirk Brounen, and Nils Kok (2020). “The capitalization of energy efficiency: Evidence from the housing market”. In: *journal of Urban Economics* 117, p. 103243.
- Ayres, Ian, Sophie Raseman, and Alice Shih (2013). “Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage”. In: *The Journal of Law, Economics, & Organization* 29.5, pp. 992–1022.
- Baldauf, Markus, Lorenzo Garlappi, and Constantine Yannelis (Feb. 2020). “Does Climate Change Affect Real Estate Prices? Only If You Believe In It”. In: *The Review of Financial Studies* 33.3, pp. 1256–1295. ISSN: 0893-9454. DOI: [10 . 1093 / rfs / hhz073](https://doi.org/10.1093/rfs/hhz073). eprint: <https://academic.oup.com/rfs/article-pdf/33/3/1256/32449257/hhz073.pdf>. URL: <https://doi.org/10.1093/rfs/hhz073>.
- Ball, Michael, Geoffrey Meen, and Christian Nygaard (2010). “Housing supply price elasticities revisited: Evidence from international, national, local and company data”. In: *Journal of Housing Economics* 19.4, pp. 255–268.
- Bayer, Patrick, Fernando Ferreira, and Robert McMillan (2007). “A unified framework for measuring preferences for schools and neighborhoods”. In: *Journal of political economy* 115.4, pp. 588–638.

-
- Bigler, Patrick and Benedikt Janzen (2023). “Green spills: Peer effects in pro-environmental behaviors”. In: *Available at SSRN 4608076*.
- Blouri, Yashar, Simon Büchler, and Olivier Schöni (2023). “The geography of mortgage interest deductions”. In: *Journal of Urban Economics* 138, p. 103604.
- Bollinger, Bryan, Jesse Burkhardt, and Kenneth T. Gillingham (Aug. 2020). “Peer Effects in Residential Water Conservation: Evidence from Migration”. In: *American Economic Journal: Economic Policy* 12.3, pp. 107–33. DOI: [10.1257/pol.20180559](https://doi.org/10.1257/pol.20180559). URL: <https://www.aeaweb.org/articles?id=10.1257/pol.20180559>.
- Bollinger, Bryan and Kenneth Gillingham (Apr. 2012). “Peer Effects in the Diffusion of Solar Photovoltaic Panels”. In: *Marketing Science* 31. DOI: [10.2307/23361417](https://doi.org/10.2307/23361417).
- Boomhower, Judson and Lucas Davis (2020). “Do energy efficiency investments deliver at the right time?” In: *American Economic Journal: Applied Economics* 12.1, pp. 115–139.
- Boomhower, Judson and Lucas W Davis (2014). “A credible approach for measuring inframarginal participation in energy efficiency programs”. In: *Journal of Public Economics* 113, pp. 67–79.
- Borenstein, Severin (2002). “The trouble with electricity markets: understanding California’s restructuring disaster”. In: *Journal of economic perspectives* 16.1, pp. 191–211.
- Bound, John, David A. Jaeger, and Regina M. Baker (1995). “Problems with Instrumental Variables Estimation When the Correlation Between the Instruments and the Endogenous Explanatory Variable is Weak”. In: *Journal of the American Statistical Association* 90.430, pp. 443–450. ISSN: 01621459. URL: <http://www.jstor.org/stable/2291055> (visited on 04/20/2022).
- Brandon, Alec et al. (2017). *Do the effects of nudges persist? Theory and evidence from 38 natural field experiments*. Tech. rep. National Bureau of Economic Research.
- Brounen, Dirk and Nils Kok (2011). “On the economics of energy labels in the housing market”. In: *Journal of Environmental Economics and Management* 62.2, pp. 166–179. ISSN: 0095-0696. DOI: <https://doi.org/10.1016/j.jeem.2010.>

-
- 11.006. URL: <https://www.sciencedirect.com/science/article/pii/S0095069611000337>.
- Cajias, Marcelo, Franz Fuerst, and Sven Bienert (2019). “Tearing down the information barrier: the price impacts of energy efficiency ratings for buildings in the German rental market”. In: *Energy Research & Social Science* 47, pp. 177–191. ISSN: 2214-6296. DOI: <https://doi.org/10.1016/j.erss.2018.08.014>. URL: <https://www.sciencedirect.com/science/article/pii/S2214629618309095>.
- Cajias, Marcelo and Daniel Piazzolo (Dec. 2012). “Green Performs Better: Energy Efficiency and Financial Return on Buildings”. In: *Journal of Corporate Real Estate* 15, pp. 53–72. DOI: [10.1108/JCRE-12-2012-0031](https://doi.org/10.1108/JCRE-12-2012-0031).
- Cajias, Marcelo and Daniel Piazzolo (2013). “Green performs better: energy efficiency and financial return on buildings”. In: *Journal of Corporate Real Estate* 15.1, pp. 53–72.
- Calabrese, Stephen et al. (2006). “Local public good provision: Voting, peer effects, and mobility”. In: *Journal of Public Economics* 90.6-7, pp. 959–981.
- Callaway, Brantly and Pedro H.C. Sant’Anna (2021). “Difference-in-Differences with multiple time periods”. In: *Journal of Econometrics* 225.2. Themed Issue: Treatment Effect 1, pp. 200–230. ISSN: 0304-4076. DOI: <https://doi.org/10.1016/j.jeconom.2020.12.001>. URL: <https://www.sciencedirect.com/science/article/pii/S0304407620303948>.
- Calonico, Sebastian, Matias D Cattaneo, and Rocio Titiunik (2014). “Robust nonparametric confidence intervals for regression-discontinuity designs”. In: *Econometrica* 82.6, pp. 2295–2326.
- Carroll, James, Eleanor Denny, and Seán Lyons (2016). “The effects of energy cost labelling on appliance purchasing decisions: trial results from Ireland”. In: *Journal of Consumer Policy* 39.1, pp. 23–40.
- Chi, Bin et al. (2021). “A new attribute-linked residential property price dataset for England and Wales, 2011 to 2019”. In: *UCL Open: Environment Preprint*.

-
- Clara, Nuno et al. (Jan. 2022). “Investments that Make our Homes Greener: The Role of Regulation”. In: *SSRN*. DOI: <http://dx.doi.org/10.2139/ssrn.4009054>.
- Comerford, David A., Ian Lange, and Mirko Moro (2018). “Proof of concept that requiring energy labels for dwellings can induce retrofitting”. In: *Energy Economics* 69, pp. 204–212. ISSN: 0140-9883. DOI: <https://doi.org/10.1016/j.eneco.2017.11.013>. URL: <https://www.sciencedirect.com/science/article/pii/S014098831730405X>.
- Costa, Dora L. and Matthew E. Kahn (2011). “Electricity Consumption and Durable Housing: Understanding Cohort Effects”. In: *American Economic Review* 101.3, pp. 88–92.
- Couture, Victor et al. (2019). *Income growth and the distributional effects of urban spatial sorting*. Tech. rep. National Bureau of Economic Research.
- Couture, Victor et al. (Apr. 2023). “Income Growth and the Distributional Effects of Urban Spatial Sorting”. In: *The Review of Economic Studies*, rdad048. ISSN: 0034-6527. DOI: [10.1093/restud/rdad048](https://doi.org/10.1093/restud/rdad048). eprint: <https://academic.oup.com/restud/advance-article-pdf/doi/10.1093/restud/rdad048/50437171/rdad048.pdf>. URL: <https://doi.org/10.1093/restud/rdad048>.
- Davis, Lucas W, Alan Fuchs, and Paul Gertler (2014). “Cash for coolers: evaluating a large-scale appliance replacement program in Mexico”. In: *American Economic Journal: Economic Policy* 6.4, pp. 207–238.
- Deng, Yongheng, Zhiliang Li, and John M. Quigley (2012). “Economic returns to energy-efficient investments in the housing market: Evidence from Singapore”. In: *Regional Science and Urban Economics* 42.3. Special Section on Asian Real Estate Market, pp. 506–515. ISSN: 0166-0462. DOI: <https://doi.org/10.1016/j.regsciurbeco.2011.04.004>. URL: <https://www.sciencedirect.com/science/article/pii/S016604621100055X>.
- Department for Levelling Up, Housing, Communities, and Communities & Local Government Ministry of Housing (2022). *Live tables on Energy Performance of Buildings Certificates*. URL: <https://www.gov.uk/government/statistical-data-sets/>

[live-tables-on-energy-performance-of-buildings-certificates#epcs-for-non-domestic-properties](#) (visited on 07/28/2022).

- Deutsch, Matthias (2010). “The effect of life-cycle cost disclosure on consumer behavior: evidence from a field experiment with cooling appliances”. In: *Energy Efficiency* 3.4, pp. 303–315.
- Diamond, Rebecca (2016). “The determinants and welfare implications of US workers’ diverging location choices by skill: 1980–2000”. In: *American Economic Review* 106.3, pp. 479–524.
- Díaz, Araceli Ortega and Kenneth B Medlock (2021). “Price elasticity of demand for fuels by income level in Mexican households”. In: *Energy Policy* 151, p. 112132.
- Dietz, Thomas et al. (2009). “Household actions can provide a behavioral wedge to rapidly reduce US carbon emissions”. In: *Proceedings of the national academy of sciences* 106.44, pp. 18452–18456.
- Dinan, Terry M. and John A. Miranowski (1989). “Estimating the implicit price of energy efficiency improvements in the residential housing market: A hedonic approach”. In: *Journal of Urban Economics* 25.1, pp. 52–67. ISSN: 0094-1190. DOI: [https://doi.org/10.1016/0094-1190\(89\)90043-0](https://doi.org/10.1016/0094-1190(89)90043-0). URL: <https://www.sciencedirect.com/science/article/pii/0094119089900430>.
- Dixit, Avinash (1991). “Irreversible investment with price ceilings”. In: *Journal of Political Economy* 99.3, pp. 541–557.
- Dubin, Jeffrey A. and Daniel L. McFadden (1984). “An Econometric Analysis of Residential Electric Appliance Holdings and Consumption”. In: *Econometrica* 52.2, pp. 345–362. ISSN: 00129682, 14680262. URL: <http://www.jstor.org/stable/1911493> (visited on 04/20/2022).
- Dubin, Jeffrey A., Allen K. Miedema, and Ram V. Chandran (1986). “Price Effects of Energy-Efficient Technologies: A Study of Residential Demand for Heating and Cooling”. In: *The RAND journal of Economics* 17.3, pp. 310–325. ISSN: DUBIN07416261. URL: <http://www.jstor.org/stable/2555713> (visited on 04/07/2022).

-
- Eeckhout, Jan, Roberto Pinheiro, and Kurt Schmidheiny (2014). “Spatial sorting”. In: *Journal of Political Economy* 122.3, pp. 554–620.
- Enrich, Jacint et al. (2024). “Measuring the impact of time-of-use pricing on electricity consumption: Evidence from Spain”. In: *Journal of Environmental Economics and Management* 123, p. 102901.
- Eurostat (2023). *Housing Horizons: Examining the UK Housing Stock in an International Context*. URL: https://www.hbf.co.uk/documents/12890/International_Audit_Digital_v1.pdf.
- Eurostat (2024). *Living conditions in Europe - energy efficiency in households*. URL: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Living_conditions_in_Europe%5C-%5C_energy_efficiency_in_households#Source_data_for_tables_and_graphs.
- Evangelista, Rui, João Andrade e Silva, and Esmeralda A Ramalho (2022). “How heterogeneous is the impact of energy efficiency on dwelling prices? Evidence from the application of the unconditional quantile hedonic model to the Portuguese residential market”. In: *Energy Economics* 109, p. 105955.
- Fabra, Natalia et al. (2021). “Estimating the elasticity to real-time pricing: evidence from the Spanish electricity market”. In: *AEA Papers and Proceedings*. Vol. 111. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 425–429.
- Faruqui, Ahmad, Sanem Sergici, and Lamine Akaba (2012). “Dynamic pricing in a moderate climate: the evidence from Connecticut”. In: *Available at SSRN 2028178*.
- Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram (2015). “Are the non-monetary costs of energy efficiency investments large? Understanding low take-up of a free energy efficiency program”. In: *American Economic Review* 105.5, pp. 201–04.
- Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram (2018). “Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Pro-

-
- gram”. In: *The Quarterly journal of Economics* 133.3, pp. 1597–1644. DOI: [10.1093/qje/qjy005](https://doi.org/10.1093/qje/qjy005). URL: <https://doi.org/10.1093/qje/qjy005>.
- Fuerst, Franz, Michel Ferreira Cardia Haddad, and Hassan Adan (2020). “Is there an economic case for energy-efficient dwellings in the UK private rental market?” In: *Journal of Cleaner Production* 245, p. 118642.
- Fuerst, Franz and Georgia Warren-Myers (2018). “Does voluntary disclosure create a green lemon problem? Energy-efficiency ratings and house prices”. In: *Energy Economics* 74, pp. 1–12.
- Fuerst, Franz et al. (2015). “Does energy efficiency matter to home-buyers? An investigation of EPC ratings and transaction prices in England”. In: *Energy Economics* 48, pp. 145–156.
- Fuerst, Franz et al. (2016). “Energy performance ratings and house prices in Wales: An empirical study”. In: *Energy Policy* 92, pp. 20–33.
- Gaigné, Carl et al. (2022). “Who lives where in the city? Amenities, commuting and income sorting”. In: *Journal of Urban Economics* 128, p. 103394.
- Gardner, Gerald T and Paul C Stern (2008). “The short list: The most effective actions US households can take to curb climate change”. In: *Environment: science and policy for sustainable development* 50.5, pp. 12–25.
- Ghosh, Arpita, Brendon McConnell, and Jaime Millán-Quijano (2024). “Do Homebuyers Value Energy Efficiency? Evidence From an Information Shock”. In.
- Gillingham, Kenneth and Karen Palmer (2014). “Bridging the Energy Efficiency Gap: Policy Insights from Economic Theory and Empirical Evidence”. In: *Review of Environmental Economics and Policy* 8.1, pp. 18–38. DOI: [10.1093/reep/ret021](https://doi.org/10.1093/reep/ret021). URL: <https://doi.org/10.1093/reep/ret021>.
- Glaeser, Edward L and Joseph Gyourko (2005). “Urban decline and durable housing”. In: *Journal of political economy* 113.2, pp. 345–375.
- González-Pampillón, Nicolás (2022). “Spillover effects from new housing supply”. In: *Regional Science and Urban Economics* 92, p. 103759. ISSN: 0166-0462. DOI: <https://doi.org/10.1016/j.regsci.2022.103759>.

-
- [//doi.org/10.1016/j.regsciurbeco.2021.103759](https://doi.org/10.1016/j.regsciurbeco.2021.103759). URL: <https://www.sciencedirect.com/science/article/pii/S0166046221001198>.
- Graziano, Marcello and Kenneth Gillingham (Oct. 2014). “Spatial patterns of solar photovoltaic system adoption: The influence of neighbors and the built environment ‡”. In: *Journal of Economic Geography* 15.4, pp. 815–839. ISSN: 1468-2702. DOI: [10.1093/jeg/lbu036](https://doi.org/10.1093/jeg/lbu036). eprint: <https://academic.oup.com/joeg/article-pdf/15/4/815/7370549/lbu036.pdf>. URL: <https://doi.org/10.1093/jeg/lbu036>.
- Grossmann, Volker et al. (2021). “Distributional consequences of surging housing rents”. In: *Journal of Economic Theory* 196, p. 105275.
- Guerrieri, Veronica, Daniel Hartley, and Erik Hurst (2013). “Endogenous gentrification and housing price dynamics”. In: *Journal of Public Economics* 100, pp. 45–60. ISSN: 0047-2727. DOI: <https://doi.org/10.1016/j.jpubeco.2013.02.001>. URL: <https://www.sciencedirect.com/science/article/pii/S0047272713000297>.
- Harold, J, J Cullinan, and Sean Lyons (2017). “The income elasticity of household energy demand: a quantile regression analysis”. In: *Applied Economics* 49.54, pp. 5570–5578.
- Hausman, Jerry A. (1979). “Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables”. In: *The Bell journal of Economics* 10.1, pp. 33–54.
- Heblich, Stephan, Alex Trew, and Yanos Zylberberg (2021). “East-Side Story: Historical Pollution and Persistent Neighborhood Sorting”. In: *Journal of Political Economy* 129.5, pp. 1508–1552. DOI: [10.1086/713101](https://doi.org/10.1086/713101). eprint: <https://doi.org/10.1086/713101>. URL: <https://doi.org/10.1086/713101>.
- Heutel, Garth (2011). *Optimal policy instruments for externality-producing durable goods under time inconsistency*. Tech. rep. National Bureau of Economic Research.
- Hilber, Christian A.L., Charles Palmer, and Edward W. Pinchbeck (2019). “The energy costs of historic preservation”. In: *journal of Urban Economics* 114, p. 103197. DOI: <https://doi.org/10.1016/j.jue.2019.103197>. URL: <https://www.sciencedirect.com/science/article/pii/S0094119019300749>.

-
- Hirst, Eric (1986). “Actual energy savings after retrofit: Electrically heated homes in the Pacific Northwest”. In: *Energy* 11.3, pp. 299–308.
- Houde, Sébastien and Joseph E Aldy (2017). “Consumers’ response to state energy efficient appliance rebate programs”. In: *American Economic Journal: Economic Policy* 9.4, pp. 227–255.
- Hyland, Marie, Ronan C Lyons, and Seán Lyons (2013). “The value of domestic building energy efficiency—evidence from Ireland”. In: *Energy economics* 40, pp. 943–952.
- Ito, Koichiro (2015). “Asymmetric incentives in subsidies: Evidence from a large-scale electricity rebate program”. In: *American Economic Journal: Economic Policy* 7.3, pp. 209–237.
- Ito, Koichiro, Takanori Ida, and Makoto Tanaka (2018). “Moral suasion and economic incentives: Field experimental evidence from energy demand”. In: *American Economic Journal: Economic Policy* 10.1, pp. 240–267.
- Jaffe, Adam B. and Robert N. Stavins (1994). “The energy-efficiency gap What does it mean?” In: *Energy Policy* 22.10. Markets for energy efficiency, pp. 804–810. ISSN: 0301-4215. DOI: [https://doi.org/10.1016/0301-4215\(94\)90138-4](https://doi.org/10.1016/0301-4215(94)90138-4). URL: <https://www.sciencedirect.com/science/article/pii/0301421594901384>.
- Jessoe, Katrina and David Rapson (2014). “Knowledge is (less) power: Experimental evidence from residential energy use”. In: *American Economic Review* 104.4, pp. 1417–1438.
- Johnson, Ruth C. and David L. Kaserman (1983). “Housing market capitalisation of energy-saving durable good investments”. In: *Economic Inquiry* 21.3, pp. 374–386. DOI: <https://doi.org/10.1111/j.1465-7295.1983.tb00639.x>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1465-7295.1983.tb00639.x>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1465-7295.1983.tb00639.x>.
- Joskow, Paul L and Donald B Marron (1992). “What does a negawatt really cost? Evidence from utility conservation programs”. In: *The Energy Journal* 13.4.

-
- Kahn, Matthew E. and Nils Kok (2014). “The capitalization of green labels in the California housing market”. In: *Regional Science and Urban Economics* 47. SI: Tribute to John Quigley, pp. 25–34. ISSN: 0166-0462. DOI: <https://doi.org/10.1016/j.regsciurbeco.2013.07.001>. URL: <https://www.sciencedirect.com/science/article/pii/S0166046213000574>.
- Kallbekken, Steffen, Håkon Sælen, and Erlend AT Hermansen (2013). “Bridging the energy efficiency gap: A field experiment on lifetime energy costs and household appliances”. In: *Journal of Consumer Policy* 36.1, pp. 1–16.
- Kholodilin, Konstantin A, Andreas Mense, and Claus Michelsen (2017). “The market value of energy efficiency in buildings and the mode of tenure”. In: *Urban Studies* 54.14, pp. 3218–3238.
- Koster, Hans R. A., Jos N. van Ommeren, and Piet Rietveld (Nov. 2014). “Historic amenities, income and sorting of households”. In: *Journal of Economic Geography* 16.1, pp. 203–236. ISSN: 1468-2702. DOI: [10.1093/jeg/lbu047](https://doi.org/10.1093/jeg/lbu047). eprint: <https://academic.oup.com/joeg/article-pdf/16/1/203/7196499/lbu047.pdf>. URL: <https://doi.org/10.1093/jeg/lbu047>.
- Lacetera, Nicola, Devin G Pope, and Justin R Sydnor (May 2012). “Heuristic Thinking and Limited Attention in the Car Market”. In: *American Economic Review* 102.5, pp. 2206–2236. DOI: [10.1257/aer.102.5.2206](https://doi.org/10.1257/aer.102.5.2206). URL: <https://www.aeaweb.org/articles?id=10.1257/aer.102.5.2206>.
- Longstreth, Molly, Anne R. Coveney, and Jean S. Bowers (June 1984). “Conservation Characteristics Among Determinants of Residential Property Value”. In: *Journal of Consumer Research* 11.1, pp. 564–571. ISSN: 0093-5301. DOI: [10.1086/208992](https://doi.org/10.1086/208992). eprint: <https://academic.oup.com/jcr/article-pdf/11/1/564/5159357/11-1-564.pdf>. URL: <https://doi.org/10.1086/208992>.
- Manski, Charles F. (1993). “Identification of Endogenous Social Effects: The Reflection Problem”. In: *The Review of Economic Studies* 60.3, pp. 531–542. ISSN: 00346527, 1467937X. URL: <http://www.jstor.org/stable/2298123> (visited on 06/23/2023).

-
- McCord, Michael et al. (2020). “Energy performance certificates and house prices: a quantile regression approach”. In: *Journal of European Real Estate Research* 13.3, pp. 409–434. DOI: <https://doi.org/10.1108/JERER-06-2020-0033>. URL: <https://www.emerald.com/insight/content/doi/10.1108/JERER-06-2020-0033/full/html>.
- McFadden, Daniel, Carlos Puig, and Daniel Kirschner (1977). “Determinants of the long-run demand for electricity”. In: *Proceedings of the American Statistical Association*. Vol. 1. 1. Business and Economics Section, pp. 109–19.
- Metcalf, Gilbert E and Kevin A Hassett (1999). “Measuring the energy savings from home improvement investments: evidence from monthly billing data”. In: *Review of economics and statistics* 81.3, pp. 516–528.
- Ministry of Housing, Communities and Local Government (2024). “English Housing Survey 2023 to 2024”. In: *Gov.UK*.
- Monte, Ferdinando, Stephen J Redding, and Esteban Rossi-Hansberg (2018). “Commuting, migration, and local employment elasticities”. In: *American Economic Review* 108.12, pp. 3855–3890.
- Mundaca, Luis and Margaret Samahita (2020). “What drives home solar PV uptake? Subsidies, peer effects and visibility in Sweden”. In: *Energy Research & Social Science* 60, p. 101319.
- Nauze, Andrea La (2023). “Motivation crowding in peer effects: The effect of solar subsidies on green power purchases”. In: *Review of Economics and Statistics* 105.6, pp. 1465–1480.
- Nevin, Rick and Gregory Watson (1998). *Evidence of rational market valuations for home energy efficiency*. MPRA Paper. University Library of Munich, Germany. URL: <https://EconPapers.repec.org/RePEc:pra:mprapa:35343>.
- Newell, Richard G and Juha Siikamäki (2014). “Nudging energy efficiency behavior: The role of information labels”. In: *Journal of the Association of Environmental and Resource Economists* 1.4, pp. 555–598.

-
- Ölander, Folke and John Thøgersen (2014). “Informing versus nudging in environmental policy”. In: *Journal of Consumer Policy* 37.3, pp. 341–356.
- Ord, J. K. and Arthur Getis (1995). “Local Spatial Autocorrelation Statistics: Distributional Issues and an Application”. In: *Geographical Analysis* 27.4, pp. 286–306. DOI: <https://doi.org/10.1111/j.1538-4632.1995.tb00912.x>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1538-4632.1995.tb00912.x>.
- Parti, Michael and Cynthia Parti (1980). “The total and appliance-specific conditional demand for electricity in the household sector”. In: *The Bell journal of economics*, pp. 309–321.
- Patacchini, Eleonora and Patricia Rice (2007). “Geography and economic performance: exploratory spatial data analysis for Great Britain”. In: *Regional Studies* 41.4, pp. 489–508.
- Patacchini, Eleonora and Giuseppe Venanzoni (2014). “Peer effects in the demand for housing quality”. In: *Journal of Urban Economics* 83, pp. 6–17.
- Pope, Devin G (2009). “Reacting to rankings: Evidence from America’s Best Hospitals”. In: *Journal of Health Economics* 28.6, pp. 1154–1165. DOI: <https://doi.org/10.1016/j.jhealeco.2009.08.006>. URL: <https://www.sciencedirect.com/science/article/pii/S0167629609000873>.
- Qiu, Yueming, Gregory Colson, and Carola Grebitus (2014). “Risk preferences and purchase of energy-efficient technologies in the residential sector”. In: *Ecological Economics* 107, pp. 216–229.
- Quentel, Milan (2024). “Gone with the Wind: Renewable Energy Infrastructure, Welfare, and Redistribution”. In: *cerge-ei.cz*.
- Reis, Ricardo (2006). “Inattentive consumers”. In: *Journal of monetary Economics* 53.8, pp. 1761–1800.
- RICS (2022). “RICS Appraisal and Valuation Standards”. In: *Royal Institution of Chartered Surveyors*.

-
- Roback, Jennifer (1982). “Wages, Rents, and the Quality of Life”. In: *Journal of Political Economy* 90.6, pp. 1257–1278. ISSN: 00223808, 1537534X. URL: <http://www.jstor.org/stable/1830947> (visited on 08/18/2022).
- Rosen, Sherwin (1974). “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition”. In: *Journal of Political Economy* 82.1, pp. 34–55. ISSN: 00223808, 1537534X. URL: <http://www.jstor.org/stable/1830899> (visited on 08/18/2022).
- Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Raymond Owens III (2010). “Housing externalities”. In: *Journal of political Economy* 118.3, pp. 485–535.
- Saiz, Albert (2010). “The geographic determinants of housing supply”. In: *The Quarterly Journal of Economics* 125.3, pp. 1253–1296.
- Sallee, James M (2014). “Rational inattention and energy efficiency”. In: *The Journal of Law and Economics* 57.3, pp. 781–820.
- Samuelson, William and Richard Zeckhauser (1988). “Status quo bias in decision making”. In: *Journal of risk and uncertainty* 1.1, pp. 7–59.
- Schmidheiny, Kurt (2006). “Income segregation from local income taxation when households differ in both preferences and incomes”. In: *Regional Science and Urban Economics* 36.2, pp. 270–299.
- Schultz, Wesley P, Azar M Khazian, and Adam C Zaleski (2008). “Using normative social influence to promote conservation among hotel guests”. In: *Social influence* 3.1, pp. 4–23.
- Sejas-Portillo, Rodolfo et al. (2020). *Limited Attention in the Housing Market: Threshold Effects of Energy-Performance Certificates on Property Prices and Energy-Efficiency Investments*. CESifo Working Paper. URL: <http://hdl.handle.net/10419/229487>.
- Severen, Christopher (2023). “Commuting, labor, and housing market effects of mass transportation: Welfare and identification”. In: *Review of Economics and Statistics* 105.5, pp. 1073–1091.

-
- Sims, Christopher A (2003). “Implications of rational inattention”. In: *Journal of monetary Economics* 50.3, pp. 665–690.
- Sokołowski, Jakub (2023). “Peer effects on photovoltaics (PV) adoption and air quality spillovers in Poland”. In: *Energy Economics* 125, p. 106808. ISSN: 0140-9883. DOI: <https://doi.org/10.1016/j.eneco.2023.106808>. URL: <https://www.sciencedirect.com/science/article/pii/S0140988323003067>.
- Sterman, John D and Linda Booth Sweeney (2007). “Understanding public complacency about climate change: Adults’ mental models of climate change violate conservation of matter”. In: *Climatic change* 80.3, pp. 213–238.
- Stigler, George J (1961). “The Economics of Information”. In: *Journal of Political Economy* 69.3, pp. 213–225. URL: <http://www.jstor.org/stable/1829263>.
- Szumilo, Nikodem (2020). “Prices of Peers: Identifying Endogenous Price Effects in the Housing Market”. In: *The Economic journal* 131.639, pp. 3041–3070.
- Taylor, Lester D (1975). “The demand for electricity: a survey”. In: *The Bell Journal of Economics*, pp. 74–110.
- The Department of Business, Energy and Industrial Strategy (2019). “Energy Company Obligation 2019-2022: Policy Guidance for obligated suppliers, manufacturers and installers on applying for Demonstration Actions, Innovation Score Uplifts and In-situ Performance”. In: *Gov.UK*.
- The Department of Business, Energy and Industrial Strategy (2020). *Energy Performance Certificates for Buildings: Action Plan*. URL: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment%5C_data/file/922660/EPC%5C_Action%5C_Plan.pdf (visited on 05/11/2022).
- The Department of Business, Energy and Industrial Strategy (2025). *2023 UK Greenhouse Gas Emissions, Final Figures*. URL: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment%5C_data/file/1051408/2020-final-greenhouse-gas-emissions-statistical-release.pdf.

-
- UK Parliament (2013). *Warm Front Scheme*. URL: <https://researchbriefings.files.parliament.uk/documents/SN06231/SN06231.pdf>.
- Urban, Jan and Milan Ščasný (2012). “Exploring domestic energy-saving: The role of environmental concern and background variables”. In: *Energy policy* 47, pp. 69–80.
- van Duijn, Mark and Jan Rouwendal (2021). “Sorting based on urban heritage and income: Evidence from the Amsterdam metropolitan area”. In: *Regional Science and Urban Economics* 90, p. 103719. ISSN: 0166-0462. DOI: <https://doi.org/10.1016/j.regsciurbeco.2021.103719>. URL: <https://www.sciencedirect.com/science/article/pii/S016604622100079X>.
- Wolske, Kimberly S, Kenneth T Gillingham, and P Wesley Schultz (2020). “Peer influence on household energy behaviours”. In: *Nature Energy* 5.3, pp. 202–212.
- Zalejska-Jonsson, Agnieszka (2014). “Stated WTP and rational WTP: Willingness to pay for green apartments in Sweden”. In: *Sustainable Cities and Society* 13, pp. 46–56. ISSN: 2210-6707. DOI: <https://doi.org/10.1016/j.scs.2014.04.007>. URL: <https://www.sciencedirect.com/science/article/pii/S2210670714000432>.