Housing Energy Efficiency over the Life Cycle

Ran Gu Eglė Jakučionytė Swapnil Singh*

October 2026

Abstract

We document that older homeowners in England live in less energy-efficient homes. Two channels could explain this. Through the *horizon channel*, older owners have fewer years to collect the return on a renovation, so they invest less. Through the *purchase window*, efficiency is bought with the home. As mobility declines with age, older owners have fewer chances to acquire efficiency cheaply. We estimate a lifecycle model of housing, mobility, and renovation on English data. The age gradient comes from the purchase window, not the horizon channel. Efficiency wears off within a few years, so the payoff horizon hardly varies with age, while older owners buy less often. Hence renovation subsidies narrow the age gap, whereas point-of-sale standards widen it.

JEL classifications: D15, J14, Q41, Q48, R21, R31.
Keywords: energy efficiency, housing, renovation, residential mobility, lifecycle, population ageing.

*Gu: City St George's, University of London and Institute for Fiscal Studies, ran.gu@citystgeorges.ac.uk. Jakučionytė: Bank of Lithuania and Vilnius University, ejakucionyte@lb.lt. Singh: Bank of Lithuania and Kaunas University of Technology, ssingh@lb.lt. We thank Richard Blundell, James Carroll, Cristina Cella, Peter Reusens, and Orkun Saka for comments on an earlier draft, and Irma Spudienė and Marius Mašalas for help with computing resources. We also thank participants at the ECB Research Cluster and the Royal Economic Society Annual Conference. Pangram classifies 89 percent of the main text as human written. The appendices were written using AI (Opus 5.5 and Sol 6.1). The views expressed are those of the authors and do not necessarily represent the views of the Bank of Lithuania or its Board.

1. Introduction

One of the basic predictions of the theory of investment is that people invest less when fewer years remain to collect the return (Becker, 1962; Ben-Porath, 1967). For instance, a longer expected life raises investment in schooling (Jayachandran and Lleras-Muney, 2009; Oster, Shoulson, and Dorsey, 2013; Fortson, 2011), and workers who expect to retire later train more (Montizaan, Cörvers, and De Grip, 2010).[1] Schooling and training stay with the worker, who can neither buy them ready-made nor sell them on. However, when we look at durable goods like housing, which change hands over time, investments are embodied in the good and pass from seller to buyer. The energy efficiency of a home is one such investment. Efficiency improvements lower the energy bill for years and stay with the house when its owner sells, so a buyer can acquire them with the home. At the same time, mortality risk rises over the lifecycle, which by the standard argument shortens the horizon (although older owners, who move less, stay longer in their homes). In this paper we study which of these forces – a horizon that shortens with age and efficiency that is bought and sold with the house – shapes the energy efficiency of homes over the lifecycle of households. In the rest of the paper, for brevity, we call the first force the horizon channel and the second the purchase window, or simply the window.

We first use survey and register data in England to document three stylized facts. First, older owners live in less efficient homes. In the English Housing Survey, among owner-occupied homes built by 1990, 21 percent of those owned by households aged 25–34 are in Energy Performance Certificate (EPC, henceforth) bands A–C, against 9.7 percent of those owned by households aged 65–75, representing a gap of about 11 percentage points. Second, following a home across successive certificates in the EPC register, we show that efficiency is not fixed: each year about 7 percent of homes built by 1990 in bands D–G move to A–C, and about 15 percent of those in A–C move down. Third, efficiency is tied to when the owner bought. Owners who bought recently live in about equally efficient homes at every age.

What do these three stylized facts say about the two forces discussed above? The answer depends on how an owner gets an efficient home. An owner can make her home efficient in two ways: (i) buy a home that is already efficient, or (ii) renovate the one she has. The difference between the two lies in what a later sale returns. In the first case, a buyer pays for efficiency in the price and gets it back when she sells, if the home is still efficient. In the second case, a renovator pays for the work, and a sale returns only part of it.[2] The horizon channel acts on the owner's renovation decision. An older owner has fewer years to recover her investment, so she might renovate less. On the other hand, the purchase window affects whether an owner can get efficiency through a purchase. Because a buyer gets back at resale most of what she paid for efficiency, buying an efficient home is the cheaper route. However, it is open only when an owner moves, and older owners in Britain move less often (Banks, Blundell, Oldfield, and Smith, 2012). Between moves, efficiency wears off, as the second fact

[1] The prediction also holds outside human capital: executives in their final years in office cut R&D (Dechow and Sloan, 1991). There is a debate about which horizon matters: Americans worked fewer lifetime hours as they lived longer, so the remaining lifetime cannot be the whole story (Hazan, 2009; Cervellati and Sunde, 2013).

[2] Prices capitalise the value of energy savings (Myers, 2019; Aydin, Brounen, and Kok, 2020; Fuerst, McAllister, Nanda, and Wyatt, 2015), and retrofits typically cost more than the savings they deliver (Fowlie, Greenstone, and Wolfram, 2018).

shows. An owner who stays therefore lives in a home that becomes less efficient each year, and older owners have usually stayed the longest since their last purchase. So, both the horizon channel and the purchase window can produce the first fact. The third fact favours the window: among owners who bought recently, older and younger owners live in about equally efficient homes, so the gap opens with the time since purchase rather than with age itself. It does not settle the question, because under the horizon channel, too, an owner who stops renovating sees her home's efficiency fall with the time since purchase.

Why does it matter whether the declining age gradient comes from the horizon channel or the purchase window? The two channels have different implications for an ageing population. If the horizon channel determined the gradient, then an increase in life expectancy would raise investment in efficiency. However, if the purchase window determined it, longer lives would do little to raise investment: what matters is how often owners buy, and older owners buy less often. This means, from a policy perspective, knowledge of the underlying mechanism is crucial to determine whether to raise efficiency by incentivizing renovation or by acting at the point of sale.

To separate the implication of the horizon channel and the purchase window, we need two quantities that are not fully observable: (i) how fast efficiency wears off relative to how fast mortality rises, and (ii) how much cheaper efficiency is at a purchase than through a renovation.[3] To obtain them, we build a lifecycle model of households with income risk, mortality risk, and borrowing limits, built on the standard ingredients of lifecycle models of housing (Cocco, 2005; Yang, 2009; Kaplan, Mitman, and Violante, 2020). Each year a household chooses consumption and liquid saving, whether to rent or own, the size of her home, and her mortgage, and, if she owns, whether her home is efficient or not. An owner keeps her home, renovates it in place, or sells; a seller or a renter rents, buys an efficient home, buys an inefficient one, or buys an inefficient one and renovates it at once. Efficiency lowers the energy bill, gives comfort, and sells at a price premium, and each year it is lost with a fixed probability. Renovating has a fixed and a variable cost, buying and selling have transaction costs, and a bequest motive values what heirs receive, including the premium on an efficient home. The model holds tastes, information, and the cost of the work fixed across ages.[4] Any age gradient in the model must therefore come from the horizon channel or from the purchase window, which opens when an owner buys and so depends on transaction costs, borrowing limits, bequests, and income.

The model has many discrete and continuous choices and occasionally binding borrowing constraints, so we proceed in two steps. In the first, we simplify the model to isolate the key mechanisms analytically. This gives three results. (i) An owner is less willing to renovate the home she lives in the more likely she is to leave it, whether by moving or by dying, because a sale returns only part of the cost of the work. This is the horizon channel. (ii) A buyer's willingness to choose

[3]The register gives the rate at which homes lose a band, but only for homes that are certified again, and a change in band is not always a physical change. The register matched to Land Registry prices gives the price premium and the proportional cost of a renovation, but not the fixed cost of the work. The model we build below estimates both the rate at which efficiency is lost and the fixed cost (Section 4.3).

[4]Other explanations could also produce the facts: older owners may value efficiency less, know less about it, find the work more burdensome, or prefer to run down their home equity. These are preference-based explanations and we do not rule them out. However, we ask whether the dynamics of the investment alone can account for the gradient.

an efficient home depends little on how long she expects to stay, because she, or her heirs, recover most of the premium when the home is sold, provided it is still efficient. (iii) Renovating is therefore worthwhile for fewer owners than buying an efficient home is, because it costs more once what a sale returns is counted. Further, this difference grows with the chance of leaving. Some owners would not renovate the home they live in, yet would choose an efficient home if they were buying; whether they live in one depends on whether they move. This is the purchase window. These results give the direction of each channel. How much each contributes to the age gradient is the second step, for which we solve and estimate the full model.

Solving and estimating this model is cumbersome. An owner with a mortgage has five continuous states — liquid wealth, the mortgage balance, two components of income, and the size of the home — two home types, and up to 25 remaining mortgage terms, and she chooses among up to six alternatives with two or three continuous choices in each. This means a grid-based solution method cannot be employed.[5] We instead approximate the value and policy functions at each age with neural networks, trained backward from the oldest age on randomly drawn states (Maliar, Maliar, and Winant, 2021; Azinovic, Gaegauf, and Scheidegger, 2022; Druedahl and Røpke, 2026). Two devices make this work. Taste shocks over the alternatives smooth the kinks that discrete choices put in the value function, so that a network can fit it (Iskhakov, Jørgensen, Rust, and Schjerning, 2017). And the parameters we estimate are inputs of the networks, so one training covers the whole estimation search (Druedahl, Huleux, and Røpke, 2026).

We estimate seven parameters by the simulated method of moments, targeting eleven moments on ownership, wealth, consumption, and the share of efficient homes by age. The estimates deliver the two quantities that the facts alone could not. First, renovating a home costs between 2.7 and 4.2 times the extra outlay of buying the same home efficient, depending on its size.[6] Second, efficiency wears off fast: an efficient home loses its efficiency with a probability of 16 percent a year, close to the 15 percent of efficient homes that drop to bands D–G each year in the register, a rate we do not target.[7] The estimated model reproduces 10 percentage points of the 11 percentage point fall in the efficient share between ages 25–34 and 65–75 (Figure 4).

Using the estimated model we first untangle whether the age gradient comes from the purchase window or from the horizon channel. The horizon channel turns out to be weak because the horizon itself hardly differs between young and old owners. What matters for a renovation is not how long the owner will live, but how long she will benefit from the work. If efficiency lasted forever, an owner aged 25 could expect to benefit for 27 discounted years and an owner aged 75 for 10. But efficiency wears off long before most owners die, so the benefit lasts about five years at 25 and about four at 75. However, the purchase window differs sharply with age. Owners of every age make similar choices

[5] A grid with 50 points per continuous state would have up to 1.6×10^{10} points at each age.

[6] The fixed cost of renovating is identified by how steeply the efficient share falls with age, because older owners rely more on renovating in place (Section 4.3).

[7] The targeted moments pin down only a combination of how long efficiency lasts and how much it is worth each year, because both raise the efficient share at every age; the register's rate separates them. Part of that rate reflects reassessment rather than physical change (Crawley, Biddulph, Northrop, Wingfield, Oreszczyn, and Elwell, 2019), so we also report the main results when efficiency lasts 25 years instead of six.

when they act. A little over two in five buyers choose an efficient home, and fewer than two in a hundred owners of inefficient homes renovate in a year. What falls with age is how often owners buy, and this fall accounts for almost all of the decline with age in the number of owners who gain an efficient home. Older owners do not renovate instead, because of the cost of the work, which a buyer of an efficient home avoids. When we remove the option to buy an efficient home, the gap between the youngest and the oldest owners falls from about ten percentage points to about one.[8]

The United Kingdom aims to bring as many homes as possible to band C by 2035 (HM Government, 2017). Households headed by someone aged 65 or over own more than a third of owner-occupied homes in England (Ministry of Housing, Communities and Local Government, 2024), and their number is projected to grow by about a fifth in the ten years to 2032, twice as fast as households overall (Office for National Statistics, 2025). We use the estimated model for a series of counterfactuals, focusing on policies which governments generally apply to improve the efficiency of homes. Because the gradient comes from the window, the route through which a policy acts decides whom it reaches. We compare steady states in partial equilibrium and report each policy as a change from the model without it. A subsidy that halves the variable cost of renovating raises the efficient share by about 8 points among owners aged 65–74 but by less than 3 points at 25–34, so it narrows the gap between them by more than 5 points, at a cost of about £290 per household a year. It reaches older owners because they are the ones who rely on renovating in place. A point-of-sale standard, which requires every home that is bought to be efficient, raises the share the most, by about 13 points among all owners, but by more among the young (about 16 points) than the old (about 10), so it widens the gap by about 6 points. It also slows turnover. A household that would have bought an inefficient home now pays more to buy, and some put off moving: owners move and buy about a quarter less often. Each move that does not happen is a window that does not open, so the standard gives back part of its own gain. Lowering the costs of buying and selling works in the opposite direction. When we halve them, owners buy more than twice as often, more windows open, and the efficient share rises by about 5 points among all owners. But owners who expect to move sooner renovate less: the rate of renovation among owners who stay falls by about a third.

The counterfactual policies we discussed above change what it costs to get an efficient home. A policy can instead change what an efficient home pays its owner: a lower energy bill each year, and a higher price when the home is sold. We vary each in turn. Doubling energy prices, as a tax on household energy would, raises the efficient share in about the same proportion among young and old owners, by about 40 percent. It lifts the whole age profile without changing its shape, but it lowers homeownership by about 3 percentage points. The price premium works differently, and against a natural conjecture. If older owners renovate little because the market pays little for efficiency, a larger premium should bring them into efficient homes. In the model, doubling the premium does raise renovation, most among older owners, because a sale then returns more of the cost of the work. But a buyer pays the premium, so buying an efficient home becomes more expensive and fewer buyers choose one. Young owners, who get most of their efficiency by buying, lose the most. The age gap

[8]The result survives a much slower loss of efficiency. If efficiency lasts 25 years rather than six, the horizon now plays a role, but purchases still account for about 90 percent of the decline.

almost closes, but because the homes of young owners become less efficient, not because those of older owners improve.

The mechanism is not specific to energy. Much of the quality of a durable is chosen when it is bought, because that is when changing it is cheap. Models of durable adjustment explain when households replace what they own (Grossman and Laroque, 1990; Eberly, 1994; Attanasio, 2000). Our results add that the quality a household holds then depends on how often it transacts, so a household that rarely transacts holds quality chosen long ago. The fuel economy of the car fleet improves as old cars are scrapped and new ones bought, so a regulation that raises the price of new cars delays the improvement (Gruenspecht, 1982; Stavins, 2006; Jacobsen and van Benthem, 2015), and when owners replace or resell a car is itself an economic choice (Adda and Cooper, 2000; Stolyarov, 2002). Homes differ from cars in two ways. First, a home is rarely replaced but often resold. Second, its efficiency can also be raised in place. The moment of gaining efficiency is still mostly the purchase, and how often homes change hands depends on the age of their owners and on the costs of transacting. Older owners therefore gain efficient homes more slowly than younger ones, even if no owner's incentive to renovate changes with age. Building codes reach only new homes (Levinson, 2016; Jacobsen and Kotchen, 2013), and most of the homes that will stand in 2035 are already built.[9]

We contribute to four literatures. The first studies how the horizon shapes investment (Becker, 1962; Ben-Porath, 1967) and finds that a longer life raises schooling (Jayachandran and Lleras-Muney, 2009; Oster et al., 2013; Fortson, 2011), though which horizon counts is still debated (Hazan, 2009; Cervellati and Sunde, 2013). We answer this question for an investment embodied in a durable that wears off and is left behind at a move: its payoff horizon is short at every age, so the remaining lifetime is not what sets the age profile of investment; the timing of transactions is. The second asks why households appear to underinvest in energy efficiency (Jaffe and Stavins, 1994; Allcott and Greenstone, 2012; Gerarden, Newell, and Stavins, 2017; Fowlie et al., 2018), and often points to high implicit discount rates or heterogeneity in them (Hausman, 1979; Newell and Siikamäki, 2015; Bento, Li, and Roth, 2012). Low efficiency among older owners need not reflect high discount rates or misperceived savings; in our model it reflects when the cheap moment to invest arrives. Buyers pay a premium for a better certificate (Fuerst et al., 2015), and ratings bunch just above band thresholds (Collins and Curtis, 2018), so the certified band carries value of its own; we take the band as the object of choice and its premium as part of the payoff. The third builds lifecycle models of housing with tenure, home size, mortgages, and moving (Cocco, 2005; Yang, 2009; Attanasio, Bottazzi, Low, Nesheim, and Wakefield, 2012; Kaplan et al., 2020) and studies saving and housing in old age (De Nardi, French, and Jones, 2010; Davidoff, 2010; Nakajima and Telyukova, 2017, 2020). We add a quality attribute of the home that can be bought with it or added in place, and that wears off. The fourth shows that transaction costs reduce mobility (van Ommeren and van Leuvensteijn, 2005; Hilber and Lyytikäinen, 2017; Best and Kleven, 2018).[10] We show that transaction costs, through mobility, also slow the upgrading of the

[9] England had 25.6 million dwellings in March 2024 and added about 210,000 in the following year (Ministry of Housing, Communities and Local Government, 2025). At that rate, homes already built will make up more than nine-tenths of the stock in 2035.

[10] A mortgage locked in at a low rate also reduces mobility (Quigley, 1987; Ferreira, Gyourko, and Tracy, 2010; Fonseca and Liu, 2024). The model has no such lock-in, but it keeps the mortgage rules under which a borrower who moves is more

quality of the housing stock. Three papers are closest to ours. Fleischer, Hambel, and Koch (2026) solve the renovation problem of a Dutch household from age 60 that stays in its home, and find that mortality risk and liquidity constraints delay renovation; our owners can also buy an efficient home, and it is this second route, not the shorter horizon, that separates young owners from old. The two findings on the horizon do not conflict. In their model a renovation lowers energy use for good, so its benefit lasts as long as the household lives and mortality cuts it short. In ours efficiency wears off in about six years, so the benefit is short at every age and the remaining lifetime changes it little. Dröes and van der Straten (2024) show in Dutch survey data that households with higher incomes both improve their homes more often and sort into homes that are already efficient; we study the same two margins by age, in a model that explains why older owners rely less on buying and how policy acts on each margin. Reusens, van Kempen, Vandenbergh, Vastmans, and Damen (2026) estimate that the Flemish obligation on buyers of the least efficient homes to renovate within five years lowered the price of these homes by about 2 percent; we simulate a stricter standard, under which every home that is bought must be efficient, and trace which owners it reaches and how it changes how often they move. Methodologically, we solve a large lifecycle model with neural networks, building on Maliar et al. (2021), Azinovic et al. (2022), and Druedahl and Røpke (2026); Druedahl et al. (2026).

Section 2 describes the data and three facts. Section 3 sets out the model and the three propositions. Section 4 describes the solution and the estimation. Section 5 separates the horizon from the window, and Section 6 studies policies. Section 7 concludes.

2. Data and Stylized Facts

In this section, first we describe different data sources we use to establish the stylized facts that motivate our model and to calibrate it. Then we describe how we measure energy efficiency. And, finally, we document three stylized facts about energy efficiency of owner-occupied homes.

2.1. Data sources and sample construction We use two main data sources. First, the English Housing Survey (EHS), and second, the register of Energy Performance Certificates (EPCs) for England and Wales. In the EHS, an annual cross section of households is surveyed in England. The survey combines a household interview with a physical inspection of subsamples of dwellings. The physical inspection is conducted by trained surveyors. The household interview collects information on income, mortgages, ownership, and other associated variables. The physical inspection collects information on the size of the dwelling, building fabric, space heating, hot water, etc. Based on this inspection the survey computes the dwelling's Standard Assessment Procedure (SAP) rating and its EPC. For our purpose, we use the interview waves of 2011–12 to 2019–20 and supplement them with physical survey waves covering that period.

In England and Wales, a certificate is typically lodged when a dwelling is built, sold, or rented.

restricted than one without a mortgage (Section 3.1). Like transaction costs, this makes moving less attractive for borrowers, and a household that moves less often has fewer chances to buy an efficient home. We would therefore expect the same channel to apply.

Lodged certificates are collected in the register of EPCs. We use this register and restrict attention to the period October 2008 to January 2023. The register provides a unique identifier for each dwelling, detailed location information, inspection and lodgement dates, dwelling construction-year bands, reason for the certificate issue, tenure of the dwelling, current and potential energy rating bands, type of the dwelling, current and potential costs of heating, hot water, and lighting, and a set of recommendations in a separate file which documents how the dwelling's energy efficiency can be improved and what can be the lower and upper bounds on the cost of those improvements. In the register we observe 24.1 million domestic certificates, of which 12.2 million are for owner-occupied dwellings. To create a panel of EPCs, within each dwelling we first order certificates by lodgement date and then form pairs of consecutive certificates from different years.

We also use four additional datasets which are used for calibrating the model and for generating data moments for the estimation. These datasets are: (i) the Wealth and Assets Survey (WAS) which provides information on household balance sheets, (ii) Understanding Society which provides information on earnings histories and address changes, (iii) national life tables which provide survival probabilities, and (iv) domestic energy prices which are combined with modelled energy use from the EPC register to compute energy bills by efficiency state.

In the EHS analysis below, our unit of observation is the household; in the EPC register analysis it is the dwelling. As we are interested in the age gradient in energy efficiency, we define age as the age of the household reference person. Further, in most of the analysis we restrict attention to the stock of homes built by 1990, which is the boundary between the older and newer construction-year bands in both data sources. We do this because newer homes are efficient largely because they were built to tighter building regulations, not because of choices made by their owners. We provide additional details about specific variables and sample selection in Appendix A.

We measure energy efficiency of a dwelling through its Energy Performance Certificate (EPC). An EPC is an assessment of a dwelling's energy performance. It is derived from the Standard Assessment Procedure (SAP) rating (BRE, 2014). The SAP rating is based on the estimated energy costs of heating, hot water, and lighting per square meter of floor area.[11] The EPC rates a dwelling on seven bands, from A (most efficient) to G (least efficient).[12] In our empirical analysis, to map the data to the model, we classify bands A–C as efficient and bands D–G as inefficient.[13]

Now we present three stylized facts about energy efficiency of owner-occupied homes using the EHS and the EPC register.

2.2. Three stylized facts

[11] These estimates are made under standard assumptions about occupancy and heating patterns. The rating therefore does not incorporate metered consumption, and does not take into account how the occupants use energy.

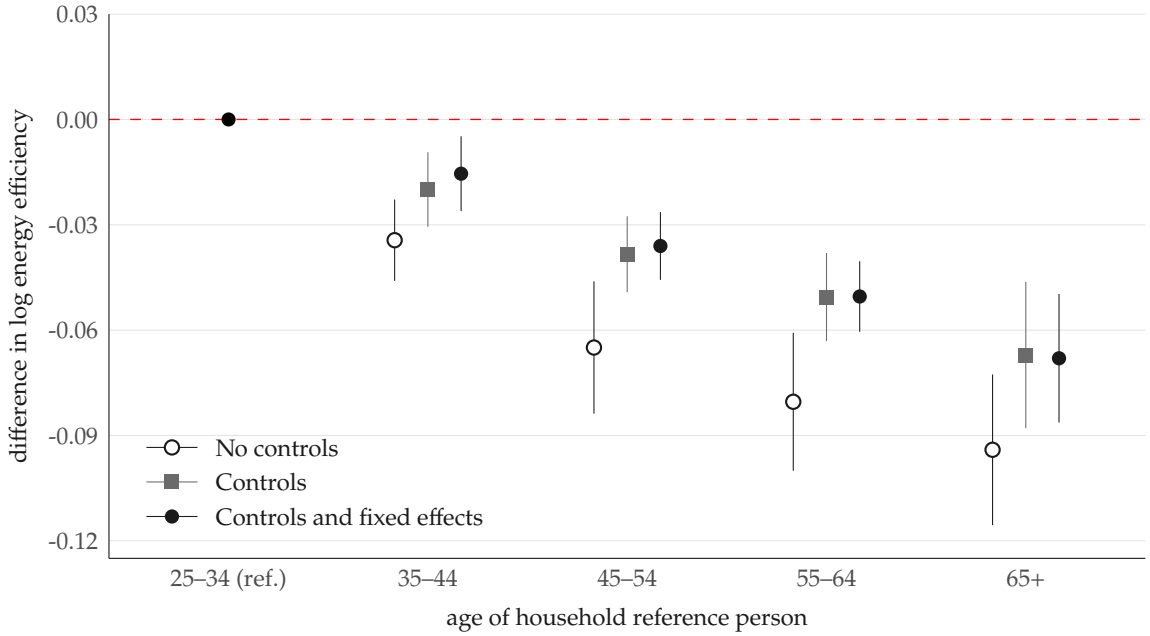
[12] It is important to note that a change in band might not reflect physical work on the dwelling. It can also reflect a different assessor, a revision of the SAP methodology, or measurement error in the rating.

[13] Band C is the policy benchmark in England. The Clean Growth Strategy sets the target that as many homes as possible reach band C by 2035 (HM Government, 2017).

2.2.1 Fact 1: older owners live in less efficient homes.

Using the EHS data we regress the log SAP rating, a continuous measure of energy efficiency, on indicators for the age group of the household, with ages 25–34 as the omitted group. Figure 1 reports coefficients with confidence intervals. The standard errors are clustered by dwelling age and region. We observe that, without controls, on average, owners aged 65+ live in homes that are about 9 percent less energy efficient than those of owners aged 25–34. Controlling for floor area, income, household size and composition, dwelling age, and labor force status reduces the gap to about 7 percent. Further, adding year and region fixed effects does not have any meaningful effect. Hence, in all three specifications efficiency decreases steadily with age. While observable differences between households and between their homes matter, the age gradient remains significant even after accounting for them.

Figure 1: The age gradient in energy efficiency, owner-occupiers



Source: English Housing Survey 2011–12 to 2019–20

Notes: This figure shows coefficients and 95 percent confidence intervals from regressions of log energy efficiency on indicators for the age of the household reference person. The sample is restricted to owner-occupiers. The age group 25–34 is omitted. Energy efficiency is derived from the SAP rating, so higher values mean better performance. "No controls" is unconditional and is denoted by a hollow circle. The second specification, denoted by a gray square, adds log floor area, log real income, household size, an indicator for a single-person household, dwelling-age categories, and the labor force status of the household reference person. The third specification, denoted by a solid circle, adds year and region fixed effects. Standard errors are clustered by dwelling age and region (81 clusters).

Figure 1 uses the continuous measure of energy efficiency. Because in the model we develop in the next section energy efficiency is characterized in a binary form, in Table 1 we show the same gradient in the binary measure. Among all owned homes, the share in bands A–C falls from 37.9 percent at ages 25–34 to 13.3 percent at ages 65–75. However, a part of this difference comes from the differences in the age of the dwelling, since homes built after 1990 tend to be more efficient. In the EHS, 29 percent of young owners live in a home built after 1990 compared to only 9 percent of owners aged 65–75. Even so, when we restrict our attention to the stock built by 1990 the age gradient

remains. The A–C share is 21 percent at ages 25–34 and declines to 16.9 percent at 35–64, and to 9.7 percent at 65–75. Hence, young owners of older homes are more than twice as likely as owners aged 65–75 to live in an efficient one.

Table 1: Energy efficiency and dwelling vintage by age, owner-occupiers

	All owned homes		Homes built by 1990
Age of household reference person	Share in bands A–C	Share built after 1990	Share in bands A–C
25–34	0.379	0.290	0.210
35–64	0.250	0.168	0.169
65–75	0.133	0.093	0.097

Source: English Housing Survey 2011–12 to 2019–20
Notes: Weighted means for owner-occupiers, English Housing Survey physical-survey waves, each wave weighted equally. The rows for ages 25–34 and 35–64 use five waves; the row for ages 65–75 uses the two waves that record exact age (2011–12 and 2012–13). The last column is the efficiency target in estimation.

As the EHS is not a panel, both Figure 1 and Table 1 confound age with cohort effects. Hence, based on this stylized fact only we cannot interpret the gradient causally. However, the structural model developed in the next section supplies the structure that separates these channels.

2.2.2 Fact 2: efficiency status changes in both directions.

While the first stylized fact compares the homes of different owners at a point in time, it does not tell us whether the efficiency of a given home changes over time. Owners can upgrade an inefficient home, and an efficient home can become inefficient due to depreciation. As the EHS does not follow dwellings over time, we use the EPC register. We first pair consecutive certificates for the same owner-occupied dwelling. We then measure how often homes cross the efficiency threshold in each direction.

Table 2 reports annual transition rates across the efficiency threshold between consecutive certificates for the same dwelling. Among homes built by 1990, 6.9 percent of homes in bands D–G move to A–C each year, and 15.2 percent of homes in bands A–C move back to D–G. The downgrade rate is about twice as large as the upward rate. But as there are more homes in bands D–G, the absolute number of homes that move up is larger than the number that move down. Newer homes move up more often (11.9 percent) and down less often (5.8 percent). Hence, there is churning in both directions.[14]

2.2.3 Fact 3: efficiency is tied to when the owner bought.

The first two stylized facts show that older owners live in less efficient homes and that the efficiency status of a home is not constant but changes over time. However, they do not tell us when these

[14] As certificates are typically issued during sales and lettings, the rates captured in Table 2 describe homes that are re-certified. Hence, this provides a partial picture of the housing stock. The appendix reports the comparison between the first and latest certificate of each home and the full matrix of band transitions (Figures A1 and A2).

Table 2: Transitions across the efficiency threshold between EPC certificates, owner-occupied homes

	D–G to A–C		A–C to D–G	
Dwelling vintage	Annual rate	Certificate pairs	Annual rate	Certificate pairs
All owner-occupied homes	0.074	1,407,637	0.111	298,267
Built by 1990	0.069	1,295,552	0.152	171,041
Built after 1990	0.119	104,678	0.058	123,400

Source: Register of Energy Performance Certificates for England and Wales, October 2008 to January 2023

Notes: Pairs of consecutive certificates for the same owner-occupied dwelling, in different years, from the register of Energy Performance Certificates for England and Wales. The annual rate averages, over pairs that start in the origin state, an indicator for ending in the other state divided by the number of years between the two certificates. The row for all homes includes pairs whose construction era is unknown. A transition is a change in the assessed band, not necessarily physical work on the home.

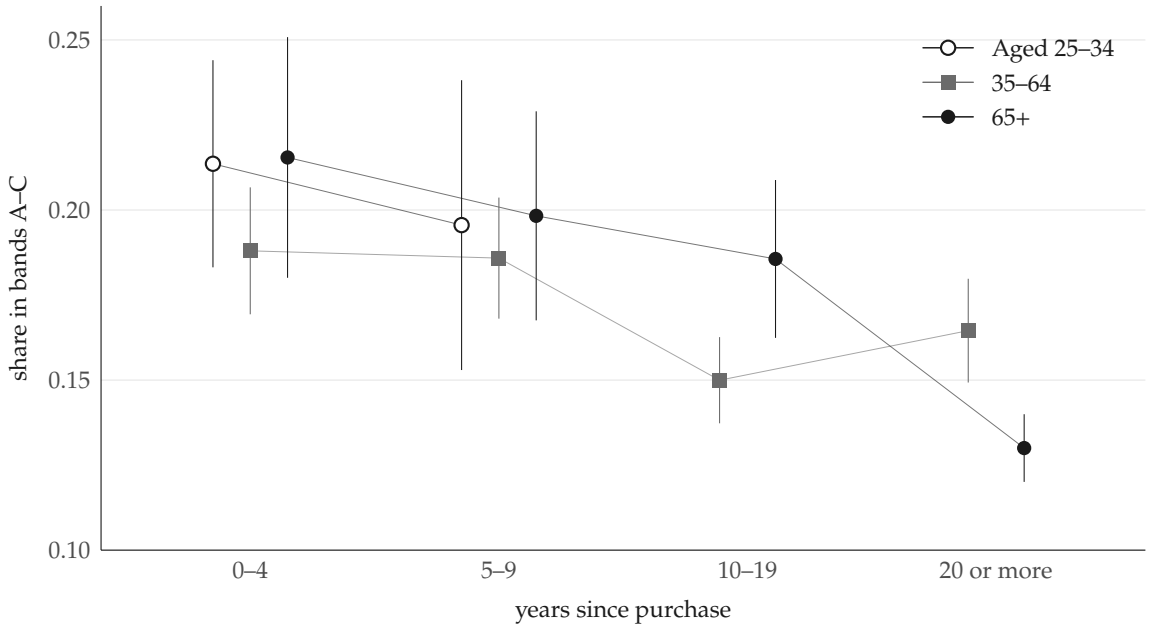
transitions across efficiency bands happen. For example, if an owner improves a home mainly around the time she buys it, and then its efficiency wears off afterwards, then the efficiency of a home should depend on how long ago its owner bought it, and not only on the owner's age. Further, as older owners have, on average, owned their homes for longer, this could account for part of the age gradient. As the EHS records how many years the household has owned its current home, it allows us to compare owners of the same age who bought at different times. For this analysis we restrict attention to homes built by 1990.

Figure 2 shows the efficient A–C share among owners of homes built by 1990 against the number of years since they bought the home, conditional on age. Among owners of age group 65+, the share is 21.5 percent for those who bought within the last five years. This reduces to 13 percent for those who bought 20+ years ago. For owners aged 35–64 the profile is flatter: about 19 percent for those who bought within five years, 15 percent at 10–19 years, and 16.5 percent at 20+ years. Owners aged 25–34 are observed only at short durations, with shares of 21.4 and 19.6 percent. The interesting takeaway from this figure is that owners who bought recently live in similarly efficient homes at every age. The difference seems to arise largely from the varying lengths of ownership.

Figure 2 provides an alternative way to interpret the age gradient. Older owners could have a lower A–C share for two reasons. First, conditional on years since purchase, they could live in less efficient homes than younger owners. Second, they could have owned their homes for longer. As over time efficiency degrades, homes owned for longer are less often in bands A–C. Figure 2 indicates that the second reason might be more important in determining the age gradient.

Till now we have shown that (i) older owners live in less efficient homes, (ii) the efficiency status of a home changes in both directions, and (iii) efficiency is associated with length of ownership. Two explanations are consistent with these stylized facts. First, because of a higher probability of mortality in later life, older owners have fewer expected years left in which to recover the cost of a renovation. Hence, a renovation is worth less to them. Second, mobility declines with age (see Appendix A.5), so

Figure 2: Energy efficiency by years since purchase and age, owners of homes built by 1990



Source: English Housing Survey 2011–12 to 2019–20
Notes: Share of owner-occupiers in homes built by 1990 whose home is in bands A–C, by the number of completed years the household has owned the home and the age of the household reference person, with 95 percent confidence intervals. English Housing Survey, five physical-survey waves (2011–12 to 2019–20), each wave weighted equally. Cells with fewer than 100 households are not shown.

older owners make fewer purchases. If efficiency is mostly acquired at purchase, fewer purchases mean fewer efficient homes, and the age gradient follows.

Fact 3 seems to be consistent with the second explanation. However, it does not settle the question. This is because years since purchase is not assigned at random. The two explanations also interact, because a short horizon is itself a reason not to move. The available data also does not allow for the separation of these effects. The EHS is a cross-section, so it mixes age effects with cohort effects. The EPC register observes a home only when it is certified, so it captures renovation by owners who do not sell only when it leads to a new certificate. We therefore develop a structural lifecycle model in which households choose when to move, what home to buy, and when to renovate.

3. Model

In this section we build a partial equilibrium lifecycle model in which households face idiosyncratic income risk, mortality risk, and borrowing constraints. Households choose consumption, liquid saving, whether to rent or own, the size of their home, mortgage debt, and the energy efficiency of the home they own. Efficiency is assumed to be binary, which keeps the household's choice set manageable and allows us to map model moments to the data moments outlined in Section 2. We borrow many ingredients which are standard in lifecycle models of housing (Cocco, 2005; Piazzesi and Schneider, 2016; Kaplan et al., 2020; Yang, 2009).

The two explanations of Section 2 correspond to two forces in the model. The first is the *horizon channel*. The idea is that an efficient home pays off only for as long as its owner lives in it and it stays efficient. The second force is the *purchase window*: a buyer can choose an efficient home and pay a small price premium, whereas an owner of an inefficient home who stays must pay for a renovation. As older owners buy less often, they have fewer of these opportunities. It is important to emphasize that we impose no preference for efficiency that changes with age.

Section 3.1 describes the environment, Section 3.2 the household problem, and Section 3.3 the forces that shape the age gradient. Parameter values are in Section 4 and further details in Appendix B.

3.1. Environment

Demographics. Time is discrete, and one period is one year. Households enter the model at age $J_0 = 23$, work until age 64, and retire at $J^{\text{ret}} = 65$. A household of age j survives to the next period with probability ϕ_j. Ages $\bar{J} = 82$ and over form a single group in which the household survives each year with the same probability, $\phi_{\bar{J}}$, and faces the same problem every year. No household is forced to die at a fixed age.[15]

Preferences. A household maximizes expected lifetime utility

$$\mathbb{E}_{J_0}\left[\sum_{j=J_0}^{\infty}\beta^{j-J_0}\Big(\prod_{i=J_0}^{j-1}\phi_i\Big)\Big(u_j(c_j,s_j)+(1-\phi_j)\,g(q_j)\Big)\right], \tag{3.1}$$

where c_j is nondurable consumption, s_j is the consumption of housing services, q_j is the bequest the household leaves if she dies at the end of age j, and $\phi_j = \phi_{\bar{J}}$ for $j \geq \bar{J}$. The product of survival probabilities is the probability of being alive at age j. Hence, as $\phi_{\bar{J}} < 1$, the sum is finite. Section 3.2 adds a taste shock to each discrete choice. Period utility is

$$u_j(c,s) = e_j\,\frac{\left(c^{\alpha}s^{1-\alpha}\right)^{1-\sigma}-1}{1-\sigma}, \tag{3.2}$$

where α is the weight on nondurable consumption, $1/\sigma$ is the elasticity of intertemporal substitution, and the equivalence scale e_j captures changes in household size with age. Bequests give warm-glow utility

$$g(q) = \upsilon\,\frac{(q+\overline{q})^{1-\sigma}-1}{1-\sigma}, \tag{3.3}$$

where υ measures the strength of the bequest motive and $\overline{q}$ makes it a luxury (De Nardi, 2004; De Nardi et al., 2010). The bequest q is liquid wealth plus the market value of the home, at its efficiency in the year of death, minus the mortgage. Hence, an efficient home passes its price premium to the heirs only if it is still efficient when the owner dies.

[15] A fixed terminal age would make the last owners sell or stop renovating for mechanical reasons. With a recurring last group, the oldest households keep a constant expected remaining lifetime, $1/(1-\phi_{\bar{J}})$ years.

Income. We follow Carroll and Samwick (1997); Gourinchas and Parker (2002); Blundell, Pistaferri, and Preston (2008) to model income. During working age, a household's gross income is the product of a deterministic age profile and two stochastic components,

$$y_j = \kappa_j \, z_j \, \nu_j, \qquad z_{j+1} = z_j \, \psi_{j+1}, \qquad j < J^{\text{ret}}, \tag{3.4}$$

where κ_j is the deterministic age profile, z_j is the permanent component, ψ_j is a permanent shock, and ν_j is a transitory shock. The shocks are lognormal and their distributions are given as

$$\log \psi_j \sim \mathcal{N}\left(-\tfrac{1}{2}\sigma_\psi^2, \, \sigma_\psi^2\right), \qquad \log \nu_j \sim \mathcal{N}\left(-\tfrac{1}{2}\sigma_\nu^2, \, \sigma_\nu^2\right), \tag{3.5}$$

so that $\mathbb{E}[\psi_j] = \mathbb{E}[\nu_j] = 1$. We also assume that the shocks are independent of each other, across households, and over time. The permanent component is therefore a random walk in logs.

At retirement the permanent component is fixed at its value at age $J^{\text{ret}} - 1$, and the transitory shock is shut off. A retired household receives the pension

$$y_j = \min\left\{\max\left\{\rho_{\text{ret}} \kappa_{J^{\text{ret}}-1} z_{J^{\text{ret}}-1}, \, \underline{y}^{\text{ret}}\right\}, \, \overline{y}^{\text{ret}}\right\}, \qquad j \geq J^{\text{ret}}, \tag{3.6}$$

where ρ_{ret} is the replacement rate of the household's permanent income in her last working year, $\underline{y}^{\text{ret}}$ is a floor, which stands for the minimum income the state guarantees to pensioners, and $\overline{y}^{\text{ret}}$ is a cap on the pensions of the highest earners. The pension is certain, and it is also the household's permanent income in retirement. Earnings and pensions are taxed according to

$$y_j^{\text{net}} = \tau_0 \, y_j^{1-\tau_1}, \tag{3.7}$$

where τ_1 measures the progressivity of the tax system and τ_0 its average level (Heathcote, Storesletten, and Violante, 2017). All money amounts in the model are in units of mean working-age gross income.

Housing and energy efficiency. A home has two characteristics. First, the quantity of housing it provides, h, and, second, its energy efficiency type, $t \in \{N, R\}$, where N denotes an inefficient home and R an efficient one. The quantity h summarizes floor space, location, and the other attributes of a home that are valued by its occupant. We simply call it the size of the home.

Prices are constant, since the model is partial equilibrium. The consumption good is the numeraire, and we choose the units of h so that an inefficient home costs one unit of consumption per unit of size, $p^N = 1$. An efficient home of the same size costs more,

$$p^R = (1 + \gamma_p) \, p^N, \tag{3.8}$$

where $\gamma_p \in (0, 1)$ is the price premium of efficiency. The premium is what a buyer pays for an efficient home over an inefficient home of the same size, and what a seller of an efficient home receives over the seller of an inefficient one. A home of size h and efficiency t is worth $p^t h$.

An owner of a home of size h and efficiency type t receives the flow of housing services given as

$$s(h,t) = \left[\omega_0 + \eta\,\mathbf{1}\{t = R\} + \omega_h \log(1 + h/h_\omega)\right] h. \tag{3.9}$$

The term in brackets is the flow of services per unit of size. For an inefficient owned home it is $\omega_0 + \omega_h \log(1 + h/h_\omega)$, which rises with size at a decreasing rate. The constant h_ω scales the size. A unit of inefficient rented housing yields one unit of services, so the difference, $\omega_0 - 1 + \omega_h \log(1 + h/h_\omega)$, is the extra value an owner places on each unit of her own home, for instance because of security of tenure and control over the home. This extra value rises with size, so owning is worth more relative to renting for larger homes, such as houses rather than flats. The parameter η is the comfort value of efficiency. An efficient home is warmer, less draughty, and less damp than an inefficient home of the same size, and so gives more services.

Although owning a home bestows benefits through higher services, it also comes with several costs. There are two direct costs. First, each period the owner pays a proportional maintenance cost $\delta_h p^t h$, which offsets physical depreciation. Second, the owner also pays an energy bill $e^t h$. We assume $e^R < e^N$ implying that an efficient home needs less energy to heat.

An efficient home pays its owner in three ways: (i) more comfort, (ii) lower energy bills, and (iii) a higher value when it is sold or bequeathed. All three gains are proportional to the size of the home.

We do not assume that efficiency is an absorbing state. An efficient owned home stays efficient the next year with probability $\varrho \in (0, 1)$. This transition from efficient to inefficient type reflects the downgrades in heating systems, the wear of insulation, and a rising standard against which a home is rated. To simplify the solution of the model we assume that the probability does not depend on the time since the home was renovated.

There are three ways to own an efficient home. A household can (a) buy a home that is already efficient, (b) buy an inefficient home and renovate it in the year of purchase, or (c) renovate the inefficient home she already owns. A renovation of a home of size h has both a fixed and a proportional cost, given as

$$C(h) = \xi_f + \xi\, p^N h \tag{3.10}$$

where ξ_f is the fixed cost and ξ the proportional cost. A buyer who chooses option (a) pays only the price premium on the new home, $\gamma_p p^N h'$, plus the transaction cost on it. An owner who stays can only take route (c) and must pay $C(h)$. When the premium is small relative to the cost of a renovation, $(1 + \zeta_b)\gamma_p p^N h < C(h)$, for a household that moves, buying an efficient home costs less than renovating.[16] This is the window of the title: a purchase is the cheaper time to obtain efficiency. Note that it does not come from a discount on renovation at purchase, as the cost of a renovation is the same whether or not the household moves. Under the same condition, route (b) costs more than route (a). Because of the taste shocks, however, some households still choose it, and its mortgage is set

[16] A household that buys a home of size h can end up in an efficient home in two ways. Buying it efficient costs $(1 + \zeta_b)p^R h$ instead of $(1 + \zeta_b)p^N h$, an extra outlay of $(1 + \zeta_b)\gamma_p p^N h$ by equation (3.8). Buying it inefficient and renovating it, at purchase or later, costs an extra $C(h)$. Both routes leave the household with the same asset, so the condition compares their outlays. For an owner who would not otherwise move, buying an efficient home also costs the transaction costs of the move itself.

against the inefficient value of the home.

It is important to note that the model has no dwelling age. We interpret its housing stock as the stock of homes built by 1990, in which efficiency is the result of retrofits rather than of construction standards.

Owning also incurs an indirect cost due to the home being an illiquid asset. Buying a home costs a share ζ_b of its value and selling one a share ζ_s. These transaction costs induce friction in mobility of households across homes (Grossman and Laroque, 1990).[17]

The other option to get housing services is through renting. We assume that a renter can choose the size of her home every year without incurring any transaction cost. She pays rent p^{rent} per unit of size, and pays no maintenance cost. To reduce the number of discrete choices, we assume that the renter does not choose efficiency. A rented home can be of efficient type with probability π_R. Hence, a renter pays the expected energy bill $\left[\pi_R e^R + (1-\pi_R)e^N\right]h$ and receives the housing services $(1+\pi_R\eta)h$. The rent gives landlords a constant gross yield χ on the value of the homes they let,

$$p^{\text{rent}} = \chi\left[\pi_R p^R + (1-\pi_R)p^N\right] = \chi\left(1+\pi_R\gamma_p\right)p^N, \tag{3.11}$$

so the ratio of the average value of rented homes to rent is constant. Whether a rented home is efficient is decided by its landlord, whom we do not model. The efficiency choice in the model is made by owners, the households for whom the facts are measured.

The size menu of homes available for renting or owning is different. A rented home has size between $\underline{h}^{\text{rent}}$ and $\overline{h}^{\text{rent}}$, and an owned home between $\underline{h}^{\text{own}}$ and $\overline{h}^{\text{own}}$, with $\underline{h}^{\text{rent}} < \underline{h}^{\text{own}} < \overline{h}^{\text{rent}} < \overline{h}^{\text{own}}$. The smallest homes can thus only be rented and the largest only owned, while homes in between can be either, as in Kaplan et al. (2020).[18]

Saving and mortgages. Households can save in a liquid asset with return r_b but cannot borrow without collateral. Hence, $b_j \geq 0$. Homes can be financed with repayment mortgages at rate r_m. A loan of balance m with n years remaining requires the annual payment

$$P(m,n) = \frac{r_m(1+r_m)^n}{(1+r_m)^n - 1}\, m. \tag{3.12}$$

We allow for prepayment. Hence, an owner may pay more than $P(m,n)$, up to the full balance. A new mortgage has a term of L years. At purchase, the loan is constrained by the value of the home and by the buyer's income:

$$m' \leq \lambda_L\, p^t h' \qquad \text{and} \qquad P(m',n') \leq \lambda_D\, y_j^{\text{perm}}, \tag{3.13}$$

[17] Higher transaction costs reduce the moves of owners, in the UK through stamp duty (Hilber and Lyytikäinen, 2017) and in the Netherlands (van Ommeren and van Leuvensteijn, 2005), and owners are less likely than renters to move for a job elsewhere (Battu, Ma, and Phimister, 2008; Head and Lloyd-Ellis, 2012). Lifecycle models of housing treat illiquidity in the same way (Attanasio et al., 2012; Bajari, Chan, Krueger, and Miller, 2013; Kaplan et al., 2020). In our model these transaction costs are the only costs of moving.

[18] In general equilibrium, as is the case with Kaplan et al. (2020), how segmented the rental and owner-occupied stocks are determines whether looser credit raises house prices or the ownership rate. Kaplan et al. (2020) let landlords convert owned homes into rentals at no cost. Our prices and rents are fixed, so the model does not take a stand on this question.

where λ_L is the maximum loan-to-value (LTV) ratio, λ_D the maximum payment-to-income (PTI) ratio, and y_j^{perm} the household's permanent income at age j, that is, her gross income without the transitory shock (the pension in retirement). An owner who moves with a mortgage keeps its remaining term and can borrow at most the balance due under the old schedule plus any increase in the value of the home, and plus the cost of a renovation at purchase. We assume that moving cannot be used to take out equity or to restart the repayment schedule.[19] An owner who renovates in place can borrow the cost of the renovation within the same limits, repaid by the existing payoff date.[20] Appendix B.2 states these limits in full.

Initial conditions. Households enter at age 23 with a joint distribution of income, liquid wealth, tenure, home value, and mortgage, which we take from the data in Section 4.

What the model leaves out. House prices, rents, and interest rates are fixed, and there are no aggregate shocks. Every cohort faces the same environment, so the model has no differences between cohorts; this is what lets it isolate the effect of age from that of cohort. We do not model landlords' decisions, so the extent to which rented homes are efficient is taken as given.

3.2. The household problem

Timing. At the start of the year the household's income is realized, and her cash on hand is $x_j = (1 + r_b)b_j + y_j^{\text{net}}$. She draws a taste shock for each of her discrete alternatives, chooses an alternative, and, within that alternative, chooses consumption, saving, and her mortgage. At the end of the year the household dies with probability $1 - \phi_j$ and leaves her estate. If she survives, she draws next year's income shocks, and an efficient owned home loses its efficiency with probability $1 - \varrho$.

State and choices. At the start of the year a household is described by her age j and the state

$$\mathbf{x}_j = (b_j, m_j, n_j, z_j, v_j, o_j, h_j), \tag{3.14}$$

where b_j is liquid wealth, m_j and n_j are the mortgage balance and its remaining term, z_j and v_j are the components of income, h_j is the size of an owned home, and $o_j \in \{0, N, R\}$ records the household's tenure: $o_j = 0$ for a renter and $o_j = t$ for an owner of a home of type t. Table 3 lists the discrete alternatives available in each tenure option and the underlying continuous choices within each of those options. A renter has four alternatives: (i) to keep renting, (ii) to buy an efficient home, (iii) to buy an inefficient home and renovate it, or (iv) to buy an inefficient home and leave it as it is. An

[19] Without these rules, a household could sell and buy an almost identical home only to take out a new mortgage, and so release equity or restart the repayment schedule. The model allows refinancing only to pay for a renovation. UK mortgages are commonly portable when the borrower moves (Liu, 2023), so keeping the remaining term, as we do, matches UK practice.

[20] When a household renovates, at purchase or in place, the loan is limited by the value of the home before the work, $\lambda_L p^N h$; the value the renovation adds, $\gamma_p p^N h$, is not collateral. A household that buys an efficient home, by contrast, can borrow a share λ_L of the premium.

owner can stay, or sell and then rent or buy. An owner of an inefficient home can also renovate in place. As rental size is chosen each year, it is not a state variable.

Table 3: Discrete alternatives and continuous choices, by tenure at the start of the year

Discrete alternative	Continuous choices
Renter: 4 alternatives	
Rent	c, b', h'
Buy an inefficient home	c, b', h', m'
Buy an efficient home	c, b', h', m'
Buy an inefficient home and renovate it	c, b', h', m'
Owner of an efficient home: 5 alternatives	
Stay	c, b', $\widetilde{P}$
Sell and rent	c, b', h'
Sell and buy an inefficient home	c, b', h', m'
Sell and buy an efficient home	c, b', h', m'
Sell, buy an inefficient home and renovate it	c, b', h', m'
Owner of an inefficient home: 6 alternatives	
Stay	c, b', $\widetilde{P}$
Renovate in place	c, b', $\widetilde{P}$
Sell and rent	c, b', h'
Sell and buy an inefficient home	c, b', h', m'
Sell and buy an efficient home	c, b', h', m'
Sell, buy an inefficient home and renovate it	c, b', h', m'

Notes: c is consumption, b' liquid saving, h' the size of the home rented or bought, m' the mortgage taken at purchase, and $\widetilde{P}$ the mortgage payment of an owner who stays. Consumption is the residual of the budget constraint. Whether an alternative is feasible depends on resources, the size limits, and the borrowing limits.

Taste shocks. The model is agnostic about many reasons why a household moves or renovates in a given period. For instance, a household might move due to a new job, a change in the family composition due to a birth, death or marriage. Or she might decide to renovate due to a failing boiler. We capture these idiosyncratic events through taste shocks. In each period, after observing her state, the household draws a shock ε_j^d where superscript d indicates the alternative she can choose. The shocks are independent across alternatives and over time and follow a Type-I extreme-value distribution with scale σ_ε. We provide more details about this in Section 4. We assume that a taste shock adds directly to the value of an alternative as a whole, so it is the same for every allocation within the alternative. Hence, it affects which alternative the household chooses, but not how she consumes, saves, and borrows once the alternative is chosen.

We divide each period into three stages. In the first stage, at the start of the period, the household knows her state $\mathbf{x}_j$, including this period's income and whether her home is still efficient. At the start of the second stage, she draws the taste shocks and chooses one of the alternatives in Table 3. After the discrete choice she chooses how much to consume and save, the size of her home if deciding to buy or rent, and her mortgage if deciding to buy a home. She pays this period's costs which can include rent or the price of the home she buys, the costs of buying and selling, maintenance, the energy bill,

mortgage payments, and the cost of any renovation. After this she receives this period's utility from consumption and from the services of the home she lives in. Note that in the second stage nothing is random. In the third stage, which is the end of the period, the household has made all choices and knows her saving, her mortgage, and her home. The household then might die with probability $1 - \phi_j$ and leave a bequest. If she survives, she draws next period's income shocks, her efficient home loses its efficiency with probability $1 - \varrho$, and the next period begins at the first stage. Uncertainty is thus resolved twice in a period: the taste shocks between the first and second stages, and death, income, and the loss of efficiency between the third stage and the next period.

We now formalize the household's problem with three value functions, one for each stage of the period described above. $V_j(\mathbf{x}_j)$ is the value of the state at the start of the period, i.e., before the taste shocks are drawn. $Q_j^d(\mathbf{x}_j)$ is the value of choosing alternative d, one row of Table 3. $W_j(\overline{\mathbf{x}}_j)$ is the value at the end of the period, when all choices are made and before the household learns whether she survives. Below, we define them in reverse order.

After her choices, the household carries the position

$$\overline{\mathbf{x}}_j = (b', m', n', z_j, \overline{o}, h') \tag{3.15}$$

out of the period, where b' is liquid saving, m' is the mortgage balance, n' is the remaining term of her mortgage, z_j the permanent component of her income, and h' is the size of her home. The state $\overline{o} \in \{0, N, R\}$ records tenure and the efficiency of an owned home. The position $\overline{\mathbf{x}}_j$ determines everything that can happen next. If the household dies, she leaves the bequest $q(\overline{\mathbf{x}}_j) = b' + p^{\overline{o}} h' - m'$ if she owns and $q(\overline{\mathbf{x}}_j) = b'$ if she rents. However, if she survives, the household starts the next period in the state $\mathbf{x}_{j+1}$, which differs from $\overline{\mathbf{x}}_j$ in two ways. First, next period's income shocks are drawn, and second, an efficient home stays efficient with probability ϱ.

The value for the household from this position is

$$W_j(\overline{\mathbf{x}}_j) = (1 - \phi_j)\, g\big(q(\overline{\mathbf{x}}_j)\big) + \beta \phi_j\, \mathbb{E}\big[V_{j+1}(\mathbf{x}_{j+1}) \mid \overline{\mathbf{x}}_j\big], \tag{3.16}$$

where the first term is the value of the bequest if the household dies at the end of the period. The second captures the value of continuing, which is discounted and weighted by the survival probability. The expectation is over next period's income shocks and over whether an efficient home stays efficient. For the final age group, i.e., $j \geq \overline{J}$, the problem is the same every period, so $V_{j+1} = V_j$ and the value function solves a fixed point.

We now come to the middle stage of the period, when the taste shocks are revealed and the discrete choice is made. Within each discrete choice d, the household chooses the continuous allocation $\mathbf{a}$ allowed by her budget, the home size limits, and the borrowing limits. The allocation determines the current period's consumption c^d, this period's housing services s^d, and the position $\overline{\mathbf{x}}_j^d$ the household carries out of the period. Appendix B.4 gives the budget constraints under different discrete choices.

The value of each alternative d is given as

$$Q_j^d(\mathbf{x}_j) = \max_{\mathbf{a}\in\mathcal{A}_j^d(\mathbf{x}_j)} \left\{ u_j(c^d, s^d) + W_j\left(\bar{\mathbf{x}}_j^d\right) \right\}. \tag{3.17}$$

Note that all alternatives are valued with the same function W_j. They differ only in this period's utility and in the future position they can lead to.

After drawing her taste shocks, the household chooses the alternative with the highest $Q_j^d(\mathbf{x}_j) + \varepsilon_j^d$. Before the shocks are drawn, the value of the state is

$$V_j(\mathbf{x}_j) = \sigma_\varepsilon \log \sum_{d\in\mathcal{D}(o_j)} \exp\left(Q_j^d(\mathbf{x}_j)/\sigma_\varepsilon\right), \tag{3.18}$$

where $\mathcal{D}(o_j)$ is the set of feasible alternatives for the given tenure o_j. Alternative d is chosen with probability $\exp(Q_j^d/\sigma_\varepsilon)/\sum_{d'}\exp(Q_j^{d'}/\sigma_\varepsilon)$ (McFadden, 1974; Rust, 1987). Section 4 describes how we use this formulation within a deep learning framework to solve the model.

3.3. What drives the age gradient in the model In each period in which she lives in an efficient home, an owner receives two benefits: (i) the comfort η, and (ii) the saving $(e^N - e^R)h$ on energy bills. When the owner sells the home or dies, that home is worth $\gamma_p p^N h$ more if it is still efficient. However, all these three gains depend on how long the owner keeps an efficient home. For an owner who owns an efficient home at age j, we measure the payoff horizon from this efficiency as the expected discounted number of years, counting the first, in which the owner is alive, still lives in the home, and the home is still efficient. The expression for the horizon is given as

$$H_j = \sum_{k\geq 0} \beta^k \prod_{i=0}^{k-1} \phi_{j+i}\, \varrho\, \mu_{j+i+1}, \tag{3.19}$$

where the k-th term is the discounted probability that, k years later, the owner is alive, still lives in the same home, and the home is still efficient. Notice that the term for $k = 0$ equals one, implying that the current year counts in full. Each additional period multiplies this probability by three factors: (i) ϕ_{j+i}, the probability that an owner of age $j + i$ survives to the next year, (ii) ϱ, the probability that the home stays efficient, and (iii) μ_{j+i+1}, the probability that the owner, now a year older, does not move. Hence, the horizon combines survival probability, the loss of efficiency, and moving probability. The survival probability falls with age, which shortens the horizon of older owners. However, this is counteracted by moving. Older owners move less, which in turn lengthens their horizon. By year k, the probability of having lost efficiency is $1 - \varrho^k$, which can be interpreted as extra discounting. The lower is ϱ, the less weight the horizon puts on later years. These are exactly those years in which the survival of younger and older owners differs most. The age pattern of the horizon then depends on how survival, moving, and ϱ act together, and Section 4.5 reports it at the estimates.

How much the payoff horizon matters depends on how the household gets the efficient home. We make this statement more concrete for the two discrete choices of Table 3 in which efficiency

is chosen. The first choice is when an owner of an inefficient home who stays decides whether to renovate in place. The second is when a household that buys a home decides whether the home is efficient. We characterize both of these choices in a simplified version of the model, which differs from the model of Section 3.2 in two ways. First, there is no income risk. Second, taste shocks enter only an owner's choice between keeping her home and selling it. Each period an owner draws one shock for keeping her home and one for selling it. As no shock enters the choice whether to renovate, an owner that keeps an inefficient home renovates if $Q_j^{\text{stay},R} \geq Q_j^{\text{stay},N}$ where superscript stay refers to the case where the owner keeps the home, and R and N refer to the renovated and non-renovated options, respectively. Similarly, a household that buys a home takes it efficient if that has the higher value. Here μ_j is the probability that an owner of age j keeps her home. For analytical characterizations, we assume that the cost of a renovation, the price premium, the saving on the bill, and the comfort of an efficient home are small relative to the household's wealth.

Renovating in place. Consider an owner of an inefficient home of size h at age j. She can leave the home inefficient or renovate it in place. Let c be her consumption if she does not renovate. Write $s_N = s(h, N)$ and $s_R = s(h, R)$ for the services of the home before and after a renovation. Further, define $\theta \equiv (s_R/s_N)^{(1-\alpha)/\alpha} - 1$. By equation (3.2), $u_j(c, s_R) = u_j\big((1+\theta)c, s_N\big)$. This means an efficient home is worth as much as a share θ more consumption. Let $\Delta e \equiv (e^N - e^R)h$ be the yearly saving on the bill, $\Delta p \equiv \gamma_p p^N h$ the price premium, and $\Lambda \equiv \varrho/(1+r)$ the value today of a unit of wealth that comes with the efficient state next year, where r is the rate at which the household finances the work. Here $r = r_b$ if she pays out of saving and $r = r_m$ if she borrows. Let u_c be the marginal utility of consumption this year, V'_{j+1} the marginal value of wealth next year, and g' the marginal utility of the estate. Then we can define the survival probability weighted by marginal utility as

$$\tilde{\phi}_j \equiv \frac{\beta\phi_j(1+r)V'_{j+1}}{\beta\phi_j(1+r)V'_{j+1} + (1-\phi_j)\,g'}, \tag{3.20}$$

where the numerator captures the marginal value of a pound saved this year that the household will receive if she survives. The denominator is the total marginal value of a pound saved this year, which is received either if the household survives or dies.

In the period of the work, an owner that stays and renovates pays

$$\overline{C} \equiv C(h) + \Delta e - \delta_h \Delta p$$

more than an owner whose home is already efficient. The additional cost includes the cost of renovation, the bill of an inefficient home, which is incurred during the period, less the maintenance on the premium. If the owner sells, the efficient state adds $(1 - \delta_h - \zeta_s)\Delta p$ to the proceeds, since the seller pays the period's maintenance and the selling cost. So $\ell \equiv \overline{C} - (1 - \delta_h - \zeta_s)\Delta p = C(h) + \Delta e - (1 - \zeta_s)\Delta p$ is the part of $\overline{C}$ that a sale does not return.

An owner of an inefficient home has two options, renovate now or wait. Renovating now means one more year of comfort. However, waiting has two advantages. The owner pays a year later. As the

payment is delayed, she learns whether she will still be alive and in the home to enjoy the efficient state. This distinction matters because the market returns only part of the cost: if the owner dies, her heirs receive only the premium, and if the owner is alive next period but moves, a sale returns only part of what she paid. So the owner renovates now if this extra comfort outweighs the cost of paying a year early and the expected losses from dying or moving. The comfort is worth more, in pecuniary terms, to a household that consumes more, so the rule translates into a threshold for consumption. The more likely the owner is to move, the higher the threshold. Proposition 1 makes this argument precise.

Proposition 1 (Renovating in place)**.** *In the simplified model, consider an owner of an inefficient home at age j who keeps her home this year, and suppose that she can pay for the work by saving less or borrowing more at the rate r. To first order, $Q_j^{\text{stay},R} \geq Q_j^{\text{stay},N}$ if and only if*

$$\theta c \;\geq\; \underbrace{C(h) - \Lambda\overline{C}}_{\text{cost of renovating now}} + \underbrace{(1-\tilde{\phi}_j)\big(\Lambda\overline{C} - \Delta p\big)}_{\text{loss at death}} + \underbrace{\tilde{\phi}_j\,(1-\mu_{j+1})\,\Lambda\ell}_{\text{loss at exit}}, \tag{3.21}$$

provided that next year an owner of the inefficient home that keeps it would renovate it. Without this proviso the condition is necessary but not sufficient.

The left side of (3.21) is the comfort this year measured in the units of consumption. The right side is composed of three terms. The first is the cost of renovating now rather than next year. Since $(1+r)\big(C(h) - \Lambda\overline{C}\big) = (r + 1 - \varrho)\,C(h) - \varrho\,(\Delta e - \delta_h \Delta p)$, the first term is, up to the factor $(1+r)$, the interest on the cost and the expected loss of the efficient state, minus next year's bill saving net of maintenance. This is a user cost, and it is higher for a borrower than for a saver. The second is the loss at death. If the owner dies at the end of the year, her heirs receive the premium Δp for the efficient state, rather than the $\Lambda\overline{C}$ that the state is worth to a living owner who would otherwise renovate next year. The term is a gain if $\Delta p > \Lambda\overline{C}$. The last term captures the loss at exit. If the owner moves next year, she receives $(1 - \delta_h - \zeta_s)\Delta p$ for the state rather than $\overline{C}$, and loses ℓ. The owner therefore renovates if and only if her consumption is at least

$$c_j^* = \frac{C(h) - \Lambda\overline{C} + (1-\tilde{\phi}_j)\big(\Lambda\overline{C} - \Delta p\big)}{\theta} + (1-\mu_{j+1})\,\frac{\tilde{\phi}_j \Lambda\ell}{\theta}, \tag{3.22}$$

a threshold that rises in a straight line with the probability of moving. The proof of Proposition 1 is in Appendix C.

Buying an efficient home. A household that buys a home of size h' (prime denotes the new home) can get the efficient state in two ways. She can buy the efficient home, for $(1+\zeta_b)\Delta p'$ more than the inefficient home, or buy it inefficient and renovate it at once, for $C(h')$. Let $C_b \equiv \min\{(1+\zeta_b)\Delta p',\, C(h')\}$ be the cost of the cheaper route. Denote by $Q_j^{\text{buy},R}$ the value of buying an efficient home and by $Q_j^{\text{buy},N}$ the value of buying an inefficient home. Further, let c denote the household's consumption if she buys the inefficient home. In the period of purchase the buyer of the efficient home pays

$C_b - \Delta e'$ more than the buyer of the inefficient home. This is because she pays the efficient bill at once, and no maintenance that year. Then in later periods she pays $\delta_h \Delta p' - \Delta e'$ more than an owner of the inefficient home. The difference, $\overline{C}_b \equiv C_b - \delta_h \Delta p'$, plays the role of $\overline{C}$, and $\ell_b \equiv \overline{C}_b - (1 - \delta_h - \zeta_s)\Delta p' = C_b - (1 - \zeta_s)\Delta p' \leq (\zeta_b + \zeta_s)\Delta p'$ that of ℓ. However, unlike the choice in place, the buyer's choice is not one of timing. Once she has bought the inefficient home, getting the efficient state later costs no less than it would have cost at the purchase, so a buyer near her threshold who passes up the efficient home would not renovate later.

Proposition 2 (Buying an efficient home)**.** *In the simplified model, consider a household that buys a home of size h' at age j. Suppose that $\tilde{\phi}$, μ, Λ, and her consumption c are constant while she owns the home, and that no borrowing limit binds before she moves again. To first order, $Q_j^{\text{buy},R} \geq Q_j^{\text{buy},N}$ if and only if*

$$\theta' c \;\geq\; \underbrace{(1-\Lambda)\overline{C}_b - (\Delta e' - \delta_h \Delta p')}_{\textit{user cost}} + \underbrace{(1-\tilde{\phi})\big(\Lambda \overline{C}_b - \Delta p'\big)}_{\textit{loss at death}} + \underbrace{\tilde{\phi}\,(1-\mu)\,\Lambda \ell_b}_{\textit{loss at exit}}. \tag{3.23}$$

The three terms mirror those of Proposition 1.[21] The user cost equals $C_b - \Delta e' - \Lambda \overline{C}_b$. The loss at death is usually a gain: the heirs receive $\Delta p'$, which is about what the buyer paid for the state and more than the $\Lambda \overline{C}_b$ it is worth to a living owner. The loss at exit is small, since ℓ_b is. The buyer therefore takes the home efficient if and only if her consumption is at least

$$c_b^* = \frac{(1-\Lambda)\overline{C}_b - (\Delta e' - \delta_h \Delta p') + (1-\tilde{\phi})\big(\Lambda \overline{C}_b - \Delta p'\big)}{\theta'} + (1-\mu)\,\frac{\tilde{\phi}\Lambda \ell_b}{\theta'}, \tag{3.24}$$

a threshold that also rises in a straight line with the probability of moving, but with a small slope. Appendix C gives the proof.

Propositions 1 and 2 have the same three terms but differ in what the market returns, and so in how far a short horizon raises each threshold. Proposition 3 compares them; Appendix C gives the proof.

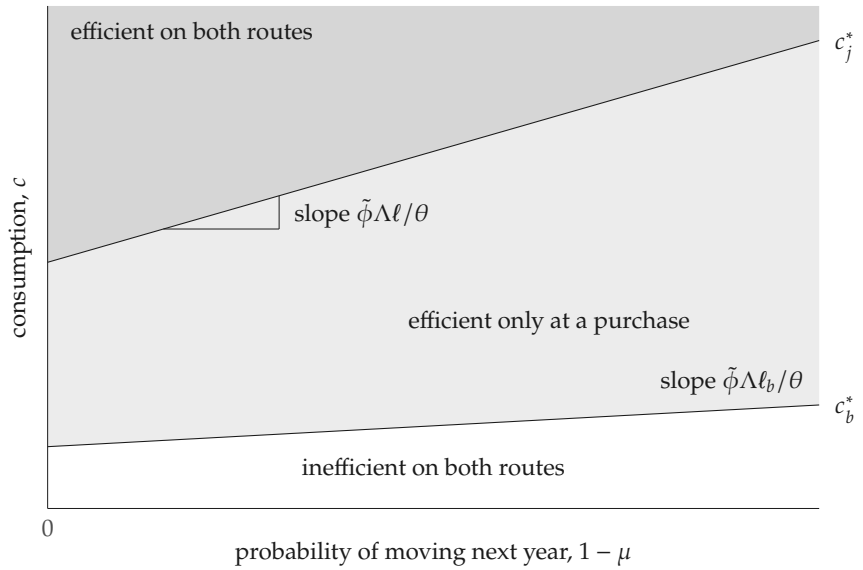
Proposition 3. *In the simplified model, compare an owner of an inefficient home of size h that keeps it with a household that buys a home of the same size. Suppose that both have the same $\tilde{\phi}$, μ, and Λ, and satisfy the assumptions of Proposition 2.*

(a) In place. *A higher probability of moving raises the threshold c_j^* in proportion to ℓ, the part of the cost of the work that a sale does not return. A higher probability of dying also raises it if $\Lambda \mu \overline{C} > \Delta p$, since the heirs then receive less than the efficient state is worth to the owner alive next year.*

(b) At a purchase. *A higher probability of moving raises the threshold c_b^* only in proportion to ℓ_b, which is at most the transaction costs on the premium, $(\zeta_b + \zeta_s)\Delta p'$. A higher probability of dying lowers it if $\Lambda(1 + \zeta_b - \delta_h) < 1$, since the heirs then receive more than the efficient state is worth to the owner alive next year.*

[21] The proposition holds survival, moving, and consumption constant while the household owns the home. Lemma C.3 in Appendix C drops this assumption, provided that an owner of the inefficient home would never renovate it.

Figure 3: The two thresholds and the probability of moving



Notes: Schematic. The consumption above which an owner renovates in place, c_j^* (equation (3.22)), and above which a buyer takes the home efficient, c_b^* (equation (3.24)), against the probability of moving next year, for a next home of the same size and the same $\tilde{\phi}$ and Λ, as in Proposition 3. The slopes are those of its parts (a) and (b), with $\theta' = \theta$ since $h' = h$. Drawn with illustrative values, so the axes carry no units.

(c) The gap. *The in-place threshold is greater than the buyer's threshold, and the gap widens as the probability of moving or of dying rises. A household whose consumption lies between the two thresholds would not renovate in place, but would take an efficient home if she bought one.*

Formally, the slopes in parts (a) and (b) are $\tilde{\phi}\Lambda\ell/\theta$ and $\tilde{\phi}\Lambda\ell_b/\theta'$, and the gap in part (c) is

$$c_j^* - c_b^* = \frac{C(h) - C_b + \Delta e}{\theta \widetilde{H}} > 0, \qquad \widetilde{H} \equiv \frac{1}{1 - \tilde{\phi}\Lambda\mu}, \tag{3.25}$$

where $C(h) - C_b + \Delta e = \overline{C} - \overline{C}_b$ is the extra one-off cost of renovating in place.

Figure 3 draws the two thresholds with the probability of moving on the horizontal axis and the consumption on the vertical axis. In equation (3.25), $\widetilde{H}$ is the payoff horizon of equation (3.19) at constant parameters, with next year weighted by its marginal utility relative to this year.[22] Since μ is a choice probability, parts (a) and (b) compare households that differ in how likely they are to move, for instance because of their stage of life. Moving and death change a threshold only in proportion to the part of the cost that the market does not return. If a sale returned the whole cost, and the heirs valued the state as a living owner does, the threshold would be the same for an owner who moves next year as for one who never moves, and the same at every age. For an owner in place, the part not returned is large.

The gap ties the payoff horizon to age. A short horizon has different sources at different ages:

[22] At constant parameters $H_j = 1/(1 - \beta\phi\varrho\mu)$, while $\tilde{\phi}\Lambda = \beta\phi\varrho V'_{j+1}/u_c$ by the Euler equation. Appendix C shows that both propositions spread a one-off cost over $\widetilde{H}$ (equation (C.3)), and Lemma C.3 extends this to parameters that vary with age.

young owners move often, and old owners are more likely to die. Either source widens the gap (Proposition 3(c)). An owner of an inefficient home whose consumption lies in the gap does not renovate; she waits for her next purchase to get efficiency. The young buy often, so their wait is short. The old rarely buy, so they keep waiting. The horizon decides who waits, and purchases decide how long.

Because an owner in the gap gets efficiency only when she next buys, the frequency of purchases matters. Moving is costly, as the household pays the selling and buying costs and must meet the borrowing limits again. Older owners move less and so buy less often, and the home they live in is more likely to have been bought long ago. If no renovation takes place, a home that was efficient at purchase is still efficient k years later with probability ϱ^k. Older owners' homes are therefore less likely to be efficient.

Three further forces affect the age gradient. First, bequests. The bequest incorporates the efficiency of the home. Hence, the heirs receive the premium. Keeping other factors constant, a stronger bequest motive lowers $\tilde{\phi}_j$, and more so at older ages, which puts more weight on the heirs' value of the state, Δp, and less on that of a living owner (Proposition 3(a) and (b)). This deters a renovation in place whose cost is well above the premium, and favours the purchase of an efficient home, whose cost is close to it. Second, mortgage credit. Λ is lower for a borrower than for a saver, since $r_m > r_b$. Hence, both thresholds are higher for a borrower. Young households buy more often but face tighter borrowing limits, and an owner who renovates in place cannot borrow against the premium of an efficient home before the work is done. Third, comfort. Although the utility value of comfort, η, is the same at every age, its monetary value, θc, moves with consumption. So comfort adds an age gradient through the lifecycle profile of consumption.

Through the lens of a simplified model we show the determinants of acquiring energy-efficient homes. These same forces operate in the full model. However, there every alternative has a taste shock, so the thresholds become probabilities. With income risk, μ and $\tilde{\phi}$ also average over next year's income. Further, the probability of moving can depend on whether the home is efficient. The propositions give the direction of each force, not its size, so they cannot say which force drives the age gradient. We solve this issue in the next section by estimating the full model.

4. Solution and Estimation

The household's problem described in Section 3 is costly to solve. This is because of three reasons. First, the state space is large. Of the seven variables in equation (3.14), five of them are continuous. Second, the discrete choice set is large. A household chooses among four, five, or six discrete alternatives depending on its tenure, leading to fifteen discrete choices in total. Within each discrete alternative it makes two or three free continuous choices.[23] Further, each alternative is valued by an expectation of next period's value over the permanent and transitory income shocks, and over whether an efficient home continues to stay efficient. Also, borrowing limits and minimum size of housing introduce kinks

[23] Liquid saving is determined by the budget constraint once consumption is chosen, so it is not a free choice.

in the value of each alternative, making first order conditions insufficient for an optimum. Third, to provide a coherent story we intend to estimate the model. That implies evaluating the model at different combinations of parameter vectors.

Due to this complexity, we solve the model with deep learning, building on Druedahl and Røpke (2026) and Druedahl et al. (2026). We extend their approach from solving the model to estimating it. Section 4.1 explains how we solve the model, Section 4.2 sets the fixed parameters, Section 4.3 discusses which moments identify the estimated ones, and Section 4.4 describes the estimator. Section 4.5 reports the estimates and the fit.

4.1. Solving the model The solution of the model consists of two functions of the state $\mathbf{x}_j$ at every age j. The first is the value function $V_j(\mathbf{x}_j)$ as defined in equation (3.18). The second is the policy function, which gives, for every alternative d, the continuous choices $\mathbf{a}_j^d(\mathbf{x}_j)$ that attain the maximum in equation (3.17). Note that the discrete choice needs no function of its own, since the probability of each alternative follows from the values Q_j^d. We stack these policy functions across alternatives into $\mathbf{a}_j(\mathbf{x}_j)$. The three value functions of Section 3.2 link one age to the next. Given V_{j+1}, equation (3.16) gives the end-of-period value function W_j. Equation (3.17) then gives the value function Q_j^d of each alternative and the policy function $\mathbf{a}_j^d$ that attains it. And, finally, equation (3.18) gives V_j and the probability of each alternative. We start from the oldest age group, solve the fixed point problem, and then use this recursion to find the solution at every age. Finally, with the value functions and the policy function in hand we simulate households and estimate the model.

The standard way to implement this recursive solution is to store V_{j+1} at the points of a grid and interpolate between them to compute the expectation in equation (3.16). However, as mentioned above, a grid is out of reach here. We instead represent V_j and $\mathbf{a}_j$ by neural networks (Maliar et al., 2021; Azinovic et al., 2022; Druedahl and Røpke, 2026; Druedahl et al., 2026). A grid method stores the function at fixed points and computes it elsewhere by interpolating between the nearest of them. The issue is that it needs points along every dimension. Hence, its number grows exponentially with the number of dimensions. On the contrary, a network is instead a single formula in its inputs. It passes them through several layers, where each of them combines linearly using the previous layer's outputs and then applies a fixed nonlinear function to them. A network's coefficients, called weights, are chosen to fit the function at a sample of points. Once the network is fitted, the formula can be evaluated at any point. Hence, networks can approximate smooth functions of many variables with far fewer coefficients than a grid in as many dimensions would need.[24]

The solution of course depends on the parameters of the model. We set some of them before estimation (Section 4.2) and estimate the others (Sections 4.3 and 4.4). The estimator needs to evaluate

[24]Consider a grid with 50 points in each continuous state. An owner with a mortgage has five continuous states, b, m, z, ν, and h, so the grid has $50^5 \approx 3.1 \times 10^8$ points for each of the two home types and each of up to 25 remaining terms, or up to 1.6×10^{10} points at each age. Storing V_j alone at these points takes up to 125 gigabytes per age in double precision. At every point, the maximization in equation (3.17) must be solved in each of up to six alternatives over two or three continuous choices, and every candidate choice requires the expectation in equation (3.16). Adding the eight parameters as dimensions – the seven we estimate and the discount factor – as in the next paragraph, raises the number of points for each home type and term to $50^{13} \approx 1.2 \times 10^{22}$. By comparison, the value network of each age has about 550,000 weights and is fitted at 20,480 sampled points.

the model at many values of the estimated parameters. For this, a grid method would solve the model multiple times at each of them. We avoid this by making these parameters inputs of the networks, as Druedahl et al. (2026) suggest. Collecting the parameters which we want to estimate in the vector ϑ we can compare the fit at different values of it, and write the value function and the policy function as $V_j(\mathbf{x}_j; \vartheta)$ and $\mathbf{a}_j(\mathbf{x}_j; \vartheta)$. At each age j, a value network and a policy network approximate

$$V_j(\mathbf{x}_j; \vartheta) \approx \widehat{V}_j\big(\mathbf{x}_j, \vartheta; \mathbf{w}_j^V\big), \qquad \mathbf{a}_j(\mathbf{x}_j; \vartheta) \approx \widehat{\mathbf{a}}_j\big(\mathbf{x}_j, \vartheta; \mathbf{w}_j^a\big), \tag{4.1}$$

where $\mathbf{w}_j^V$ and $\mathbf{w}_j^a$ are the weights of the two networks. Note that both networks have the same inputs which are the state $\mathbf{x}_j$, comprising liquid wealth b, the mortgage balance m and its remaining term n, the components of income z and ν, tenure o, and the size of the home h, along with the parameters ϑ which we want to estimate and the discount factor β. Appendix D.1 lists the inputs in full, including transformations of the state that help the fit.

The value network returns the value of the state whereas the policy network returns the continuous choices in every alternative. The two have distinct roles. The value network of age $j+1$ supplies V_{j+1}, which enters the end-of-period value function W_j and hence the value function of every alternative at age j. The policy network approximates the maximizing choices, so that simulating a household does not require solving the maximization in equation (3.17) again. The choices must satisfy the budget and borrowing constraints, so the policy network does not output them directly. It outputs consumption and the shares of the feasible ranges of home size and of the mortgage, which a fixed mapping turns into a feasible allocation with liquid saving as the residual (Appendix D.1).

Backward induction with networks. We do not put age as an input as each age has its own pair of networks. We train the networks age by age, backward from the last age group. At ages 82 and above the problem is the same every year, and we iterate on it as a fixed point problem. Then moving backwards, at age j, the networks of age $j+1$ are already trained and are held fixed. We draw a sample of states and parameter vectors. For each draw, we first compute the end-of-period value W_j in equation (3.16) by integrating $\widehat{V}_{j+1}$ over next year's income shocks by quadrature and over whether an efficient home stays efficient. After this we solve the maximization problem in equation (3.17) for every feasible alternative. This step gives the value of each alternative, Q_j^d, and its maximizing choices. The logsum expression in equation (3.18) turns these values into the value of the state, V_j. The value network is fitted to these values and the policy network to the maximizing choices, and the algorithm moves to age $j-1$. Each network is essentially learning a function that has been computed at the sample points, and its job is to interpolate between them, across states and across parameter vectors.

We cannot employ this method if the value function is not smooth. The taste shocks are what make this step work. Without them, the value of the state would be the largest of the Q_j^d. This maximum has a kink wherever the household is indifferent between two alternatives, and the kinks will move with the parameters. The logsum actually smooths each kink over a range of values of width about σ_ε (Iskhakov et al., 2017), so the value of the state is smooth in the state and in the parameters, which a network fits well.

Simulation. To simulate the model, for each household we draw her income shocks, whether her home stays efficient if she is an owner of an efficient home, whether she survives, and a taste shock for every alternative, every period. The policy network provides the allocation in every alternative. The value of each alternative follows from equation (3.17) with $\widehat{V}_{j+1}$, and the household chooses the alternative with the largest value plus taste shock. All candidates in the estimation use the same draws. Appendices D.1–D.3 describe the networks, the training, and the accuracy of the solution.

4.2. Parameters set before estimation Table 4 lists the parameters we set before estimation along with their sources. Nominal amounts are in units of mean working-age gross income.

Demographics and income. Survival probabilities come from the ONS National Life Tables for 2015–17. In the 82+ group a household survives each year with probability 0.880, which matches the remaining life expectancy at 82 in the same tables. The equivalence scale, the square root of household size as in Kinnerud (2025), follows household size by age in the Wealth and Assets Survey and lies between 1.07 and 1.54.

We estimate the income process (3.4) from waves 1–10 (2009–2019) of Understanding Society. We take the gross income of the household in 2015 prices as our income measure. The sample consists of households with reference persons aged 20–64 who work at least 10 hours a week and are observed in at least five waves. The age profile κ_j is calculated from a regression of log income on age indicators, marital status, the number of children, and education with household fixed effects. We smooth the age effects with a cubic in age. The profile rises from £31,200 at age 23 to a peak of £58,300 at age 58. The residual of this regression is the sum of the permanent and transitory components. Since the permanent component is a random walk and the transitory shock is independent over time, the variance of the change in the residual between two interviews k years apart is $k\sigma_\psi^2 + 2\sigma_\nu^2$ (Carroll and Samwick, 1997). We regress the squared changes for $k = 1, \ldots, 9$ on k. The slope estimates σ_ψ^2 and half the intercept estimates σ_ν^2, which gives $\sigma_\psi = 0.096$ and $\sigma_\nu = 0.206$. A permanent shock of one standard deviation thus changes the household's income by about 10 percent permanently, and a transitory shock changes it by about 21 percent for one year. Appendix A.5 gives the details. We estimate the tax function (3.7) on the same sample, by regressing the log of net household income on the log of gross household income, both measured as deviations from the age profile of gross income. The parameter τ_1 is one minus the slope and τ_0 is the exponential of the intercept, which gives $\tau_0 = 0.777$ and $\tau_1 = 0.064$.

The pension floor in equation (3.6), $\underline{y}^{\text{ret}} = 0.30$ which translates to about £15,000 a year, is a minimum pension income. It lies between the Pension Credit guarantee for a couple, £237.55 a week or about £12,400 a year in 2016/17, and two full new State Pensions, about £16,200 a year. The cap, $\overline{y}^{\text{ret}} = 3.64$, binds only for the highest earners. It follows the maximum benefit in Kinnerud (2025), but we set it high: at the estimated replacement rate it binds only when permanent income in the last working year exceeds 4.7 times mean income. The replacement rate ρ_{ret} between them is estimated.

Table 4: Parameters set before estimation

Parameter	Interpretation	Value	Source
Demographics			
a_0, T	Entry age; number of ages (the last is 82+)	23; 60	Model choice
a^{ret}	Retirement age	65	Model choice
ϕ_j	Survival probabilities, ages 23–81	2015–17 table	ONS National Life Tables, UK
ϕ_{82+}	Annual survival in the 82+ group	0.880	ONS life expectancy at 82
e_j	Equivalence scale	1.07–1.54	WAS household size by age
Preferences			
σ	Relative risk aversion	2	Kaplan et al. (2020)
β	Discount factor	0.97	Criterion profile; see text
α	Weight on non-housing consumption	0.75	Kinnerud (2025)
$\bar{q}$	Bequest shifter	1	Normalization
σ_ε	Scale of taste shocks	0.10	Maintained
ω_0, h_ω	Owner-service intercept; size scale	1.015; 2	Maintained
Income, taxes and pensions			
$\bar{Y}$	Model unit: mean working-age income	£49,914	Normalization
κ_j	Earnings profile, ages 23–64	£31k–£58k	Understanding Society
σ_ψ, σ_ν	S.d. of permanent; transitory income shocks	0.096; 0.206	Understanding Society
τ_0, τ_1	Net income $\tau_0 y^{1-\tau_1}$	0.777; 0.064	Understanding Society
$\underline{y}^{\text{ret}}$, $\bar{y}^{\text{ret}}$	Pension floor; cap	0.30; 3.64	Between a couple's Pension Credit and State Pension; set high
Housing and renovation			
p^N	Price of an inefficient home (numeraire)	1	Normalization
γ_p	Price premium of an efficient home	0.03	In EPC–Price Paid estimates, 2.0–3.5%
ξ	Proportional renovation cost	0.08	In EPC cost-to-price means, 7.6–9.0%
θ_{move}	Fixed-cost multiplier, renovation at purchase	1	No discount (maintained)
δ_h	Maintenance cost, share of value	0.02	ONS capital stock accounts
ζ_b	Buying cost, share of value	0.039	Stamp duty and fees, 2018
ζ_s	Selling cost, share of value	0.050	Attanasio et al. (2012)
$\bar{h}^{\text{own}}$	Maximum owned house size	31.1	WAS home values (1.2 × p99)
Energy			
e^N, e^R	Annual energy bill per unit of housing: D–G; A–C	0.0053; 0.0032	EPC register; DESNZ fuel prices
Rental market			
p^{rent}	Annual rent per unit of housing	0.0503	Gross yield of 5% (Bank of England)
π_R	A–C share of rental homes	0.175	English Housing Survey 2016
$\bar{h}^{\text{rent}}$	Maximum rental size	14.9	EHS rents (1.2 × p99)
Credit			
r_b, r_m	Return on liquid saving; mortgage rate	0.008; 0.034	Bank of England; ONS
λ_L	Maximum loan-to-value at origination	0.90	Attanasio et al. (2012)
λ_D	Maximum payment-to-gross-income ratio	0.33	Carro et al. (2023)
L	Term of a new mortgage (years)	25	Standard UK term
Initial conditions at age 23			
π_0^R	Share of owners in an A–C home	0.218	EHS, owners aged 16–24
	Income, wealth, tenure and mortgages	Records	WAS, heads aged 20–29

Notes: Nominal amounts are in units of mean working-age gross income. Rates are annual.

Housing, renovation, and energy. We set the price premium of an efficient home, γ_p, and the proportional cost of a renovation, ξ, using the EPC register matched to Land Registry Price Paid data, which cover sales in England and Wales in 2008–2024. For the price premium, we regress the log price per square metre on an indicator for a home in bands A–C along with controls for floor area and the number of rooms and fixed effects for the month of sale, the postcode, the property type, the built form and the construction period. The premium lies between 2.0 and 3.5 percent across different specifications. We set $\gamma_p = 0.03$.

Each certificate lists the measures recommended for the home, an indicative cost for each, and the band the home would reach if all were installed. We use this information to calibrate the variable cost of renovation. We divide the total cost by the price of the home. For homes in bands D–G that the measures would bring to A–C, the ratio is 9.0 percent on average, whereas for those it would bring only to band C it is 7.6 percent. The medians are lower, 6.8 and 5.2 percent. We set $\xi = 0.08$, between the two means. Since the ratio covers the whole recommended cost, the fixed cost ξ_f, which we estimate, comes on top of it.

Buying costs $\zeta_b = 3.9$ percent of the value of the home. It covers stamp duty, legal and survey fees, and mortgage fees.[25] Selling costs $\zeta_s = 5.0$ percent of the value of the home, the value that Attanasio et al. (2012) use for selling in a lifecycle model of UK housing. Since $(1 + \zeta_b)\gamma_p = 0.031$ is less than ξ, the extra outlay of buying an efficient home is smaller than the cost of renovating one of the same size, whatever its size and whatever ξ_f: the condition in Section 3.1 holds for every home, and across the ranges above: at a premium of 3.5 percent, $(1 + \zeta_b)\gamma_p = 0.037$, below the lower of the two medians, 5.2 percent.

We set the energy bill per unit of housing at $e^N = 0.0053$ for a home in bands D–G and $e^R = 0.0032$ for a home in bands A–C. We compute them from the energy use that the EPC register models for each band combining it with DESNZ fuel prices (Appendix A.6). Then the rest of the return to efficiency comes from comfort, η, and from the higher price of the home at sale or bequest, which depends on how long efficiency lasts, ϱ. We estimate both η and ϱ.

Renting. Rent per unit of housing follows from equation (3.11). We set the gross yield to $\chi = 5.0$ percent, the ratio of annual rent to value for rented homes in England and Wales in 2016 (Bank of England, 2017, Annex 2, Table A.3). We set $\pi_R = 0.175$ which is the share of homes in bands A–C among privately rented homes built before 1990 in the English Housing Survey.[26] With $\gamma_p = 0.03$, these values give $p^{\text{rent}} = 0.0503$: a renter pays 5.03 percent of the value of an inefficient home of the same size a year, and pays the energy bill on top.

Since $p^N = 1$ and nominal values are measured in units of $\overline{Y}$, the size of a home is its value as an inefficient home in units of $\overline{Y}$, and the size of a rented home is its annual rent divided by $p^{\text{rent}}\overline{Y}$. We

[25] The buying cost is the sum of stamp duty (2.94 percent, the receipts of stamp duty divided by the value of transactions), legal fees (0.34 percent), a survey (0.21 percent), and mortgage fees (0.37 percent). Stamp duty: HMRC, UK stamp tax statistics, 2018. Legal fees and survey: Reallymoving Conveyancing and Surveying Price Indices, 2018, as a share of the average transaction value of £284,864. Mortgage fee: the average mortgage fee of £1,047 in December 2018 (Moneyfacts).

[26] We leave out social housing, whose rents are not set in the market.

set the largest rented home at 1.2 times the 99th percentile of rents in the same sample, an annual rent of £37,400, which gives $\overline{h}^{\text{rent}} = 14.9$. We set the largest owned home at 1.2 times the 99th percentile of home values in the Wealth and Assets Survey, £1.55 million, which gives $\overline{h}^{\text{own}} = 31.1$. Fewer than 1 percent of renters and of owners in the data live in larger homes. The smallest rented home, $\underline{h}^{\text{rent}} = 0.01$, has a rent of about £25 a year. It is a numerical floor that keeps renting feasible for every household, not a feature of the data. The smallest owned home, $\underline{h}^{\text{own}}$, is estimated.

Savings and credit. Liquid saving earns $r_b = 0.8$ percent a year which is the average real return on a long term riskfree bond from January 2009 to December 2019. To compute this, in each month we take the 25 year yield on the Bank of England's nominal gilt yield curve and subtract the twelve month rate of CPI inflation.

The mortgage interest rate is set at $r_m = 3.4$ percent, which is the average over 2009–2019 of the Bank of England's quoted rate on a five year fixed rate mortgage. We set the new mortgage term at 25 years, and at origination its loan to value can be at most 90 percent of the value of the home (Attanasio et al., 2012) and its payment to income at most 33 percent of the household's permanent gross income.[27]

An owner pays a per-period maintenance cost of $\delta_h = 2$ percent of the value of the home. This is the average over the period 2009–2019 of the depreciation of dwellings as a share of their gross capital stock excluding land, both at current prices, in the capital stock accounts of the Office for National Statistics. Note that, given these calibrated numbers, interest on the value of the home adds 0.8 percent forgone on her own savings or 3.4 percent on a mortgage. Hence, owning costs 2.8 percent of the value a year in all if she owns outright and 5.1 percent if she borrows 90 percent of the value.[28] At the same time an owner also receives more services from each unit than a renter. Therefore, owning is cheaper than renting for a household with equity. Buying and selling together cost 8.9 percent of the value, which an owner without a mortgage recovers in about four years at the 2.2 percent a year it saves over renting.

Preferences. Relative risk aversion is $\sigma = 2$, as in Kaplan et al. (2020). The discount factor is $\beta = 0.97$.[29] The weight on nondurable consumption is $\alpha = 0.75$, so housing services take a quarter of the Cobb–Douglas aggregate, and the bequest shifter is $\overline{q} = 1$. The intercept and scale of owner services in equation (3.9) are fixed at $\omega_0 = 1.015$ and $h_\omega = 2$, and the size gradient ω_h is estimated. The taste shocks have scale $\sigma_\varepsilon = 0.10$. This scale sets how strongly choices respond to differences in

[27] Carro, Hinterschweiger, Uluc, and Farmer (2023) report a hard maximum loan-to-income ratio of 5.4 to 5.6. We take 5.5. Repaid over 25 years at 3.4 percent, a loan of 5.5 times income costs about 33 percent of income each year.

[28] The 5.1 percent is $0.9 \times 3.4 + 0.1 \times 0.8 + 2$: 90 percent of the value is borrowed at 3.4 percent, the remaining 10 percent forgoes 0.8 percent, and maintenance is 2 percent.

[29] When β is estimated jointly with the seven parameters of Section 4.3, the estimate is 0.98. That is the upper end of the range $[0.88, 0.98]$ on which the networks are trained, so it is not an interior minimum. The criterion of the estimator in Section 4.4 rises steeply below 0.97: it is 0.048 at $\beta = 0.96$ and 0.109 at 0.95. Between 0.97 and 0.98 it is nearly flat, at 0.016 and 0.013–0.014. The data therefore rule out a lower discount factor but barely distinguish 0.97 from 0.98. The small gain at 0.98 comes from retirement income and the wealth of households over 75. The three renovated shares together fit better at 0.97, and the renovation parameters ϱ, η and ξ_f differ by less than 6 percent between the two values. The value 0.97 is close to the 0.964 in Kaplan et al. (2020).

their values. We do not estimate it, and the estimates are conditional on it.

Initial conditions. Households enter at age 23 with the income, liquid wealth, tenure, home value, and mortgage of households matched to the Wealth and Assets Survey whose head is aged 20–29. The WAS does not record energy efficiency, so an owner starts in an efficient home with probability $\pi_0^R = 0.218$, the share in bands A–C among owners aged 16–24 in the English Housing Survey.

4.3. Estimated parameters and identification We estimate seven parameters, which are: (i) the minimum owned house size $\underline{h}^{\text{own}}$, (ii) the owner-service size gradient ω_h, (iii) the bequest scale υ, (iv) the pension replacement rate ρ_{ret}, (v) the fixed cost of a renovation ξ_f, (vi) the retention of efficiency ϱ, and (vii) the comfort value of efficiency η. To estimate these parameters, we target eleven data moments (Appendix A.7 gives the sample behind each), which are: homeownership at ages 25–34, 35–64, and 65–75; the share of housing in the wealth of households aged 35–64; the ratio of income at ages 65–74 to income at ages 55–64; liquid wealth relative to income at ages 65–75 and 75+; consumption relative to income at ages 50–64; and the share of owners in an efficient home at ages 25–34, 35–64, and 65–75. Below we discuss the identification of these parameters based on the moments we target.

First, homeownership at ages 25–34, 35–64, and 65–75 and the share of housing in the wealth of households aged 35–64 inform $\underline{h}^{\text{own}}$ and ω_h. This is because a larger minimum owned home makes ownership harder to reach. Hence, it has a level effect and lowers ownership in every age group. At the same time, a steeper size gradient makes owning more valuable relative to renting. Hence, it raises ownership and the housing share of wealth. However, compared to $\underline{h}^{\text{own}}$, its effect is not uniform across age groups. It increases ownership mostly among the young because most of them rent.

The ratio of income at 65–74 to income at 55–64 pins down the replacement rate ρ_{ret}. Given the pension income, liquid wealth relative to income at 65–75 and at 75+ pin down the bequest scale υ. A stronger bequest motive makes old households save relatively more for future generations. So, a higher pension lowers the wealth as it crowds out saving for old age, and a stronger bequest motive offsets it. Consumption relative to income at 50–64 summarizes saving before retirement.

The shares of owners in an efficient home at 25–34, 35–64, and 65–75 inform ξ_f, ϱ, and η. All three shares are positively correlated with ϱ and η and negatively with ξ_f. However, proportions are different. The fixed cost lowers the share proportionally most among older owners. Older owners buy less often, so more of their efficiency comes from renovating in place, which pays the fixed cost, while buying an efficient home avoids it (Section 5). The fixed cost therefore governs how steeply the share falls with age, through the width of the gap in Proposition 3.

One issue with jointly estimating the retention of efficiency, determined by ϱ, and its comfort value, η, is that both raise the share at every age in nearly the same proportions. This is because the stock of efficient homes can be large because efficiency is worth a lot while it lasts or because it lasts long. It is hard for the three targeted shares to tell the two apart.[30] The shares therefore are pinning down a

[30] At the estimation point, doubling η raises the share to 27.9, 23.1, and 16.6 percent at 25–34, 35–64, and 65–75 (Section 5),

combination of the two. We deal with this problem in two ways after the estimation. First, we will show that away from the estimate the two parameters shape the age profile of the share differently (Section 5). Second, a moment that we do not target, the rate at which homes lose efficiency, will actually separate these two parameters near the estimate.

4.4. Estimator We denote parameters to estimate by ϑ. These parameters are estimated using the method of simulated moments. The method needs the model at every value of ϑ that the search visits. As mentioned above, we expand the state space of the model to include ϑ. Hence, networks take ϑ as an input. We solve the model only once, for all values of $\vartheta \in \Theta$ at the same time, where Θ is the box of search ranges. Evaluating a new value then requires only a simulation. Training the networks over it took 9.7 hours on one L40S GPU, and Table D3 in Appendix D.1 lists their architecture and training settings, and Appendix D.4 describes the search.

We also compute the standard errors of estimated parameters using the sandwich formula

$$V = A\,(\Omega + \Sigma_S)\,A', \qquad A = (G'WG)^{-1}G'W, \tag{4.2}$$

where G is the derivative of the model moments with respect to the parameters, computed by central differences at ±3 percent of the width of each search range, Ω is the covariance of the data moments, from a cluster bootstrap of the surveys, Σ_S is the covariance of the simulation noise in the model moments, and W is the weighting matrix in the simulated method of moments. Appendix D.5 gives the details.

4.5. Estimates and fit Panel A of Table 5 reports the estimates. The smallest owned home is worth $\underline{h}^{\text{own}} = 3.14$ times mean income. The pension replaces about 77 percent of permanent income in the last working year.

The fixed cost of a renovation is $\xi_f = 0.159$, about £7,900. With the proportional cost, renovating the smallest owned home costs 4.2 times the extra outlay of buying it efficient, about £20,500 against £4,900, and the largest 2.7 times. The ratio falls as house size grows because the fixed cost is spread over a larger home. An efficient home stays efficient with probability $\varrho = 0.838$ a year, so efficiency lasts about six years on average, which at first appears to be low. However, we do not target the rate at which homes lose efficiency, but the model's rate, 16 percent a year, is close to the 15.2 percent of homes built by 1990 that drop from bands A–C to D–G each year between certificates (Table 2).[31] The comfort value $\eta = 0.061$ gives an efficient home about 5 percent more services than an inefficient home of the same size, which is worth about 1.5 percent of nondurable consumption.

The parameters associated with ownership, saving, and pensions are precisely estimated, with relatively low standard errors. However, the three parameters of efficiency are not. The standard

and raising ϱ from 0.84 to 0.88 instead raises it to 28.0, 23.5, and 16.7 percent.

[31] Part of this rate reflects reassessment rather than physical change: Crawley et al. (2019, Table 3) estimate that a home truly in band C has a 19 percent chance of being rated D or below. Efficiency in the model is therefore the certified band. It is the band on which the price premium γ_p is estimated and the band that regulation targets. A household can lose it through wear, a change in the rating method, or reassessment, and ϱ measures how long a home keeps it.

Table 5: Estimated parameters and targeted moments

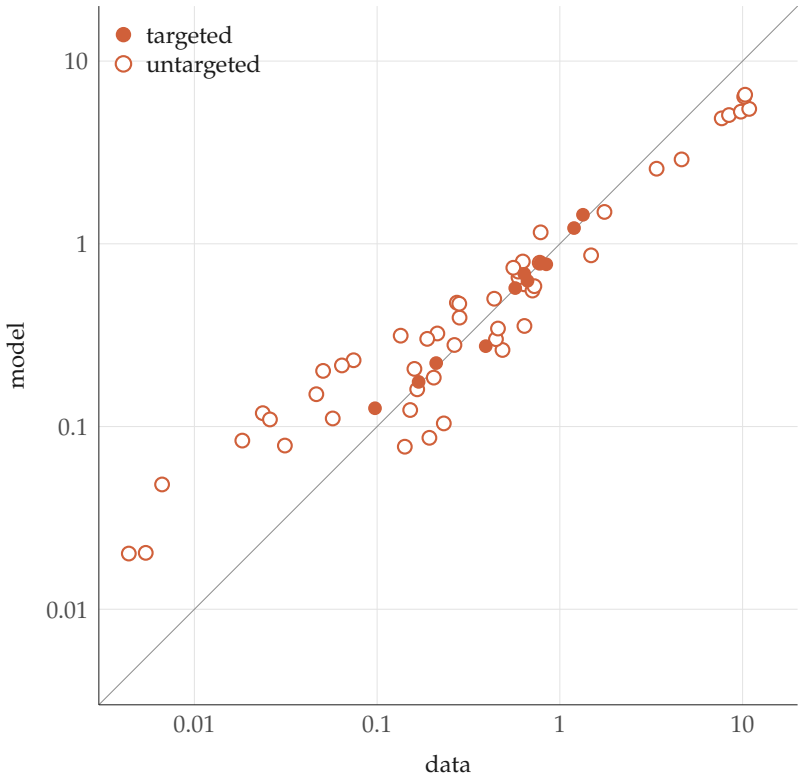
A. Estimated parameters			
Parameter	Interpretation	Estimate	Search range
$\underline{h}^{\text{own}}$	Minimum owned house size	3.14	[1, 6]
		(0.18)	
ξ_f	Fixed renovation cost	0.159	[0, 0.2]
		(0.169)	
υ	Bequest-utility scale	66.5	[5, 500]
		(3.5)	
ω_h	Owner-service size gradient	0.129	[0, 0.15]
		(0.020)	
ϱ	Retention of the efficient state	0.838	[0.8, 0.99]
		(0.151)	
η	Comfort value of an efficient home	0.061	[0, 0.15]
		(0.175)	
ρ_{ret}	Pension replacement rate	0.769	[0.3, 1.1]
		(0.024)	

B. Targeted moments		
Moment	Data	Model
Homeownership		
25–34	0.393	0.275
35–64	0.665	0.623
65–75	0.782	0.783
Owners in an A–C home, built by 1990		
25–34	0.210	0.225
35–64	0.169	0.178
65–75	0.097	0.127
Wealth, consumption and income		
Housing share of wealth, 35–64	0.570	0.569
Liquid wealth to income, 65–75	1.196	1.218
Liquid wealth to income, 75+	1.338	1.449
Consumption to income, 50–64	0.639	0.684
Retirement income, 65–74 to 55–64	0.843	0.773

Sources: Panel B: English Housing Survey (homeownership, efficient shares), Wealth and Assets Survey (wealth and income), and ONS Family Spending (consumption)
Notes: Panel A: standard errors in parentheses, from equation (4.2). They include sampling error in the data moments and simulation error in the model moments. The search range is the box over which the networks are trained and the estimator searches; it is not a confidence interval. Panel B: owners in an A–C home are shares of owner-occupiers of homes built by 1990. Model moments are simulated with 65,536 households at the estimates in Panel A.

errors of ξ_f, ϱ, and η are about as large as their search ranges. This is the weak separation we described in Section 4.3. Hence, the data pin down a combination of how long efficiency lasts and how much it is worth, but not each on its own. Another way to see this is that if ϱ were known, the standard error of η would be 0.02 instead of 0.175. Further, ϱ lies 0.04 above the lower end of its range, which is about

Figure 4: Model and data, targeted and untargeted moments



Notes: Each point is one moment with a counterpart of the same definition in the model, on log scales. Solid points are targeted and hollow points untargeted. The line is the 45-degree line. One moment, the change in liquid wealth at 65–74, is not drawn because it is zero in the model.

a quarter of a standard error. This is where the normal approximation behind the standard errors is poor. As networks are trained only within the search range, we cannot look beyond it. For these specific reasons Section 5 also reports the results at a higher retention, $\varrho = 0.96$, at which efficiency lasts 25 years.

Panel B of Table 5 compares the targeted moments with the data. The model reproduces the fall of the efficient share with age. The share falls from 22.5 percent of owners at 25–34 to 12.7 percent at 65–75. In the data these numbers are 21.0 and 9.7 percent. Hence, the model accounts for 9.8 of the 11.3-point decrease. The model matches well the ownership at 65–75 and at 35–64, but at 25–34 only 27.5 percent of households own, against 39.3 percent in the data. Further, it matches the housing share of wealth and liquid wealth at 65–75, and overstates liquid wealth at 75+ slightly. Similarly, consumption at 50–64 is slightly higher in the model, whereas the ratio of retirement income relative to income at 55–64 is 77 percent in the model and 84 percent in the data.

To assess how the model performs for non targeted moments, in Figure 4 we plot the model with the data for all 61 moments that have a counterpart with the same definition. Further, we contrast them with respect to the 45-degree line. A large fraction of moments, both targeted and not targeted are close to the 45-degree line.

5. Results

Section 2 proposed two probable explanations for why older owners live in less efficient homes. These were the horizon channel and the purchase window. In Section 3 the model was set up to include these two forces. This section asks which of the two or whether their combination produces the age gradient at the estimates of Section 4. We first show the lifecycle that the estimates imply.

Figure 5 plots the lifecycle profiles of homeownership, the share of owners in an efficient home, mean consumption and mortgage status at the estimates. Households rent when young and buy in their late twenties and thirties. Hence, the homeownership rate increases from 22 percent at 25 to 52 percent at 40 and 75 percent at 60 and remains constant at about 78 percent after that. Most of the ownership is through a mortgage. The share of households with a mortgage rises from about 20 percent to 60 percent by age 55, and decreases after that. By age 70, nearly three quarters of households own outright. Consumption has the usual hump shape. It rises with income until the late fifties and falls after retirement. The share of owners in an efficient home peaks at 23 percent around 30, falls to 18 percent at 50, and to 11 percent at 75. The share starts to fall soon after the ages at which most households buy.

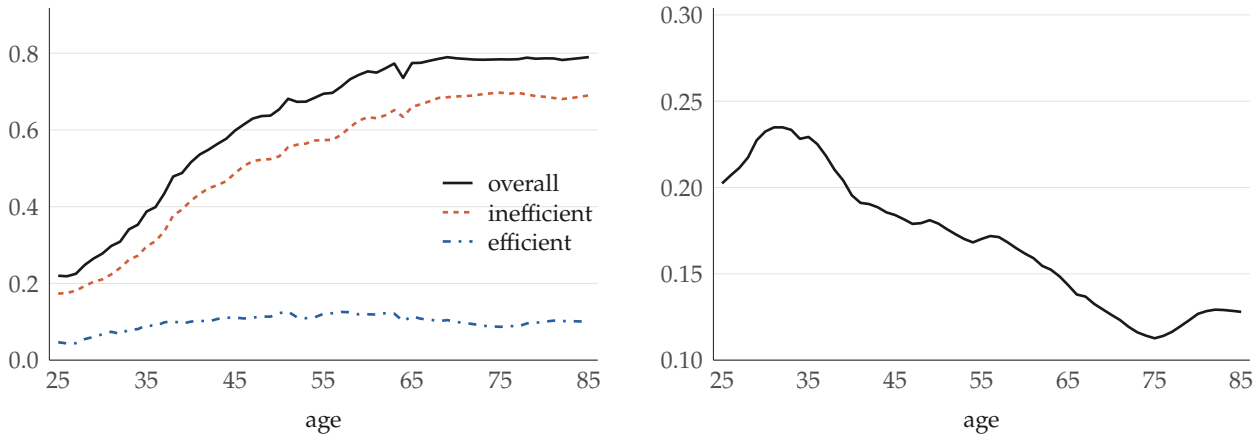
Panel (b) of Figure 5 shows how the efficient share falls with age. However, from this we cannot deduce why this share is falling. As outlined above, there are two channels. The first is the horizon channel. Under this channel an efficient home pays off only while its owner lives in it and it stays efficient. As survival probability reduces as households age, older owners have fewer expected years left to enjoy the benefit of an efficient home. The second channel is the purchase window. A buyer can choose an efficient home for a small price premium whereas an owner who stays must pay for the renovation. As the probability to buy reduces over age, this channel can lead to a declining efficient home share with age. In the rest of this section we measure the size of both these channels at estimated parameters. First, in Section 5.1 we compute the payoff horizon by age and ask whether a larger or longer-lasting payoff changes the gradient. Second, in Section 5.2 we explore the purchase window. Here we follow owners into and out of efficient homes, find the margin that produces the fall with age, and close the purchase window by taking efficient homes off the market. Finally, we ask why older owners do not renovate in place instead, and which factors determine how often they buy. We show that almost all of the gradient comes from purchases.

5.1. The horizon channel For the ease of exposition, we rewrite the expression of the payoff horizon H_j

$$H_j = \sum_{k \geq 0} \beta^k \prod_{i=0}^{k-1} \phi_{j+i} \varrho \mu_{j+i+1}$$

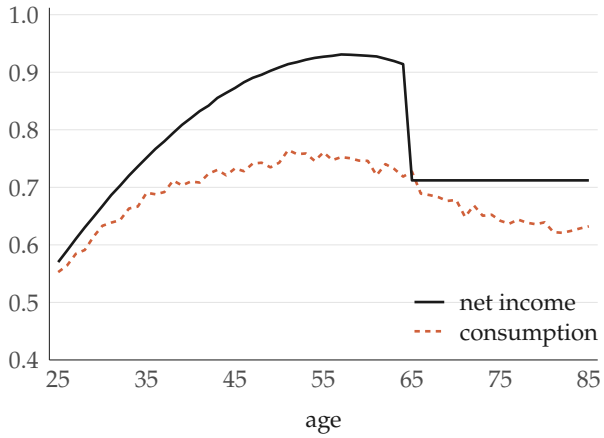
where β is the discount factor, ϕ_{j+i} the probability that an owner of age $j+i$ survives to the next year, ϱ the probability that an efficient home stays efficient, and μ_{j+i+1} the probability that an owner of age $j+i+1$ does not move. Table 6 computes the payoff horizon at different ages and illustrates the effect

Figure 5: Lifecycle profiles at the estimates

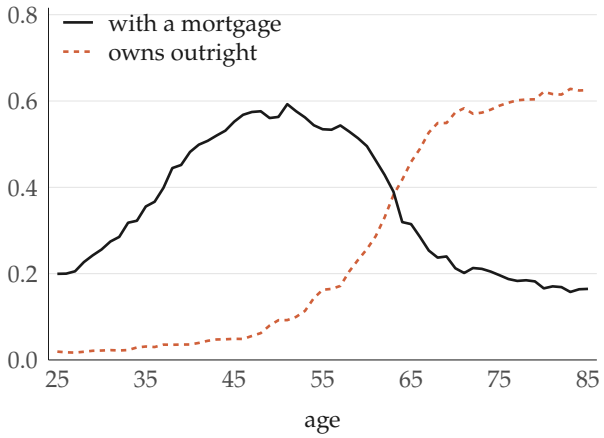


(a) Homeownership

(b) Owners in an A–C home



(c) Net income and consumption



(d) Owners by mortgage status

Notes: Simulated at the estimates of Table 5. Panel (a) plots homeownership, overall and split by whether the home is efficient (A–C) or not, as shares of all households. Panel (b) plots owners of an efficient home as a share of owners. Panel (c) plots mean income and mean consumption. Panel (d) plots owners with a mortgage and owners who own outright, as shares of all households.

of these factors by adding one factor at a time. In column (1) we keep discounting and survival only, setting ϱ and μ to one. In column (2) we add the loss of efficiency at the estimate $\varrho = 0.84$. Lastly in columns (3) and (4) we add moving, with μ_j either the model's probability that an owner of an efficient home does not move or the probability that an owner does not move in the Understanding Society survey.

With survival alone (column (1)), the horizon falls steeply with age. At age 25, the horizon is 27 years which falls to ten years at age 75. This is the intuition behind the first explanation. However, once we include the loss of efficiency, the picture changes. At $\varrho = 0.84$ efficiency lasts about six years on average, and a home that is efficient today is still efficient five years later with probability $\varrho^5 \approx 0.41$. Once this loss is added in column (2), the horizon is 5.3 years at 25, 5.3 at 45, 5.0 at 65, and 4.4 at 75. In short, the horizon is essentially flat. Adding the model's probabilities of moving (column (3)) lowers

Table 6: Payoff horizon of an efficient home, by age

Age at renovation	Discounting and survival (1)	+ decay of A–C status (2)	+ moving: model (3)	+ moving: observed rates (4)
25	26.9	5.3	3.4	3.8
35	24.8	5.3	3.4	4.2
45	22.0	5.3	3.0	4.7
55	18.6	5.2	2.8	4.7
65	14.6	5.0	3.1	4.6
75	10.0	4.4	2.9	4.2
Age 75 relative to age 25	37%	83%	86%	112%

Notes: The horizon H_j of equation (3.19), in years, for an owner of an efficient home at age j, at the estimates ($\beta = 0.97$, $\varrho = 0.84$). Each column adds one force: column (1) has discounting and survival only, for an owner who never moves in a home that never loses efficiency; column (2) adds the loss of A–C status at rate $1 - \varrho$; columns (3) and (4) add moving, at the model's probability that an owner of an A–C home moves (3) or at the probability that an owner moves in Understanding Society (4). The last row is the horizon at 75 as a percentage of the horizon at 25.

the horizon to 3.4, 3.0, 3.1, and 2.9 years. With the moving probabilities of owners in Understanding Society (column (4)), the horizon is 3.8 years at 25, 4.7 at 45, 4.6 at 65, and 4.2 at 75. As older owners move less, it offsets their higher probability of dying.[32]

The main message of Table 6 is that the horizon is essentially flat once we take decay and moving into account. This has implications for choices for which horizon matters. In Section 3.3 we show that each choice depends on a level of consumption. This is because the comfort of an efficient home is worth more to a household that consumes more. So, an owner renovates in place if its consumption is at least c_j^* (equation (3.22)), and a buyer takes an efficient home if its consumption is at least c_b^* (equation (3.24)). We call the difference between the two the threshold gap. An owner whose consumption lies in it will not renovate, but would buy an efficient home at its next purchase. By equation (3.25), the threshold gap is the extra one time cost of renovating in place spread over the horizon. If the horizon changes little with age, it implies that the threshold gap for a given home is about as wide for an owner of 75 as for an owner of 35. Also, by Proposition 2, a higher probability of moving or dying raises a buyer's threshold only in proportion to the transaction costs on the premium. Therefore, the share of buyers who take an efficient home should not fall much with age either. Hence, the two margins on which the horizon acts, what a buyer chooses and whether an owner who stays renovates, should therefore be close to flat in age.

We next change the payoff of an efficient home by varying the comfort value η and the probability

[32] One caveat with these results is that ϱ appears to be low and imprecisely estimated. The question arises: what if ϱ were set to a higher value? In Appendix E.1 we recompute the horizon at $\varrho = 0.96$, at which efficiency lasts 25 years, with the other parameters at the estimates. With loss of efficiency alone, survival now matters and the horizon falls from 14.1 years at 25 to 7.8 at 75 (column (5) of Table E4). Once we include moving it is again close to flat: 6.2 years at 25, 5.7 at 45, 6.1 at 65, and 5.4 at 75 with the model's probabilities of moving (column (6)), and 6.7, 9.8, 8.8, and 7.0 years with those in Understanding Society (column (7)). Hence, most owners move long before a rating that lasts 25 years wears off. Section 5.2 returns to this case once it has set out the accounting.

of maintaining renovation status ϱ and ask whether the age profile of the efficient share shifts or tilts. Figure 6 plots the result. For each run we keep the other parameters at their estimates and let households reoptimize. Each point in the figure plots the efficient share of owners at 25–34 against the share at 65–75. We superimpose dashed lines through the origin, along each of which the ratio of the two shares is constant. At the estimated point, the ratio is 1.75. A change that moves a point along its line changes the level of the profile, not its slope. A change that moves it toward the 45-degree line, on which young and old owners are alike, flattens the gradient.

By varying η we set how large the payoff is each year by keeping its horizon fixed. If young households choose the efficient home because they expect to enjoy it for a longer period, a larger payoff per year would be worth more to them. This would raise the share proportionally more among the young. However, it does not. Without comfort the efficient share falls to 17.2 percent at 25–34 and 9.7 percent at 65–75. With comfort doubled it rises to 27.9 and 16.6 percent. The resulting ratios are 1.76 and 1.68, against 1.75 at the estimates. Hence, the comfort points stay close to the line of the estimates. In short, a larger payoff per year actually raises the level of the profile, but does not alter its slope much.

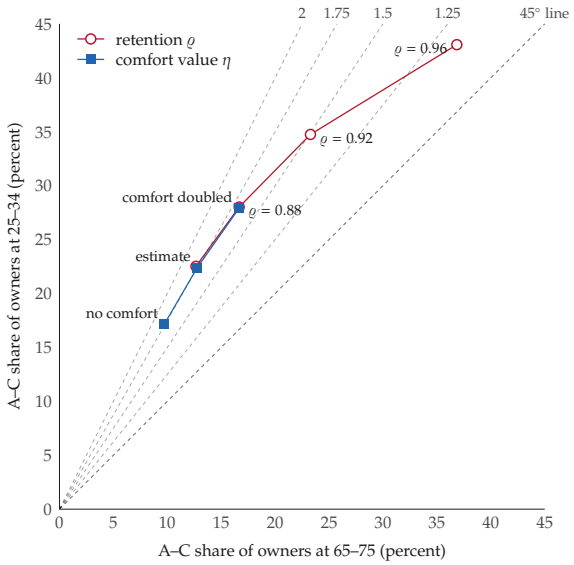
Retention ϱ sets how long the payoff lasts and it does so through two effects. First, it lengthens the horizon and this should steepen the gradient, since survival starts to matter once efficiency lasts. Second, it slows the loss of efficiency in homes bought in the past, and this can flatten the gradient. This is because older owners bought longer ago, so they lose more of their efficiency when it wears off fast. The figure shows the net of the two effects. Raising ϱ from its estimate of 0.84 to 0.88 acts like a larger payoff as it lands almost exactly on the point of doubled comfort. However, raising it further, to 0.92 and to 0.96, bends the path toward the 45-degree line, and the ratio consistently falls from 1.75 at the estimates to 1.49 and 1.17. The gradient flattens, so the second effect wins. Purchases at 65–75 do not change (Table F7 in Appendix F), so older owners do not catch up by buying more often. When efficiency lasts, they catch up because they keep more of the efficiency they bought.[33]

In conclusion, the horizon channel does not produce the gradient. Once we take into account efficiency wearing off and owners' moving probability, the horizon is about the same at every age. Figure 6 emphasizes this point. If young owners chose efficient homes because they expect to enjoy them for longer, a larger payoff each year would raise their share the most, and a payoff that lasts longer would steepen the gradient. However, neither happens. A larger payoff raises the share at every age proportionally. A payoff that lasts longer actually flattens the gradient, as older owners then keep more of the efficiency they purchased long ago. Therefore, what separates older owners from younger ones is how long ago they bought as well as how often they buy. Section 5.2 takes this up.

5.2. The purchase window Section 5.1 shows that the horizon is about the same at every age, so the thresholds on which it acts barely change with age. The choices themselves can still change, because whether a buyer takes an efficient home, or an owner who stays renovates, also depends on how

[33] Appendix E.2 separates the two effects at $\varrho = 0.96$. The share of buyers who choose an efficient home now falls with age, from 58 to 45 percent, which is the horizon at work. But older owners lose less of the efficiency they bought when young, and this outweighs it.

Figure 6: The efficient share of young and old owners when the payoff changes



Notes: Each point is a simulation that changes one parameter while keeping the others at the estimates of Table 5 and allows for households to reoptimize. A–C owners as a percentage of owners aged 25–34 (vertical) and 65–75 (horizontal). Along each dashed line through the origin the share at 25–34 is a fixed multiple of the share at 65–75, given at the end of the line; on the 45-degree line the two shares are equal. Comfort value η takes three values: none, the estimate of 0.061, and twice the estimate. Retention ϱ takes four values: the estimate of 0.84, 0.88, 0.92, and 0.96.

much the household consumes, and consumption falls after retirement (panel (c) of Figure 5). Table 7 therefore measures the two choices directly. If both are flat in age, the gradient must come from how often households buy, and the table checks this too. It follows owners from one year to the next and records how they enter and leave the stock of efficient homes, per 100 owners a year. Panel A shows that at every age the main way in is to buy an efficient home: 6.9 per 100 owners at 25–34, 4.1 at 35–64, and 1.9 at 65–75. Few buyers renovate at purchase, 0.3 to 0.4 per 100 owners. Renovating in place adds 1.2 to 1.4 per 100 owners at every age. Buying is therefore the only route that falls with age. As it falls, renovating in place becomes a larger part of a smaller inflow: 14 percent of the inflow at 25–34, 24 percent at 35–64, and 36 percent at 65–75. This is why the fixed cost of renovating, which a buyer of an efficient home avoids, weighs most on older owners (Section 4.3). On the way out, the largest flow at every age is the loss of efficiency: 4.6, 3.5, and 2.3 per 100 owners.[34] Selling to rent and selling to buy an inefficient home follow. Death matters only at 65–75, and even there it removes 0.24 per 100 owners. The net inflow is 1.5 per 100 owners at 25–34, close to zero at 35–64, and −0.5 at 65–75. The stock of efficient homes grows among young owners and shrinks among old ones, because purchases no longer replace what wears off.

To find which margin produces the fall in the inflow, we write it as

$$\text{inflow} = P\,(c_b + c_r) + N\,r,$$

[34]The loss is 18 to 21 percent of the stock in the first row, more than $1 - \varrho = 16$ percent, because it applies to the stock after the year's purchases and sales, which includes the homes bought that year.

Table 7: Flows into and out of efficient ownership, by age

	Age		
	25–34 (1)	35–64 (2)	65–75 (3)
A. Flows into and out of A–C ownership, per 100 owners a year			
A–C owners (stock)	22.52	17.72	12.66
Inflows			
Bought an A–C home	6.94	4.08	1.87
Renovated at purchase	0.28	0.42	0.29
Renovated in place	1.20	1.39	1.21
Total	8.42	5.89	3.37
Outflows			
Lost A–C status	4.63	3.46	2.34
Sold, now renting	1.35	1.16	0.59
Sold, bought a D–G home	0.88	1.13	0.72
Died	0.02	0.07	0.24
Total	6.88	5.82	3.89
Net inflow	1.54	0.07	−0.52
B. Margins of the inflow			
Purchases by renters and D–G owners, per 100 owners	14.96	10.12	4.52
Percent of those purchases			
A–C home	46.4	40.3	41.4
Renovated at purchase	1.9	4.2	6.4
In-place renovation, percent of D–G owners	1.6	1.7	1.4

C. Decomposition of the fall in the inflow to 65–75

	From 25–34	From 35–64
Fall in the inflow, per 100 owners	5.05	2.52
Percent of the fall due to		
Purchases	99	103
A–C share of purchases	10	−3
Share renovated at purchase	−9	−6
D–G owners	−3	−3
In-place renovation rate	3	10

Notes: Simulated at the estimates of Table 5. Panel A: flows per 100 owners in the age group a year. The stock is the A–C share of owners at the start of the year. Inflows are households outside the A–C stock at the start of the year. Lost A–C status applies to the stock after the year's purchases and sales, so as a share of the start-of-year stock it exceeds $1 - \varrho$. Owners who move from one A–C home to another stay in the stock and are not shown. Panel B: the inflow is $P(c_b + c_r) + Nr$, where P is the number of purchases by renters and owners of D–G homes per 100 owners, c_b and c_r are the shares of those purchases that are of A–C homes and that are renovated at purchase, N is 100 less the stock, and r is the in-place renovation rate of D–G owners. Panel C: the change in each margin is weighted by the average of the other margins in the two age groups, so the parts add up to 100.

where P is the number of purchases by renters and owners of inefficient homes per 100 owners, c_b and c_r are the fractions of those purchases that are of efficient homes and that are renovated at purchase, N is the number of owners of inefficient homes per 100 owners, and r is the fraction of them that renovate in place. Purchases by owners of efficient homes are left out, since they cannot add to the stock. Panel B shows that P falls from 15.0 at 25–34 to 10.1 at 35–64 and 4.5 at 65–75. The two choices are close to flat, so the fall in consumption with age does not undo the flat horizon. Between 40 and 46 percent of purchases are of efficient homes at every age, and between 1.4 and 1.7 percent of owners of inefficient homes renovate in place each year. The flat share of efficient purchases is the model's counterpart of Fact 3: owners who bought recently live in similarly efficient homes at every age. Panel C decomposes the fall in the inflow from 25–34 to 65–75, 5.05 per 100 owners, and from 35–64 to 65–75, 2.52 per 100 owners. It evaluates the change in each margin at the average of the other margins in the two age groups, so the parts add up exactly to the fall. Purchases account for 99 percent of the first fall and 103 percent of the second. The other margins offset one another, and together they account for 1 percent of the first fall and −3 percent of the second. Owners over 65 renovate in place a little less often than those aged 35–64, 1.4 against 1.7 percent, which accounts for a tenth of the second fall. But older buyers more often renovate at purchase, 6.4 against 4.2 percent of purchases, and with the other margins this more than offsets it.[35]

The accounting shows where the gradient appears, not why. Purchases are a choice, and a short horizon is itself a reason not to move (Section 2). We therefore change the choices themselves. The most direct test closes the window. In the first experiment of Table 8, efficient homes cannot be bought: a buyer can still buy an inefficient home and renovate it at once, and an owner can still renovate in place, and households reoptimize. The efficient share falls to 6.9, 5.5, and 5.6 percent at 25–34, 35–64, and 65–75, and the gap between the youngest and the oldest group shrinks from 9.6 points to 1.3. The payoff horizon still differs by age in this economy, but without efficient homes to buy almost no gradient remains. Almost all of the gradient needs the window, which is the causal counterpart of the accounting.

If owners who stop buying do not renovate in place instead, something must make renovating in place costly relative to buying. In equation (3.25) it is the extra one-off cost of renovating in place, and the fixed cost ξ_f is part of it. Halving ξ_f narrows the threshold gap: the efficient share rises at every age, most at 65–75 (15.0 against 12.8 percent), where more of the inflow comes from renovating in place, and the ratio of the share at 25–34 to the share at 65–75 falls from 1.75 to 1.57. Without a fixed cost the ratio is 1.36. The fixed cost governs how steeply the share falls with age, as Section 4.3 argued, but even without it the threshold gap narrows and does not close.

What, then, sets how often older households buy? The last rows of Table 8 vary the strength of the bequest motive, v, which is pinned down by the wealth of households over 65 (Section 4.3), not by any moment of efficiency. With a weak bequest motive, $v = 5$ instead of the estimate of 66.5, households aged 65–75 make 23.7 purchases a year per 100 owners instead of 6.1, and the gradient reverses: the

[35]The same conclusion holds for the stock. Rebuilding the stock year by year from these margins, with purchases held at their rate at 25–34, the efficient share at 65–75 would be 25.3 percent instead of 12.7. Holding any other margin at its rate at 25–34 changes it by at most 1.1 points.

Table 8: The efficient share of owners when the window changes

	A–C share of owners (%)			Gap, 25–34	Ratio, 25–34	Purchases,
	25–34	35–64	65–75	less 65–75	to 65–75	65–75[a]
	(1)	(2)	(3)	(4)	(5)	(6)
At the estimates	22.4	17.8	12.8	9.6	1.75	6.1
A–C homes come with a purchase						
No A–C homes for sale	6.9	5.5	5.6	1.3	1.24	4.8
Fixed cost of renovation ($\xi_f = 0.159$)						
None	28.0	25.0	20.6	7.4	1.36	7.2
Halved	23.6	19.9	15.0	8.6	1.57	6.5
Bequest motive ($\upsilon = 66.5$)						
Weaker, $\upsilon = 5$	23.3	28.6	30.8	−7.5	0.76	23.7
Stronger, $\upsilon = 200$	29.2	14.8	6.9	22.3	4.24	2.8
EHS (target)	21.0	16.9	9.7	11.3	2.16	

Notes: Each row changes one element of the model, keeps the other parameters at the estimates of Table 5 (values in parentheses), and lets households reoptimize. Shares are A–C owners as a percentage of owners in each age group, from 65,536 simulated households. The gap is the share at 25–34 less the share at 65–75, in percentage points; the ratio is the share at 25–34 over the share at 65–75. No A–C homes for sale: a buyer cannot buy an efficient home, but can still buy an inefficient home and renovate it at once, and an owner can still renovate in place. EHS: the targets of Table 5.
[a] All home purchases, by renters and owners, per 100 owners aged 65–75 a year.

efficient share is 23.3, 28.6, and 30.8 percent at 25–34, 35–64, and 65–75. With a strong one, $\upsilon = 200$, they make 2.8 purchases per 100 owners, and the gap between 25–34 and 65–75 widens to 22.3 points. The share at 65–75 moves with purchases: it rises by a factor of 2.4 when they rise by a factor of 3.9, and about halves when they halve. This is what the window predicts. It gives the bequest force of Section 3.3 a size, and it is consistent with the bequest motive shaping the gradient through how often the old buy.

The accounting and the experiment without efficient homes for sale give the same answer when efficiency lasts 25 years (Appendices E.2 and E.3). At $\varrho = 0.96$, with the other parameters at the estimates, purchases account for 90 percent of the fall in the inflow from 25–34 to 65–75 and 89 percent of the fall from 35–64, against 99 and 103 percent at the estimates (Table E5). The horizon now has a role: the share of purchases that are of efficient homes falls from 58 percent at 25–34 to 45 percent at 65–75, and this margin, on which the horizon acts, accounts for 17 and 13 percent of the two falls. Without efficient homes for sale, the gap between 25–34 and 65–75 shrinks from 6.2 points to 1.4, against 9.6 to 1.3 at the estimates (Table E6). The gap that remains without the window is about the same at both values of ϱ. A longer-lasting payoff gives the horizon a role, but the window still produces most of the gradient.

The gradient in the model is a purchase gradient. The payoff horizon is short at every age, because efficiency wears off, so it hardly differs between young and old owners, and the choices on which it acts hardly vary with age. Older owners live in less efficient homes because they buy less often, and buying is how most efficient homes are acquired. This is what Fact 3 shows in the data. The distinction matters for policy. A subsidy that lowers the cost of renovating in place narrows the threshold gap,

while a policy that acts at the point of sale works through the window. Section 6 compares the two.

6. Counterfactual Policies

Section 5 shows that older owners live in less efficient homes because they buy less often. In this section we analyze the effects of different policies on the efficient share among young and old owners, as well as overall ownership of efficient homes. We focus on three types of policies, distinguishing them by the route through which they act. The first policy focuses on lowering the cost of renovating in place, the fixed cost ξ_f and the variable cost ξ. This policy lowers the threshold c_j^* in equation (3.22), the consumption at which an owner who stays renovates. The second policy acts at a purchase. A point-of-sale standard requires every home that is bought to be made efficient. Similarly, lower costs of buying and selling, ζ_b and ζ_s, make households buy more often. The last group, higher energy prices, raises the payoff of efficiency each year without changing either cost. For each policy we report how the efficient share changes among all owners, among young and old owners, and in the gap between them. For the subsidies to renovation, we also report the public expenditure they need.

We compute the counterfactuals with the approach of Section 4.1, in which parameters are inputs of the networks. Here the inputs are six parameters that a policy can change: the fixed and variable costs of a renovation, ξ_f and ξ; the costs of buying and selling, ζ_b and ζ_s; and the maximum loan-to-value and payment-to-income ratios, λ_L and λ_D. Each ranges over an interval that contains its value in Tables 4 and 5 and every value we consider, and every other parameter is held at its value in those tables. We call this the counterfactual solution. Higher energy prices change two parameters that are not inputs, the energy bills per unit of size e^R and e^N. The point-of-sale standard changes the choices: a buyer can no longer buy an inefficient home without renovating it at once. Every simulation draws the same 65,536 households and the same income shocks, taste shocks, and draws that decide whether a home stays efficient, so two simulations differ only in the policy. We report the effect of a policy as the change, in percentage points, from the baseline of the counterfactual solution.

Before discussing the implications of these policies, we outline three caveats which apply to these experiments. First, the model is partial equilibrium. Hence, the price premium of efficient homes does not respond to a policy and a buyer can always buy an efficient home at the same premium irrespective of how many buyers want one. This matters most for a policy that raises the demand for efficient homes a great deal, such as the standard. Second, we compare two steady states, one with the policy and one without: each household lives its whole life under the policy, from age 23. We do not compute the transition due to computation constraints. Third, we do not compute welfare. The model has no cost of energy use beyond the bill a household pays, so household welfare would leave out any gain from lower emissions, and it would need an assumption about how each subsidy is financed. We report the efficient share and the flows behind it.

6.1. Lowering the cost of renovating in place A renovation costs $C(h) = \xi_f + \xi p^N h$ irrespective of whether the owner renovates at a purchase or in place. A cut in either part i.e., the fixed cost ξ_f

Table 9: Lowering the cost of renovating in place

	Variable cost halved ($\xi = 0.04$) (1)	Fixed cost halved ($\xi_f = 0.079$) (2)
Change in the A–C share of owners (percentage points)		
All owners	+6.3	+1.7
25–34	+2.5	+0.9
65–74	+8.1	+1.5
Gap, 25–34 less 65–74	−5.6	−0.6
Public outlay (% of mean income per household a year)	0.57	0.12

Notes: Each column halves one part of the cost of a renovation, $C(h) = \xi_f + \xi p^N h$, on the counterfactual solution and keeps every other parameter at its value in Tables 4 and 5. Changes from the baseline of the counterfactual solution, from 65,536 simulated households with the same draws in the baseline and in both columns. Shares are A–C owners as a percentage of owners in each age group; the gap is the share at 25–34 less the share at 65–74. Public outlay: the subsidy paid on all renovations, including those that would take place without it.

or the variable cost ξ lowers the threshold c_j^* in equation (3.22) and narrows the threshold gap in equation (3.25) which is the extra cost of renovating in place over buying an efficient home. We halve each part separately. The variable cost ξ is reduced from 8 to 4 percent of the value of the home, and the fixed cost ξ_f from 0.159 to 0.079. The variable cost is the larger part of a renovation. Renovating the smallest owned home, combining fixed and variable parts, costs about £20,500, of which £12,600 is the variable cost (Section 4.5), and the variable cost rises with the size of the home. Halving the variable cost saves £6,300 on the smallest owned home and more on a larger one. Halving the fixed cost saves £3,950 on every home, irrespective of its size.

Table 9 shows the effects of halving each cost component. Halving the variable cost raises the efficient share by 6.3 percentage points among all owners, by 2.5 points at 25–34 and by 8.1 points at 65–74. As old owners are more likely to renovate in place, the gap between young and old owners narrows by 5.6 points. The subsidy roughly triples the fraction of owners of inefficient homes who renovate in place each year. This increase occurs at both age groups. But it adds most where renovating in place supplies most of the inflow into efficient ownership. As older owners buy less often, it supplies a larger part of theirs: 36 percent at 65–75 against 14 percent at 25–34 in the estimated model (Section 5.2).

Halving the fixed cost acts in the same way but by less magnitude. This is because it saves less on each renovation. It raises the efficient share by 1.7 points among all owners, by 0.9 at 25–34 and by 1.5 at 65–74, and narrows the gap by 0.6 points (Table 9, column (2)).

To compute the public expenditure with this policy, we assume that the government pays the part of the cost that the policy removes. That means $\xi p^N h/2$ when the variable cost is halved and $\xi_f/2$ when the fixed cost is halved, where h is the size of the renovated home. We count households of every age and tenure, renters included, with each age weighted by the probability of surviving to it, and express it as a percentage of mean working-age gross income. The expenditure is therefore the

Table 10: Acting at a purchase

	Point-of-sale standard (1)	Buying and selling costs halved (2)
Change in the A–C share of owners (percentage points)		
All owners	+13.2	+5.0
25–34	+16.3	+6.6
65–74	+9.9	+4.5
Gap, 25–34 less 65–74	+6.4	+2.1
Owner moves (change per 100 owners a year)	−2.1	+10.0

Notes: Column (1) requires every home that is bought to be A–C; the model is solved again with the standard in place, on one training of the networks. Two trainings of the baseline of the counterfactual solution differ by at most 0.8 points in each share. Column (2) halves ζ_b, from 0.039 to 0.0195, and ζ_s, from 0.050 to 0.025, on the counterfactual solution, and reports the average of the effects on two trainings of its networks. Every other parameter is at its value in Tables 4 and 5. Changes from the baseline of the same training, from 65,536 simulated households with the same draws in the baseline and with the policy. Shares and the gap as in Table 9. Owner moves are sales by owners.

number of renovations per household a year times the average subsidy per renovation. Without a subsidy there are 1.0 renovations per 100 households a year. Halving the variable cost raises this to 2.7, and the average subsidy is about £10,700 for the smallest owned home. The expenditure turns out to be 0.57 percent of mean income per household a year which is about £290. Halving the fixed cost raises renovations to 1.5 per 100 households a year at £3,950 each, an outlay of 0.12 percent, which is about £60.[36]

6.2. Acting at a purchase The second group of policies acts at a purchase, which, as established in Section 5.2, is where most of the inflow into efficient ownership happens. Table 10 reports results for two such policies. The first policy is a point-of-sale standard. We impose the condition that every home that is bought must be efficient. Hence, a buyer either buys an efficient home or renovates an inefficient one at purchase. As this imposition is not determined by a parameter of the networks, we solve the model again with the standard in place. Households therefore know of this policy throughout their lives and plan for it. In the second policy we halve the costs of buying and selling, ζ_b from 3.9 to 1.95 percent of the value of the home and ζ_s from 5.0 to 2.5 percent. Table 10 reports the results.

In column (1) we report the results of the standard policy. The standard raises the efficient share by 13.2 points among all owners. Looking at age groups, it raises the share by 16.3 points at 25–34 and by 9.9 points at 65–74. Consequently, it widens the gap by 6.4 points. In the estimated model, between 40 and 46 percent of purchases are of efficient homes, and a few more are renovated at purchase. As under the standard every purchase adds to the efficient stock, the gain is largest where households buy most often. Renovation in place changes very little, since the standard does not apply to owners

[36] In the model, taxes and pensions depend only on earnings (equations (3.6) and (3.7)), so no other item of the public budget responds to the subsidy.

who stay.

It is important to note that the standard makes buying expensive for a household that would have bought an inefficient home. This is because it must pay at least $(1 + \zeta_b)\Delta p'$ more, the premium of an efficient home and the buying cost on it. Such a buyer's consumption is below the threshold c_b^* of Proposition 2, so the efficient state is worth less to it than this cost, and therefore some of these households put off moving. Consequently, owner moves reduce by 2.1 per 100 owners a year, or 27 percent, and purchases fall by 23 percent. From this perspective we observe a trade-off. Each move that does not take place is a purchase window that does not open. Hence, the fall in moves partly offsets the gain which the standard intended to generate. However, overall homeownership barely changes. It falls by 0.3 points as households know of the standard in advance and plan for it. When the standard is imposed without warning instead, so that households make their plans as if there were no standard and meet it only when they buy, ownership actually falls by 6.1 points. Note that the standard needs no public expenditure as it is a regulation and its cost falls on buyers.

Halving the costs of buying and selling makes households buy more often. Purchases rise by 118 percent and owner moves rise by 10.0 per 100 owners a year, from 7.6 (Table 10, column (2)). As households move more often, the efficient share rises by 5.0 points among all owners. The share rises by 6.6 points at 25–34 and by 4.5 at 65–74.[37] Ownership rises by 2.9 points. However, again, there is a trade-off. Owners who expect to move sooner renovate in place less, since the threshold c_j^* in equation (3.22) rises with the probability of moving. Therefore, the fraction of owners of inefficient homes who renovate in place each year falls by about a third.

6.3. Changing the payoff The policies in Sections 6.1 and 6.2 alter the cost associated with achieving the efficient state. However, a policy can instead change the payoff from owning an efficient home. In the model, an efficient home pays its owner in three ways: (i) comfort, (ii) a lower energy bill, and (iii) a higher price when it is sold. Section 5.1 varied comfort. In this section we vary the other two. The energy bill is a part of the payoff that a policy, such as a tax on household energy, can alter. A higher energy price raises every bill and, with them, the yearly saving $\Delta e = (e^N - e^R)h$. We double energy prices for both qualities, which doubles Δe. As energy prices are not a parameter of the networks, we solve the model again with doubled prices. We then double the price premium. Table 11 reports both experiments.

Doubling energy prices raises the efficient share by 5.2 points among all owners (Table 11, column (1)). Looking at age groups, the efficient share increases by 7.1 points at 25–34 and by 4.3 points at 65–74. Hence, the gap widens by 2.8 points. In proportion, however, the share rises by approximately the same amount at both ages, 40 percent at 25–34 and by 37 percent at 65–74. Like comfort in Figure 6, the bill saving sets the level and not the slope of the age profile.

The saving does not enter the two thresholds alike. A buyer who takes the efficient home pays the lower bill from its first year, and the threshold c_b^* in equation (3.24) falls by the whole saving, $\Delta e'/\theta'$. An owner who renovates now rather than next year gains only next year's saving, which it

[37] The two trainings of the networks agree on the gain among all owners, 4.9 and 5.0 points, but less on how it splits by age: the gap widens by 3.8 points on one and by 0.5 on the other.

Table 11: Changing the payoff

	Energy prices doubled ($2e^R$, $2e^N$) (1)	Price premium doubled ($\gamma_p = 0.06$) (2)
Change in the A–C share of owners (percentage points)		
All owners	+5.2	−4.7
25–34	+7.1	−8.3
65–74	+4.3	−2.7
Gap, 25–34 less 65–74	+2.8	−5.6
Routes into efficiency (percent change)		
Efficient homes bought	+28	−43
In-place renovation rate, D–G owners	+21	+41
Ownership rate (change in percentage points)	−2.9	+0.1
Owner moves (change per 100 owners a year)	+0.8	−0.1

Notes: Column (1) doubles the energy bill per unit of size of both qualities, e^R and e^N. Column (2) doubles the price premium of an efficient home, γ_p, from 0.03 to 0.06.

collects only if it survives, stays, and the home remains efficient. The threshold c_j^* in equation (3.22) falls by $\Lambda\tilde{\phi}_j\mu_{j+1}\Delta e/\theta$.[38] A larger saving therefore widens the gap between the two thresholds in equation (3.25), and buying an efficient home gains more than renovating in place. The number of efficient homes bought rises by 28 percent, and the share of purchases that are efficient rises from 38 to 48 percent, whereas the rate of renovation in place rises by 21 percent. At every age most efficient homes are acquired at a purchase, 89 percent of the inflow at 25–34 and 70 percent at 65–74. Hence the share rises in about the same proportion at both ages.

Doubling energy prices also has associated costs. Ownership falls by 2.9 points, and by between 2.5 and 4.3 points in every age group from 35 on. Owners and renters alike pay the bill on every unit of size, so housing becomes dearer for both. Renters rent smaller homes, as the median rent falls by 9 percent, and buyers buy smaller ones, as the median value of an owned home relative to income falls by 5 percent at 35–64. The fall in ownership is consistent with the floor on the size of an owned home. An owned home cannot be smaller than $\underline{h}^{\text{own}}$, 3.14 times mean income, whereas a rented home can. A household that now wants less housing than the smallest owned home therefore rents instead. Fewer renters buy: the share of renters aged 35–64 who own two years later falls from 17.4 to 15.5 percent, and slightly more owners sell and rent. Owners also move more often, by 0.8 per 100 owners a year, and at least 70 percent of the extra moves end in another purchase.

The third part of the payoff is the price premium, which an owner collects when it sells an efficient home and its heirs collect if it dies. A policy that makes buyers value the certificate more, for example by making the rating more visible when a home is sold, raises the premium γ_p. We double it, from 3 to 6 percent, and, as with energy prices, solve the model again. The buyers of efficient homes pay

[38] The owner pays the bill of the inefficient home in the year of the work, so $\overline{C}$ and ℓ both rise one for one with Δe. Differentiating equation (3.22) then gives $\partial c_j^*/\partial\Delta e = -\Lambda\tilde{\phi}_j\mu_{j+1}/\theta$, and differentiating equation (3.24) gives $\partial c_b^*/\partial\Delta e' = -1/\theta'$.

the premium to their sellers, so it needs no public outlay. Doubling the premium lowers the efficient share by 4.7 points among all owners (column (2)), by 8.3 points at 25–34 and by 2.7 points at 65–74, and the gap narrows by 5.6 points, almost nine-tenths of it. In proportion, the share falls by 46 percent at 25–34 and by 23 percent at 65–74. Unlike the bill saving, the premium changes the slope of the age profile. It barely changes who owns or how often owners move.

The premium enters the two routes to efficiency with opposite signs. A buyer pays it. Doubling it doubles the cost of buying the efficient state, $C_b = (1 + \zeta_b)\Delta p'$, and so lowers the extra one-off cost of renovating in place, $C(h) - C_b + \Delta e$, which sets the gap between the two thresholds in equation (3.25). Renovating the smallest owned home now costs 2.1 times the extra outlay of buying it efficient, rather than 4.2 times, and renovating the largest 1.4 times rather than 2.7. An owner who renovates, by contrast, gets more of its cost back, since the premium is what a sale returns and what the heirs receive: a larger Δp shrinks both the loss at exit and the loss at death in equation (3.21). The routes in column (2) show both responses. The number of efficient homes bought falls by 43 percent, and the share of purchases that are of efficient homes falls from 38 to 22 percent. The fraction of owners of inefficient homes who renovate in place each year rises by 41 percent, and the share of purchases renovated at purchase rises from 4.4 to 6.4 percent.

The fall in efficient purchases outweighs the rise in renovation, and most so among the young, who acquire efficiency mainly by buying it. At 25–34, homes bought efficient make up 84 percent of the inflow into efficient homes, against 59 percent at 65–74. At 25–34 the number of efficient homes bought falls by 52 percent, whereas the in-place rate rises by 31 percent. Older owners respond more on the other margin: at 65–74 the in-place rate rises by 61 percent, and the inflow into efficient homes rises by 4 percent. Their share nevertheless falls, because they reach 65 with fewer efficient homes. The share at 55–64 falls by 5.0 points, and efficiency lasts about six years. After 75 the share does not change by more than the variation between two trainings of the networks.

The experiment bears on a natural conjecture, that the gradient reflects a small premium. When the market pays little for efficiency, an owner who expects to leave the home soon, by moving or dying, gets little of the cost of a renovation back. A larger premium would then bring older owners into efficient homes and close the gap from below. In the model, a larger premium does raise renovation, most among older owners, but it lowers their efficient share, and the gap closes from above, through what the young buy. The premium therefore acts on the gradient through the purchase window of Section 5.2: the larger the premium, the smaller the advantage of buying efficiency over renovating, and the young use that route most. For the same reason, the size of the gradient depends on the premium measured in the data. Because the premium is fixed in the model, the experiment compares two markets that price efficiency differently; the supply of efficient homes for sale does not respond.

7. Conclusion

In England, older owners live in less energy-efficient homes than younger owners. This paper asks why. While keeping the preferences of older and younger owners the same, it considers two channels

through which age could affect investment in efficiency. The first is the horizon channel: an older owner has fewer years to collect the return on a renovation. The second is the purchase window: a buyer gets efficiency with the home for a small premium and recovers it when she sells, but this route is open only to owners who move.

By solving and estimating a lifecycle model on English data we find that the age gradient comes from the purchase window, not from the horizon channel. At the estimated parameters, the horizon is short at every age, because efficiency wears off long before most owners die. Owners of every age choose alike when they buy, and few renovate in place at any age. However, owners buy less often as they age. When buying an efficient home is ruled out, the gap in the efficient share between the youngest and the oldest owners falls from about ten percentage points to about one. The result holds even when efficiency lasts 25 years instead of six. The horizon then falls with age, but purchases still account for most of the decline.

The age gradient is therefore a purchase gradient. There are two implications associated with this finding. First, longer lives do little to raise investment in efficiency. As the population ages, a growing share of homes is held by owners who rarely buy, and these homes are seldom upgraded. Second, the route through which a policy acts decides whom it reaches. For example, a subsidy for renovation reaches older owners, because they are primarily the ones who rely on renovating in place. This narrows the age gap. A standard imposed at the point of sale raises efficiency the most, but more so among the younger households, and steepens the age gradient. Also, by making a purchase dearer it slows the turnover through which efficiency enters the stock. Reducing the transaction costs associated with buying and selling opens more purchase opportunities but weakens the incentive to renovate in place. A higher price of energy raises efficiency at every age in about the same proportion and leaves the gradient in place, while a larger price premium on efficient homes closes the age gap by making the homes of young owners less efficient, not those of older owners more efficient.

Although we study energy efficiency, the mechanism is not specific to energy. Whenever the quality of a durable is cheapest to change at the moment it is bought, the quality a household holds depends on how often it transacts. Hence, households that rarely transact hold quality chosen long ago. Homes are rarely replaced but often resold, and building codes reach them only when they are new, so for the existing stock the sale is the main moment at which efficiency improves.

References

Adda, J. and R. Cooper (2000). Balladurette and juppette: A discrete analysis of scrapping subsidies. *Journal of Political Economy 108*(4), 778–806.

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

Attanasio, O. P. (2000). Consumer durables and inertial behaviour: Estimation and aggregation of (S,s) rules for automobile purchases. *Review of Economic Studies 67*(4), 667–696.

Attanasio, O. P., R. Bottazzi, H. W. Low, L. Nesheim, and M. Wakefield (2012). Modelling the demand for housing over the life cycle. *Review of Economic Dynamics 15*(1), 1–18.

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

Azinovic, M., L. Gaegauf, and S. Scheidegger (2022). Deep equilibrium nets. *International Economic Review 63*, 1471–1525.

Bajari, P., P. Chan, D. Krueger, and D. Miller (2013). A dynamic model of housing demand: Estimation and policy implications. *International Economic Review 54*(2), 409–442.

Bank of England (2017, November). Financial stability report, November 2017. Technical report, Bank of England.

Banks, J., R. Blundell, Z. Oldfield, and J. P. Smith (2012). Housing mobility and downsizing at older ages in Britain and the United States. *Economica 79*(313), 1–26.

Battu, H., A. Ma, and E. Phimister (2008). Housing tenure, job mobility and unemployment in the UK. *Economic Journal 118*(527), 311–328.

Becker, G. S. (1962). Investment in human capital: A theoretical analysis. *Journal of Political Economy 70*(5, Part 2), 9–49.

Ben-Porath, Y. (1967). The production of human capital and the life cycle of earnings. *Journal of Political Economy 75*(4, Part 1), 352–365.

Bento, A. M., S. Li, and K. Roth (2012). Is there an energy paradox in fuel economy? a note on the role of consumer heterogeneity and sorting bias. *Economics Letters 115*(1), 44–48.

Best, M. C. and H. J. Kleven (2018). Housing market responses to transaction taxes: Evidence from notches and stimulus in the UK. *Review of Economic Studies 85*(1), 157–193.

Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *American Economic Review 98*(5), 1887–1921.

BRE (2014). The government's standard assessment procedure for energy rating of dwellings, 2012 edition (version 9.92). https://www.bregroup.com/sap/standard-assessment-procedure-sap-2012. Published on behalf of DECC by BRE, Garston, Watford.

Carro, A., M. Hinterschweiger, A. Uluc, and J. D. Farmer (2023). Heterogeneous effects and spillovers of macroprudential policy in an agent-based model of the UK housing market. *Industrial and Corporate Change 32*, 386–432.

Carroll, C. D. and A. A. Samwick (1997). The nature of precautionary wealth. *Journal of monetary Economics 40*(1), 41–71.

Cervellati, M. and U. Sunde (2013). Life expectancy, schooling, and lifetime labor supply: Theory and evidence revisited. *Econometrica 81*(5), 2055–2086.

Cocco, J. F. (2005). Portfolio choice in the presence of housing. *Review of Financial Studies 18*(2), 535–567.

Collins, M. and J. Curtis (2018). Bunching of residential building energy performance certificates at threshold values. *Applied Energy 211*, 662–676.

Crawley, J., P. Biddulph, P. J. Northrop, J. Wingfield, T. Oreszczyn, and C. Elwell (2019). Quantifying the measurement error on England and Wales EPC ratings. *Energies 12*(18), 3523.

Davidoff, T. (2010). Home equity commitment and long-term care insurance demand. *Journal of Public Economics 94*(1–2), 44–49.

De Nardi, M. (2004). Wealth inequality and intergenerational links. *Review of Economic Studies 71*(3), 743–768.

De Nardi, M., E. French, and J. B. Jones (2010). Why do the elderly save? The role of medical expenses. *Journal of Political Economy 118*(1), 39–75.

Dechow, P. M. and R. G. Sloan (1991). Executive incentives and the horizon problem: An empirical investigation. *Journal of Accounting and Economics 14*(1), 51–89.

Dröes, M. I. and Y. van der Straten (2024). Home improvement, wealth inequality, and the energy-efficiency paradox. Discussion Paper 24-026/IV, Tinbergen Institute.

Druedahl, J., R. Huleux, and J. Røpke (2026). Deep learning solutions of large non-convex life-cycle models. Working paper.

Druedahl, J. and J. Røpke (2026). Deep learning algorithms for solving convex finite-horizon models. Working paper.

Eberly, J. C. (1994). Adjustment of consumers' durables stocks: Evidence from automobile purchases. *Journal of Political Economy 102*(3), 403–436.

Ferreira, F., J. Gyourko, and J. Tracy (2010). Housing busts and household mobility. *Journal of Urban Economics 68*(1), 34–45.

Fleischer, C., C. Hambel, and M. Koch (2026). Energy-efficiency investments and consumption-portfolio choice. Academic Paper AP 2026-11, Netspar.

Fonseca, J. and L. Liu (2024). Mortgage lock-in, mobility, and labor reallocation. *Journal of Finance 79*(6), 3729–3772.

Fortson, J. G. (2011). Mortality risk and human capital investment: The impact of HIV/AIDS in sub-Saharan Africa. *The Review of Economics and Statistics 93*(1), 1–15.

Fowlie, M., M. Greenstone, and C. Wolfram (2018). Do energy efficiency investments deliver? Evidence from the Weatherization Assistance Program. *Quarterly Journal of Economics 133*(3), 1597–1644.

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

Gerarden, T. D., R. G. Newell, and R. N. Stavins (2017). Assessing the energy-efficiency gap. *Journal of Economic Literature 55*(4), 1486–1525.

Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the life cycle. *Econometrica 70*(1), 47–89.

Grossman, S. J. and G. Laroque (1990). Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods. *Econometrica 58*(1), 25–51.

Gruenspecht, H. K. (1982). Differentiated regulation: The case of auto emissions standards. *American Economic Review 72*(2), 328–331.

Hausman, J. A. (1979). Individual discount rates and the purchase and utilization of energy-using durables. *The Bell Journal of Economics 10*(1), 33–54.

Hazan, M. (2009). Longevity and lifetime labor supply: Evidence and implications. *Econometrica 77*(6), 1829–1863.

Head, A. and H. Lloyd-Ellis (2012). Housing liquidity, mobility, and the labour market. *Review of Economic Studies 79*(4), 1559–1589.

Heathcote, J., K. Storesletten, and G. L. Violante (2017). Optimal tax progressivity: An analytical framework. *Quarterly Journal of Economics 132*(4), 1693–1754.

Hilber, C. A. L. and T. Lyytikäinen (2017). Transfer taxes and household mobility: Distortion on the housing or labor market? *Journal of Urban Economics 101*, 57–73.

HM Government (2017). The clean growth strategy: Leading the way to a low carbon future. Technical report, Department for Business, Energy and Industrial Strategy, London.

Iskhakov, F., T. H. Jørgensen, J. Rust, and B. Schjerning (2017). The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks. *Quantitative Economics 8*(2), 317–365.

Jacobsen, G. D. and M. J. Kotchen (2013). Are building codes effective at saving energy? Evidence from residential billing data in Florida. *Review of Economics and Statistics 95*(1), 34–49.

Jacobsen, M. R. and A. A. van Benthem (2015). Vehicle scrappage and gasoline policy. *American Economic Review 105*(3), 1312–1338.

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

Jayachandran, S. and A. Lleras-Muney (2009). Life expectancy and human capital investments: Evidence from maternal mortality declines. *The Quarterly Journal of Economics 124*(1), 349–397.

Kaplan, G., K. Mitman, and G. L. Violante (2020). The housing boom and bust: Model meets evidence. *Journal of Political Economy 128*(9), 3285–3345.

Kinnerud, K. (2025). The effects of monetary policy through housing and mortgage choices on aggregate demand. *Quantitative Economics 16*(2), 659–703.

Levinson, A. (2016). How much energy do building energy codes save? Evidence from California houses. *American Economic Review 106*(10), 2867–2894.

Liu, L. (2023, January). The demand for long-term mortgage contracts and the role of collateral. Staff Working Paper 1009, Bank of England.

Maliar, L., S. Maliar, and P. Winant (2021). Deep learning for solving dynamic economic models. *Journal of Monetary Economics 122*, 76–101.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers in Econometrics*, pp. 105–142. New York: Academic Press.

Ministry of Housing, Communities and Local Government (2024). English housing survey 2023 to 2024: Headline findings on demographics and household resilience. Technical report, Ministry of Housing, Communities and Local Government, London.

Ministry of Housing, Communities and Local Government (2025). Housing supply: Net additional dwellings, England: 2024 to 2025. Technical report, Ministry of Housing, Communities and Local Government, London.

Montizaan, R., F. Cörvers, and A. De Grip (2010). The effects of pension rights and retirement age on training participation: Evidence from a natural experiment. *Labour Economics 17*(1), 240–247.

Myers, E. (2019). Are home buyers inattentive? Evidence from capitalization of energy costs. *American Economic Journal: Economic Policy 11*(2), 165–188.

Nakajima, M. and I. A. Telyukova (2017). Reverse mortgage loans: A quantitative analysis. *Journal of Finance 72*(2), 911–950.

Nakajima, M. and I. A. Telyukova (2020). Home equity in retirement. *International Economic Review 61*(2), 573–616.

Newell, R. G. and J. Siikamäki (2015). Individual time preferences and energy efficiency. *American Economic Review 105*(5), 196–200.

Office for National Statistics (2025). Household projections for England: 2022-based. Technical report, Office for National Statistics, Newport.

Oster, E., I. Shoulson, and E. R. Dorsey (2013). Limited life expectancy, human capital and health investments. *American Economic Review 103*(5), 1977–2002.

Piazzesi, M. and M. Schneider (2016). Housing and macroeconomics. In J. B. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Volume 2, pp. 1547–1640. Elsevier.

Quigley, J. M. (1987). Interest rate variations, mortgage prepayments and household mobility. *Review of Economics and Statistics 69*(4), 636–643.

Reusens, P., T. van Kempen, J. Vandenbergh, F. Vastmans, and S. Damen (2026). Impact on house prices of an energy renovation obligation for homebuyers. *Energy Economics 158*, 109317.

Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica 55*(5), 999–1033.

Stavins, R. N. (2006). Vintage-differentiated environmental regulation. *Stanford Environmental Law Journal 25*(1), 29–63.

Stolyarov, D. (2002). Turnover of used durables in a stationary equilibrium: Are older goods traded more? *Journal of Political Economy 110*(6), 1390–1413.

van Ommeren, J. and M. van Leuvensteijn (2005). New evidence of the effect of transaction costs on residential mobility. *Journal of Regional Science 45*(4), 681–702.

Yang, F. (2009). Consumption over the life cycle: How different is housing? *Review of Economic Dynamics 12*(3), 423–443.

A. Data Construction

This appendix documents the samples, definitions, and calculations behind the facts in Section 2 and the estimation targets in Section 4, in the order in which the main text uses them. Subsection A.1 describes the English Housing Survey files and definitions, Subsection A.2 the age-gradient regressions, Subsection A.3 the EPC register, Subsections A.4 and A.5 the Wealth and Assets Survey and Understanding Society, Subsection A.6 mortality and energy bills, and Subsection A.7 the samples behind each estimation target.

A.1. English Housing Survey: files and definitions.

Files. The EHS is released in two forms. A housing-stock release pools two consecutive years of fieldwork and contains the households whose dwelling received a physical inspection, together with their interview records and a weight for the pooled sample. A fieldwork release contains the full interview sample of a single year, without physical data. Table A1 lists the releases we use. Statistics that need the physical survey (the EPC band, the SAP rating, and dwelling age) use housing-stock releases only.

Table A1: English Housing Survey releases

				Used for		
Release	File	UKDA SN	Fieldwork years	Homeownership	A–C share	Regression
Housing stock 2011–12	11plus12	7511	2011–12, 2012–13	✓	✓	✓
Housing stock 2012–13	12and13	7802	2012–13, 2013–14	✓	✓	✓
Fieldwork 2013–14	fs13	7801	2013–14	✓		
Housing stock 2013–14	13plus14	8010	2013–14, 2014–15			✓
Fieldwork 2014–15	fs14	8009	2014–15	✓		
Housing stock 2014–15	14plus15	8186	2014–15, 2015–16	✓	✓	✓
Housing stock 2015–16	15plus16	8350	2015–16, 2016–17			✓
Housing stock 2016–17	16_17	8494	2016–17, 2017–18	✓	✓	✓
Housing stock 2019–20	19plus20	9058	2019–20, 2020–21	✓	✓	✓

Notes: "Homeownership" and "A–C share" are the estimation targets (Table A2); "Regression" is the age-gradient regression of Figure 1. Fieldwork years run from April to March.

Definitions. Age is the age of the household reference person (HRP). All releases report it in bands (25–34, 35–44, 45–54, 55–64, and 65 and over); exact age is available only in 11plus12, 12and13, and fs13, so statistics for ages 65–75 (inclusive) use those releases. A household is an owner if it owns outright or with a mortgage. A dwelling is built by 1990 if the EHS dwelling-age variable places it in one of the bands before 1919, 1919–44, 1945–64, 1965–80, or 1981–90. A dwelling is in A–C if its EPC band is A, B, or C; the EHS reports A and B as one band. The band is based on SAP 2009 in 11plus12 and on SAP 2012 in the later releases; the earlier release has only partial SAP 2012 recalculations, so we use its published SAP 2009 band. Years since purchase are the completed years the household has owned the home: exact in 11plus12 and 12and13, and in bands (less than one, one, two, 3–4, 5–9,

10–19, 20–29, and 30 or more years) in later releases. A recent buyer has owned the home for at most four completed years.

Weights and pooling. All EHS statistics use the household weight supplied with each release. When a statistic pools releases, we first normalize the weights so that they sum to one within each release and then average across releases. Each release therefore receives equal weight, whatever its sample size.

Overlap between releases. Because each housing-stock release pools two fieldwork years, adjacent releases contain the same households. In our files, 11plus12 and 12and13 share 6,058 households (the 2012–13 fieldwork year). The fieldwork release fs13 repeats 2013–14 households that also appear in 12and13, and fs14 repeats 2014–15 households that also appear in 14plus15. Pooled point estimates count these households twice. Our bootstrap standard errors (Section 4) draw each of the 6,058 households shared by 11plus12 and 12and13 once. Household identifiers in releases after 2013 are anonymized separately in each file, so the other overlaps cannot be linked, and the standard errors of the homeownership targets are somewhat understated.

A.2. The age gradient in the SAP rating. Figure 1 reports estimates of

$$\log SAP_i = \alpha + \sum_{j=2}^{5} \beta_j \, \mathbf{1}\{\text{age group}_i = j\} + \gamma' X_i + \delta_{t(i)} + \eta_{r(i)} + \varepsilon_i, \tag{A.1}$$

where SAP_i is the SAP rating of household i's dwelling, our measure of energy efficiency; a higher rating means a more efficient dwelling. The age groups are 35–44, 45–54, 55–64, and 65 and over, with 25–34 as the reference group. X_i contains log floor area, log equivalized household income before housing costs (deflated by CPIH, 2015 prices), and indicators for a single-person household, dwelling age, household size (top-coded at five), and the HRP's labor-force status. δ_t and η_r are release and region (nine regions) fixed effects. The three specifications in the figure are: age indicators only; adding X_i; and adding X_i and the fixed effects. The figure plots β_j; the text reports the percentage difference $\exp(\beta_j) - 1$. Standard errors are clustered by dwelling age and region (81 clusters), and the intervals use the t distribution.

The sample stacks the seven housing-stock releases marked "Regression" in Table A1, keeps owner-occupiers with an HRP aged 25 or over, and contains 34,767 observations. A version that adds private renters has 48,705 observations. The SAP rating is based on SAP 2009 in 11plus12 and on SAP 2012 in later releases. The regressions are unweighted. For owners aged 65 and over, the SAP rating is 9.0 percent lower than at 25–34 without controls, 6.5 percent lower with X_i, and 6.6 percent lower with X_i and the fixed effects.

A.3. EPC register.

Data. We use the bulk download of all domestic certificates in England and Wales from the Energy Performance of Buildings register, with one file per local authority (343 authorities). Lodgement dates run from 1 October 2008 to 31 January 2023. The register contains 24,133,533 certificates and no duplicate certificate numbers.

Each certificate reports the dwelling's unique property reference number (UPRN), the lodgement date, the current EPC band and SAP score, the construction-age band, property type, floor area, modelled energy use, the reason for the assessment (for example, a marketed sale or a new letting), and the tenure stated by the assessor. A separate file lists recommended improvements with indicative costs. We classify tenure from its text as owner-occupied (12,153,496 certificates), rented (9,255,763), or other or unknown (2,724,274). Construction-age bands are mapped to their final year. Homes are built by 1990 if that year is 1990 or earlier, which includes the bands "before 1900" and "1983–1990".

Consecutive pairs and annual rates (Table 2). We keep owner-occupied certificates with a UPRN (12,006,376 certificates). Of these, 425,058 have a construction-age band with no year (mostly recorded as missing or invalid) and so have no vintage. We order them by lodgement year within each UPRN, and pair each certificate with the preceding one. Certificates of other tenures are removed before pairing, so a pair can span an intervening letting. We keep pairs whose lodgement years differ by at least one, leaving 1,705,904 pairs; a pair lodged within a single calendar year is dropped and not linked further back. A pair takes the vintage of its later certificate. For a pair k with gap $g_k \geq 1$ in calendar years, the annual upgrade rate is the mean over pairs whose earlier certificate is in D–G of $\mathbf{1}\{\text{later certificate in A–C}\}/g_k$. The annual decay rate is the mean over pairs whose earlier certificate is in A–C of $\mathbf{1}\{\text{later certificate in D–G}\}/g_k$. Dividing by the gap treats the probability of a change over g years as g times an annual probability. This holds approximately when changes are rare, and it misses changes that are reversed between two certificates. Among homes built by 1990, the upgrade rate is 6.9 percent (1,295,552 pairs) and the decay rate 15.2 percent (171,041 pairs). The mean gap is 5.2 years. For homes built after 1990 the rates are 11.9 percent (104,678) and 5.8 percent (123,400).

Pairs around a sale. The split by sale timing in Section 2.2.2 classifies pairs by the reason for assessment recorded on each certificate, not by a match to transaction records. A pair is *at sale* if its later certificate was lodged for a marketed or non-marketed sale. It is *after a purchase* if its earlier certificate was lodged for a sale, its later certificate was not, and the two are at most two calendar years apart. All other pairs are *away from a sale*. This split is computed by a separate script that reads the raw register and orders certificates within a year by their lodgement time. Its pooled rates are 7.3 percent for upgrades and 11.2 percent for decay, against 7.4 and 11.1 percent in Table 2. For homes built by 1990, the upgrade rate is 7.3 percent at sale (835,104 pairs), 6.6 percent after a purchase (133,087), and 5.5 percent away from a sale (326,260). The decay rate is 10.7 percent at sale (124,902), 49.0 percent after a purchase (10,852), and 20.9 percent away from a sale (36,616). Pairs after a purchase are short (1.4 years on average) and few, and a sale certificate followed by a different rating within two years points to measurement error at either assessment as much as to physical change.

First and latest certificates. Figures A1 and A2 compare, for each owner-occupied dwelling with at least two certificates at least one calendar year apart (1,557,929 dwellings), the first and the latest certificate. The mean gap is 5.9 years. Among homes built by 1990 that start in D–G (1,178,403), 22.9 percent are in A–C at their latest certificate. The share is 29.1 percent for homes starting in D and between 15 and 18 percent for homes starting in E, F, or G. Of the moves from D–G into A–C, 95 percent end in band C.

Selection. The register is not a random sample of the housing stock. A certificate is lodged when a home is sold, let, or built, or when an energy-efficiency scheme requires an assessment, so homes that change hands often are over-represented, and the timing of a certificate is itself a choice. Assessment conventions and default inputs change over time and differ across assessors, and they can move a home across the C/D boundary without physical work. We therefore use the register to describe the direction and rough size of changes in efficiency status. The model is not estimated on it.

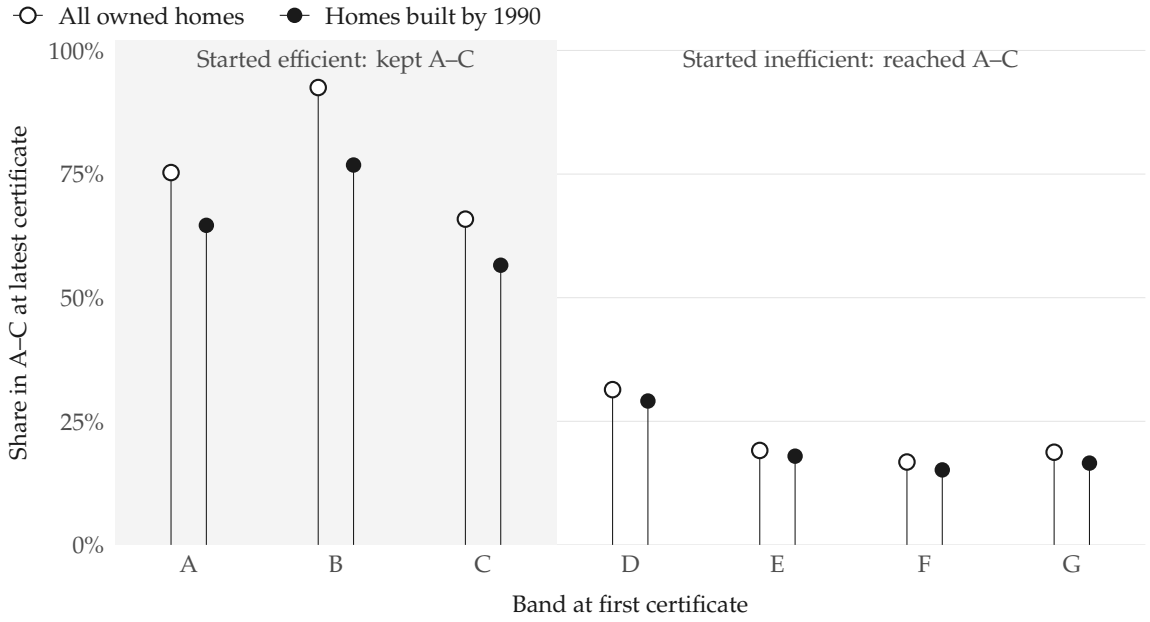


Figure A1: Efficiency at the latest certificate by band at the first certificate, owner-occupied homes
Notes: For each EPC band at a dwelling's first owner-occupied certificate, the share of dwellings in bands A–C at the latest owner-occupied certificate. Dwellings with at least two such certificates at least one calendar year apart, 1,557,929 dwellings (mean gap 5.9 years). Left: all homes. Right: homes built by 1990, based on the latest certificate. EPC register, England and Wales, October 2008 to January 2023.

A.4. Wealth and Assets Survey. We use the household files of WAS rounds 5 to 8, which cover Great Britain between April 2014 and March 2022. The sample keeps households with a positive cross-sectional weight and drops shared owners. Age is the HRP's age band; WAS does not release exact age. Pooled statistics give each round equal weight. Negative asset values are set to missing.

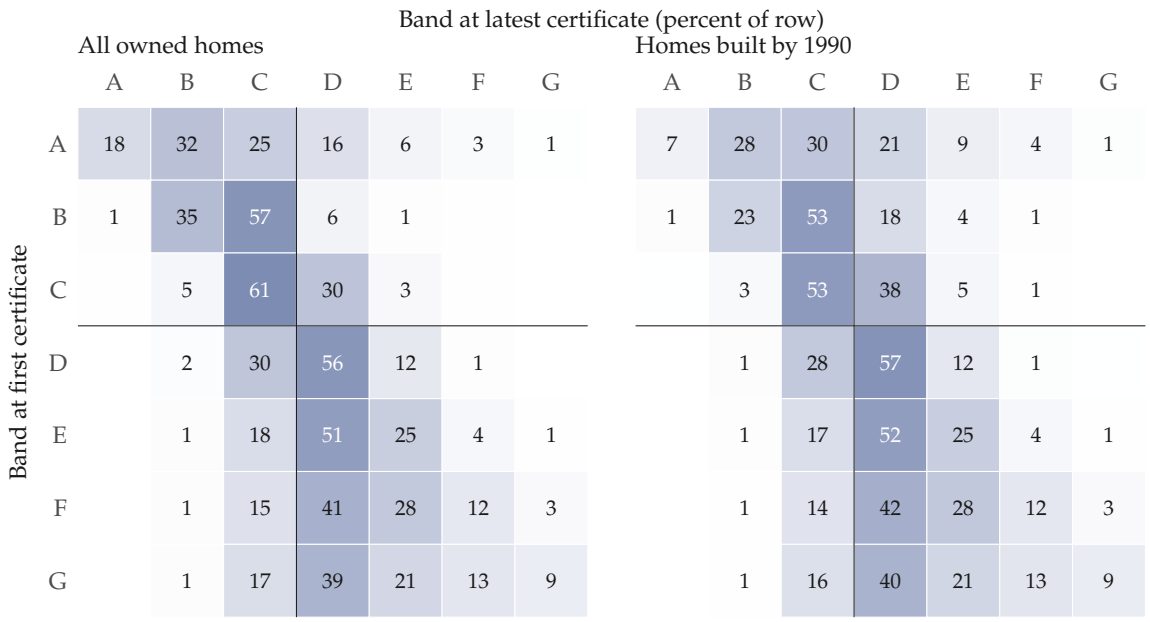


Figure A2: Transitions between EPC bands, first to latest certificate, owner-occupied homes

Notes: Row percentages: the distribution of the band at the latest certificate for each band at the first certificate. Lines mark the boundary between bands C and D. Sample as in Figure A1.

Housing share of wealth. Housing wealth is the value of the main residence for owners and zero for others. The target is the weighted mean, over households aged 35–64 with positive gross wealth, of housing wealth divided by the sum of housing wealth and gross non-housing financial wealth excluding endowments. Both numerator and denominator are gross of mortgage debt. Renters with financial wealth enter with a share of zero.

Liquid wealth to income. Liquid wealth is gross non-housing financial wealth excluding endowments, the closest harmonized household measure in WAS. It includes all financial assets, not only cash and deposits, and is not net of financial liabilities. Income is net regular household income before housing costs: `DVTotNIRR5` in round 5 and `DVTotInc_BHC` in rounds 6–8. The targets are weighted medians of the household-level ratio for HRPs aged 65–74 and 75 and over, among households with positive income. The model counterparts are ages 65–75 and 75 to the terminal age. Medians of household ratios differ from ratios of medians; the corresponding median incomes are £24,900 at 65–74 and £18,200 at 75 and over, and median liquid wealth is £29,936 and £23,499.

Consumption and retirement income. The consumption target uses published expenditure rather than microdata. Its numerator is mean annual household expenditure for an HRP aged 50–64 from ONS Family Spending (Table A11, United Kingdom, financial years ending 2016 and 2017). We count COICOP divisions 1–12 and spending on holidays abroad, but exclude housing (division 4, apart from water and other dwelling services) and moving costs; the model treats rent, energy, and maintenance separately. The weekly figures are multiplied by 52 and expressed in 2016 pounds with the CPI. The

denominator is the weighted mean net household income in WAS rounds 5 and 6 of households interviewed in the same two financial years, with an HRP aged 50–64 (2,628 and 2,721 households). The target is the ratio of the two equal-year means, 0.639. The retirement-income target is the ratio of weighted mean net household income at HRP ages 65–74 to that at 55–64 in the same WAS sample, 0.843; it compares households, not a pension replacement rate for individuals. Both targets use exact five-year HRP age bands from the person file. The numerator covers the United Kingdom and the denominator Great Britain.

Equivalence scale. The household-equivalence scale e_j in the model is the square root of mean household size by HRP age in WAS rounds 5–8. Means by age band are fitted with a cubic in age over ages 23–82, and the scale is rescaled so that its mean over ages 35–64 matches a reference value.

A.5. Understanding Society.

Income process. The income process is estimated from waves 1–10 (2009–2019) of Understanding Society, restricted to England. Interviews from wave 10 conducted in 2020 are excluded. The sample consists of male household reference persons aged 20–64 who work 10–112 hours a week and are not retired, students, caring for family, or long-term sick. It further requires positive gross household income, hourly household earnings between £2 and £400 in 2015 prices, and at least five observations per person. This leaves 32,698 person-years. Income is gross monthly household income times twelve, deflated to 2015 prices. The regressions are unweighted. We regress log income on age indicators, marital status, the number of children, and education with person fixed effects (31,711 observations, 4,265 persons). The deterministic age profile is a cubic in age fitted to the predicted profile. The residual is the sum of a permanent random walk and a transitory shock. Their variances come from the growth of the residual's variance with the lag $d = 1, \ldots, 9$: the slope identifies the permanent variance (0.0092) and half the intercept the transitory variance (0.0424). The initial permanent variance is the cross-sectional variance at age 23 net of the two (0.168).

Owner mobility. The annual moving rates of owners used as an external check pair consecutive waves from 2–3 to 12–13. The base sample is household reference persons in England who own outright or with a mortgage at wave t, aged 23–90, with a positive cross-sectional weight. A move means the person's address changed between t and $t + 1$, among those observed at $t + 1$. Weights are normalized to mean one within each pair of waves. The rates are 9.6 percent at ages 25–34 (6,975 person-years), 3.6 percent at 35–64 (60,170), and 2.2 percent at 65–75 (22,605).

A.6. Mortality and energy bills.

Mortality. Survival probabilities come from the ONS National Life Tables for the United Kingdom, 2015–17, averaging the male and female number of survivors l_x at each age. The model uses the implied annual survival probabilities for ages 23–81. Beyond the table, hazards are extrapolated with

a constant growth rate equal to the median of the last seven hazard ratios, bounded between 1 and 1.25. The survival probability in the recurring terminal state (82 and over) is one minus the inverse of the expected number of remaining periods at 82, which gives 0.880.

Energy bills. Energy bills by efficiency state are based on owner-occupied certificates in the EPC register for homes built by 1990, lodged between 2015 and 2019, using the latest certificate of each building (2,833,945 buildings). We trim modelled consumption to 1–1,500 kWh per square meter and floor area to 10–1,000 square meters. The register reports modelled primary energy use. Unweighted means are 190.9 kWh per square meter a year in bands A–C and 322.3 in D–G. We convert primary to delivered energy by dividing by 1.59, the SAP 2012 primary-energy factors (1.22 for gas, 3.07 for electricity) averaged with an assumed 80/20 gas/electricity mix. We then multiply by a common floor area of 101.75 square meters and a price of £0.0602 per kWh. The price is the 80/20 average of DESNZ domestic unit prices, including VAT and excluding standing charges, averaged over 2015–2019: £0.0381 for gas (Great Britain) and £0.1485 for electricity (United Kingdom). Annual bills are £735 in A–C and £1,242 in D–G.

A.7. Samples behind the estimation targets. Table A2 lists the source, sample, statistic, number of observations, and number of releases or rounds behind each of the eleven targets. Different targets describe different populations. The homeownership rates are shares of all households. The A–C shares condition on ownership and on homes built by 1990. The wealth ratios cover all households in the relevant age band.

Homeownership. The three homeownership targets pool the seven releases marked in Table A1 and include every household with a positive weight, whatever its tenure. Ages 25–34 and 35–64 use the banded age. Ages 65–75 use exact age and so rest on 11plus12, 12and13, and fs13.

A–C shares. The A–C targets use the five housing-stock releases with a physical survey. The sample consists of owners of homes built by 1990 with a recorded EPC band. Ages 65–75 use exact age and so rest on 11plus12 and 12and13. These two releases cover fieldwork from 2011–12 to 2013–14, and the 2012–13 households appear in both. The recent-buyer shares quoted in Section 2 use the same five releases and the same definitions. They are not targets.

Table A2: Samples behind the estimation targets

Target	Source	Sample and statistic	Obs.	Waves
Homeownership				
Ages 25–34	EHS	All households; weighted owner share	12,533	7
Ages 35–64	EHS	All households; weighted owner share	46,055	7
Ages 65–75	EHS	All households, exact age; weighted owner share	5,801	3
Owners in an A–C home, built by 1990				
Ages 25–34	EHS	Owners of homes built by 1990; weighted A–C share	1,646	5
Ages 35–64	EHS	Owners of homes built by 1990; weighted A–C share	12,279	5
Ages 65–75	EHS	As above, exact age	1,831	2
Wealth, consumption and income				
Housing share of wealth, 35–64	WAS	Households with positive gross wealth; weighted mean of the household's housing share	34,152	4
Liquid wealth to income, 65–75	WAS	HRP aged 65–74, positive income; weighted median ratio	15,587	4
Liquid wealth to income, 75+	WAS	HRP aged 75 and over, positive income; weighted median ratio	12,911	4
Consumption to income, 50–64	ONS, WAS	Mean core expenditure (published tables) over WAS mean net income	5,349	2 years
Retirement income, 65–74 to 55–64	WAS	Ratio of weighted mean net household income	7,652	2 years

Notes: Observations are pooled households before equal-wave averaging. EHS waves are the releases in Table A1; WAS waves are rounds 5–8. The last two targets use WAS households interviewed in the financial years ending in 2016 and 2017. The count for the consumption ratio refers to the income denominator.

B. Model Details

B.1. The 82+ group and the simulation. The model has 60 ages, $J_0 = 23$ to $\bar{J} = 82$. The last stands for all ages from $\bar{J}$ on: a household in this group survives each year with probability $\phi_{\bar{J}}$, and its income, equivalence scale, and choice set do not change. Its value function is therefore the same every year and solves a fixed point. Simulated households are followed for 70 years from age J_0; those still alive after $\bar{J}$ follow the policy of age $\bar{J}$.

B.2. Mortgages. Let t_0 be the efficiency of an owner's current home and $B = (1 + r_m)m - P(m, n)$ the balance due next year under the existing schedule. The rate r_m is the same for all loans. The remaining term falls by one each year. A loan taken by a household without a mortgage has term L. A mover with a positive balance keeps its remaining term, which falls by one as for a stayer. When one year remains, the payment $P(m, 1) = (1 + r_m)m$ repays the balance. The repayment schedule continues in retirement; there is no interest-only phase.

Purchase. A buyer's loan satisfies

$$m' \le \min\{\lambda_L p^t h',\ P^{-1}(\lambda_D y^{\text{perm}}; n')\},$$

where $P^{-1}(\cdot; n')$ is the largest balance whose payment over n' years equals its argument, n' is the term of the new loan, and y^{perm} is gross income with the transitory shock set to one (the pension in retirement). For a purchase with renovation, $t = N$. An owner who moves also faces

$$m' \le B + \max\{p^t h' - p^{t_0} h,\ 0\} + C(h')\,\mathbf{1}\{\text{renovates at purchase}\}.$$

A mover can therefore borrow for a more valuable home and for a renovation at purchase, but cannot take equity out; a move to a home of equal value continues the existing schedule.

Renovation in place. An owner who renovates in place receives the advance

$$A = \min\{C(h),\ \max\{K - B,\ 0\}\}, \qquad K = \min\{\lambda_L p^N h,\ P^{-1}(\lambda_D y^{\text{perm}}; \hat{n})\},$$

where $\hat{n} = n - 1$ for an owner with a mortgage and $\hat{n} = L$ for an owner without one. The advance is added to next year's balance and is repaid by the existing payoff date. The collateral is the inefficient value of the home.

B.3. Initial conditions. Households enter at age 23 with the income, liquid wealth, tenure, home value, mortgage balance, and remaining term of records of household heads aged 20–29 in the Wealth and Assets Survey; Appendix A.4 describes the sample. Records are drawn in proportion to their weights. Income in the record sets the permanent component, with a floor, and the transitory component starts at one. Negative net financial wealth is set to zero, since the model has no unsecured debt. An owner's home size is its value in the record, kept within the limits $[\underline{h}^{\text{own}}, \bar{h}^{\text{own}}]$. An owner is placed in an efficient home with probability π_0^R; the market value of its home is kept, so the size of an

efficient home is its value divided by p^R.

B.4. Budget constraints. Let x be cash on hand. Consumption is the residual in each budget constraint below.

Renting. A renter who keeps renting, or an owner who sells and rents, pays

$$c + b' + \left[p^{\text{rent}} + \pi_R e^R + (1 - \pi_R) e^N\right] h' = x + S,$$

where $S = 0$ for a renter and, for an owner who sells, S is the sale proceeds

$$S = (1 - \delta_h - \zeta_s)\, p^{t_0} h - (1 + r_m) m.$$

The seller pays the year's maintenance on the home it sells and repays the mortgage in full. Rental size lies between $\underline{h}^{\text{rent}}$ and $\bar{h}^{\text{rent}}$.

Buying. A household that buys a home of size h' and state t pays

$$c + b' + (1 + \zeta_b)\, p^t h' + e^t h' = x + S + m'.$$

A household that buys an inefficient home and renovates it pays $(1 + \zeta_b) p^N h' + C(h') + e^R h'$ for the home, the work, and the energy bill. A buyer pays the energy bill of the new home in the year of purchase and no maintenance on it until the next year.

Staying. An owner who stays pays

$$c + b' + \tilde{P} + \delta_h p^{t_0} h + e^{t_0} h = x, \qquad P(m, n) \le \tilde{P} \le (1 + r_m) m,$$

and its balance next year is $m' = (1 + r_m) m - \tilde{P}$. An owner who renovates in place also pays $C(h)$, and the lower limit on its payment falls to $P(m, n) - A$, where A is the renovation advance of Appendix B.2. In the year of the renovation it pays maintenance and the energy bill of an inefficient home, and it receives the services of an efficient home.

C. The Simplified Model: Proofs

Setup. We use the simplified model of Section 3.3. We suppress the parts of the state that do not change, and write $V_{j+1}(a, t)$ for the value next year, before the taste shocks are drawn, of an owner of a home of type t that leaves this year with financial wealth a, saving less mortgage. Next year the owner draws one shock for keeping its home and one for selling it, distributed as in Section 3.2. Let $Q^k_{j+1}(a, t)$ be the value of keeping the home: that of Stay for an efficient home, and the higher of $Q^{\text{stay},N}_{j+1}$ and $Q^{\text{stay},R}_{j+1}$ for an inefficient one. Let $Q^m_{j+1}(a, t)$ be the value of selling, the highest value among the alternatives that sell the home. Then

$$V_{j+1} = \sigma_\varepsilon \log\left[\exp\left(Q^k_{j+1}/\sigma_\varepsilon\right) + \exp\left(Q^m_{j+1}/\sigma_\varepsilon\right)\right], \tag{C.1}$$

and the owner keeps its home with probability $\partial V_{j+1}/\partial Q^k_{j+1}$. V'_{j+1} is the marginal value of cash on hand next year, so that $\partial V_{j+1}/\partial a = (1+r)V'_{j+1}$, where r is the rate on the asset the household adjusts at the margin; we write V'^{k}_{j+1} and V'^{m}_{j+1} for its values in the two branches. When the household can adjust that asset without hitting a borrowing limit, its Euler equation is

$$u_c = \beta\phi_j(1+r)V'_{j+1} + (1-\phi_j)\,g', \tag{C.2}$$

since by equation (3.16) the estate is not discounted. With $\tilde{\phi}_j$ from equation (3.20), this gives $\tilde{\phi}_j = \beta\phi_j(1+r)V'_{j+1}/u_c$, $1-\tilde{\phi}_j = (1-\phi_j)g'/u_c$, and $\tilde{\phi}_j\Lambda = \beta\phi_j\varrho V'_{j+1}/u_c$. We write $R_s \equiv (1-\delta_h-\zeta_s)\Delta p$ for what the efficient state adds to the proceeds of a sale, and $q \equiv \tilde{\phi}\Lambda$.

Outline. The proof of Proposition 1 compares renovating this year with renovating next year. Lemma C.1 writes the gain from renovating now as this year's comfort, less the cost, plus the value of the efficient state next year. Lemma C.2 bounds that value. An owner of the inefficient home can match an efficient owner next year: if the efficient owner keeps its home, by renovating, which costs $\overline{C}$ more; if it sells, by selling too, with R_s less. Whether an owner sells can itself depend on whether its home is efficient. But the owners whose choice turns on efficiency are close to indifferent between keeping and selling, so to first order the state is worth the same to them either way.

The proof of Proposition 2 compares the buyer with an owner of the inefficient home who lives as the buyer does, but without the state. It needs no proviso about renovating later. The state is cheaper to buy than to make, $\overline{C}_b < C(h') + \Delta e' - \delta_h\Delta p'$, so near the buyer's threshold an owner of the inefficient home would not renovate it.

Two routes, one horizon. Let C be the extra cost in the year the household gets the efficient state, over the extra cost of holding it for a year: $C = \overline{C}$ in place and $C = \overline{C}_b$ at a purchase. Rearranged, conditions (3.21) and (3.23) both say that the household gets the efficient state if and only if

$$\theta c + (\Delta e - \delta_h\Delta p) + (1-\tilde{\phi})\,\Delta p + \tilde{\phi}\,(1-\mu)\,\Lambda\,(1-\delta_h-\zeta_s)\Delta p \;\geq\; \frac{C}{\widetilde{H}}, \qquad \widetilde{H} \equiv \frac{1}{1-\tilde{\phi}\Lambda\mu}, \tag{C.3}$$

with θ, Δe, and Δp those of the home concerned. The left side is what the efficient state yields in a year: the comfort, the bill saving net of the maintenance on the premium, the premium that the heirs receive if the owner dies, and the premium net of costs that the owner receives if it moves. The right side spreads the cost C over the horizon $\widetilde{H}$. At constant parameters $H_j = 1/(1-\beta\phi\varrho\mu)$ in equation (3.19), while $\tilde{\phi}\Lambda = \beta\phi\varrho V'_{j+1}/u_c$ by the Euler equation. So $\widetilde{H}$ is H_j with next year weighted by its marginal utility relative to this year. The two routes face the same horizon, but spread different costs over it.

Lemma C.1 (Decomposition)**.** *Take an owner of an inefficient home of size h at age j that keeps its home this year. Let c and a' be its consumption and end-of-year wealth if it does not renovate, $C \equiv C(h)$, and suppose that*

it can reduce a' by C at the rate r. Then $G_j \equiv Q_j^{\text{stay},R} - Q_j^{\text{stay},N}$ satisfies, to first order,

$$\frac{G_j}{u_c} \approx \theta c - C + \tilde{\phi}_j \Lambda D_{j+1} + (1-\tilde{\phi}_j)\Delta p, \qquad D_{j+1} \equiv \frac{V_{j+1}(a'-C,R) - V_{j+1}(a'-C,N)}{V'_{j+1}}, \tag{C.4}$$

where D_{j+1} is the money value of being efficient next year.

Proof. If it does not renovate, $Q_j^{\text{stay},N} = u_j(c, s_N) + (1-\phi_j)g(a' + p^N h) + \beta\phi_j V_{j+1}(a', N)$. Renovating, consuming c, and leaving the year with $a' - C$ is feasible. In the year of the work the renovator receives s_R and pays the inefficient bill and maintenance (Appendix B.4). If it dies, the estate values the home at $p^R h$. If it survives, the home is efficient next year with probability ϱ. So

$$\begin{aligned} G_j \geq \big[u_j(c,s_R) - u_j(c,s_N)\big] - \big[(1-\phi_j)\big(g(a'+p^N h) - g(a'-C+p^N h)\big) + \beta\phi_j\big(V_{j+1}(a',N) - V_{j+1}(a'-C,N)\big)\big] \\ + \beta\phi_j\varrho\big[V_{j+1}(a'-C,R) - V_{j+1}(a'-C,N)\big] + (1-\phi_j)\big[g(a'-C+p^R h) - g(a'-C+p^N h)\big]. \end{aligned}$$

To first order the four terms are $u_c\theta c$, $-Cu_c$ by the Euler equation (C.2), $\beta\phi_j\varrho V'_{j+1}D_{j+1} = u_c\tilde{\phi}_j\Lambda D_{j+1}$, and $(1-\phi_j)g'\Delta p = u_c(1-\tilde{\phi}_j)\Delta p$. The renovator's best allocation differs from $(c, a'-C)$ by a first-order amount, so by the envelope theorem the inequality holds with equality to first order. □

Lemma C.2 (Matching)**.** *Let*

$$D^k \equiv \frac{Q_{j+1}^k(a,R) - Q_{j+1}^k(a,N)}{V'^k_{j+1}}, \qquad D^m \equiv \frac{Q_{j+1}^m(a,R) - Q_{j+1}^m(a,N)}{V'^m_{j+1}}$$

be the money values of the efficient state next year to an owner that keeps its home and to one that sells it. To first order:

(i) $D_{j+1} = \mu_{j+1}D^k + (1-\mu_{j+1})D^m$*, where μ_{j+1} is the probability of keeping the home, averaged between the inefficient and the efficient state and weighted by marginal utility;*

(ii) $D^k \leq \overline{C}$*, with equality if an owner of the inefficient home that keeps it would renovate it;*

(iii) $D^m = R_s$*;*

(iv) $D_{j+1} \leq U_{j+1}$*, the money value of the state along the efficient owner's path: the comfort and the bill saving net of maintenance in each year it holds the state, the premium its heirs receive if it dies, and R_s if it sells. Equality holds if an owner of the inefficient home would never renovate it.*

Proof. (i) By (C.1), the derivatives of V_{j+1} with respect to Q_{j+1}^k and Q_{j+1}^m are the probabilities of keeping and of selling the home. Integrating along the line from the values of the inefficient state to those of the efficient one gives, exactly,

$$V_{j+1}(a,R) - V_{j+1}(a,N) = \bar{\mu}\big[Q_{j+1}^k(a,R) - Q_{j+1}^k(a,N)\big] + (1-\bar{\mu})\big[Q_{j+1}^m(a,R) - Q_{j+1}^m(a,N)\big],$$

where $\bar{\mu}$ is the probability of keeping the home averaged along the line. Dividing by V'_{j+1} gives (i), with $\mu_{j+1} = \bar{\mu}V'^{k}_{j+1}/\left[\bar{\mu}V'^{k}_{j+1} + (1-\bar{\mu})V'^{m}_{j+1}\right]$; the denominator differs from V'_{j+1} by a first-order amount, which multiplies a first-order term. The probabilities of keeping the home in the two states also differ by a first-order amount, so to first order it does not matter at which state they are taken. (ii) An owner of the inefficient home that keeps it can renovate and live as an efficient owner that stays. It gets the same services s_R and the same home next year, and its borrowing limit is no tighter, since it may borrow for the work. It pays $C + (\delta_h p^N + e^N)h$ instead of $(\delta_h p^R + e^R)h$, that is, $\overline{C}$ more. So $D^k \leq \overline{C}$, with equality when this is its best choice. (iii) An efficient home adds R_s to the proceeds of a sale (Appendix B.4), and after the sale the two owners face the same choices. (iv) An owner of the inefficient home can follow the efficient owner's path: given the same taste shocks, it keeps or sells its home when the efficient owner does, with the same consumption, but without the state. Its taste shocks are then the same, so to first order this costs it U_{j+1}, computed with the efficient owner's probabilities of keeping its home. So $D_{j+1} \leq U_{j+1}$, with equality when never renovating is its best choice; its own choices between keeping and selling then differ from the efficient owner's only at the margin, where they cost nothing to first order. □

Proof of Proposition 1. Let $\ell_{j+1} \equiv \overline{C} - D_{j+1}$. Substituting $D_{j+1} = \overline{C} - \ell_{j+1}$ in (C.4),

$$\frac{G_j}{u_c} \approx \theta c - \left(C - \Lambda\overline{C}\right) - (1-\tilde{\phi}_j)\left(\Lambda\overline{C} - \Delta p\right) - \tilde{\phi}_j \Lambda\, \ell_{j+1}. \tag{C.5}$$

By Lemma C.2(i)–(iii), $\ell_{j+1} = \mu_{j+1}(\overline{C} - D^k) + (1-\mu_{j+1})\ell \geq (1-\mu_{j+1})\ell$, so $G_j \geq 0$ implies (3.21). Under the proviso, $D^k = \overline{C}$ and $\ell_{j+1} = (1-\mu_{j+1})\ell$, so (C.5) has the sign of (3.21), which is then also sufficient. □

Remark C.1 (Constant parameters). Suppose that $\tilde{\phi}$, μ, Λ, and consumption are constant over the following years, and that no borrowing limit binds. Then (3.21) is necessary and sufficient without the proviso. Let $Z \equiv \theta c + (\Delta e - \delta_h \Delta p) + (1-\tilde{\phi})\Delta p$, the yield of the state in a year in which the owner keeps it. Since $\theta c - C + (1-\tilde{\phi})\Delta p = Z - \overline{C}$, (C.4) reads $G_j/u_c \approx Z - \overline{C} + qD_{j+1}$. Next year, an owner of the inefficient home can renovate the first time it keeps its home, which values the state at $\overline{D} \equiv \mu\overline{C} + (1-\mu)R_s$; never renovate, which values it at U with $U = \mu(Z + qU) + (1-\mu)R_s$; or renovate after k years, which values it at $U + (\mu q)^k(\overline{D} - U)$. It takes the best of these, so $D_{j+1} = \min\{U, \overline{D}\}$. Let $E \equiv Z - (1-\mu q)\overline{C} + q(1-\mu)R_s$; $E \geq 0$ is (C.3) with $C = \overline{C}$, that is, (3.21) rearranged. If $D_{j+1} = \overline{D}$, then $G_j/u_c \approx E$. If $D_{j+1} = U$, then $(1-\mu q)\, G_j/u_c \approx (1-\mu q)(Z - \overline{C}) + q\left[\mu Z + (1-\mu)R_s\right] = E$. Either way G_j has the sign of E.

Proof of Proposition 2. Write $R'_s \equiv (1 - \delta_h - \zeta_s)\Delta p'$, $Z' \equiv \theta' c + (\Delta e' - \delta_h \Delta p') + (1-\tilde{\phi})\Delta p'$, and $\overline{C}' \equiv C(h') + \Delta e' - \delta_h \Delta p'$ for the in-place cost of h'. Compare the two purchases at the consumption c and the size h' of the buyer of the inefficient home. Getting the state costs C_b more this year and saves the bill $\Delta e'$, so the buyer of the state leaves the year with $C_b - \Delta e'$ less. If it dies, the estate values the home at $p^R h'$. As in Lemma C.1, to first order the gain $G_b \equiv Q_j^{\text{buy},R} - Q_j^{\text{buy},N}$ satisfies

$$\frac{G_b}{u_c} \approx \theta' c + \Delta e' - C_b + qD' + (1-\tilde{\phi})\Delta p' = Z' - \overline{C}_b + qD',$$

where D' is the money value of the state next year. As in Remark C.1, an owner of the inefficient home h' can renovate it the first time it keeps it, which values the state at $\overline{D}' \equiv \mu\overline{C}' + (1-\mu)R_s'$, or never renovate it, which values it at $U' = [\mu Z' + (1-\mu)R_s']/(1-\mu q)$, or wait; so $D' = \min\{U', \overline{D}'\}$. Both Z' and U' rise with c, so G_b is continuous and strictly increasing in c. Where $D' = U'$,

$$(1-\mu q)\,\frac{G_b}{u_c} \approx E_b \equiv Z' - (1-\mu q)\overline{C}_b + q(1-\mu)R_s',$$

and $E_b \geq 0$ is (C.3) with $C = \overline{C}_b$, that is, (3.23) rearranged. At the consumption where $E_b = 0$, $U' = \mu\overline{C}_b + (1-\mu)R_s' \leq \overline{D}'$, since $\overline{C}_b \leq C(h') - \delta_h\Delta p' < \overline{C}'$. So $D' = U'$ and $G_b = 0$ there. As G_b and E_b both increase in c, $G_b \geq 0$ if and only if $E_b \geq 0$. □

Lemma C.3 (The payoff horizon)**.** *Let a household buy a home of size h' at age j, and let $\tilde{\phi}_{j+k}$, Λ_{j+k}, μ_{j+k+1}, and c_{j+k} vary with k. Suppose that an owner of the inefficient home h' would never renovate it, and that no borrowing limit binds. Let $P_0 = 1$ and $P_k = \prod_{i<k}\tilde{\phi}_{j+i}\Lambda_{j+i}\mu_{j+i+1}$, the discounted probability of still holding the state k years after the purchase. Then, to first order, the household buys the home efficient if and only if*

$$\sum_{k\geq 0} P_k\Big[\theta' c_{j+k} + (\Delta e' - \delta_h\Delta p') + (1-\tilde{\phi}_{j+k})\Delta p' + \tilde{\phi}_{j+k}\Lambda_{j+k}(1-\mu_{j+k+1})R_s'\Big] \;\geq\; \overline{C}_b.$$

The weights sum to $\widetilde{H}_j \equiv \sum_k P_k$. Since $\tilde{\phi}_{j+i}\Lambda_{j+i} = \beta\phi_{j+i}\varrho V'_{j+i+1}/u_{c,j+i}$ and $V'_{j+i+1} \approx u_{c,j+i+1}$,

$$P_k \approx \beta^k\Big(\prod_{i<k}\phi_{j+i}\,\varrho\,\mu_{j+i+1}\Big)\frac{u_{c,j+k}}{u_{c,j}},$$

so $\widetilde{H}_j$ is the payoff horizon H_j of equation (3.19), with each year weighted by its marginal utility relative to the year of purchase. At constant parameters the condition is (C.3) with $C = \overline{C}_b$.

Proof. Write Z_{j+k} for the first three terms in the bracket and $q_{j+k} \equiv \tilde{\phi}_{j+k}\Lambda_{j+k}$. As in the proof of Proposition 2, $G_b/u_c \approx Z_j - \overline{C}_b + q_j D'_{j+1}$, and under the assumptions $D'_{j+1} = U'_{j+1}$ by Lemma C.2(iv), where $U'_{j+k} = \mu_{j+k}\big(Z_{j+k} + q_{j+k}U'_{j+k+1}\big) + (1-\mu_{j+k})R_s'$. Substituting forward, and since $P_k \to 0$, $q_j U'_{j+1} = \sum_{k\geq 1} P_k Z_{j+k} + \sum_{k\geq 0} P_k q_{j+k}(1-\mu_{j+k+1})R_s'$. Adding $Z_j - \overline{C}_b$ gives the condition. □

Proof of Proposition 3. Under the assumptions of Proposition 2, the in-place condition (3.21) is necessary and sufficient by Remark C.1. So the thresholds c_j^* and c_b^* are the consumption at which (C.3) holds with equality, with $C = \overline{C}$ and $C = \overline{C}_b$, and $\theta c^* = (1-\tilde{\phi}\Lambda\mu)\,C - (\Delta e - \delta_h\Delta p) - (1-\tilde{\phi})\Delta p - \tilde{\phi}(1-\mu)\Lambda R_s$. A higher probability of moving lowers μ, and a higher probability of dying lowers $\tilde{\phi}$.

Moving. The derivative of c^* in $1-\mu$ is $\tilde{\phi}\Lambda(C - R_s)/\theta$, and $C - R_s$ is ℓ in place and ℓ_b at a purchase. Since $C_b \leq (1+\zeta_b)\Delta p'$, $\ell_b \leq (\zeta_b + \zeta_s)\Delta p'$.

Death. The derivative of c^* in $1-\tilde{\phi}$ is $[\Lambda(\mu C + (1-\mu)R_s) - \Delta p]/\theta$. Here $\Lambda(\mu C + (1-\mu)R_s)$ is what the state is worth to the owner alive next year, by Lemma C.2 in place and by the proof of Proposition 2 at a purchase, and Δp is what the heirs receive. In place the derivative is at least $(\Lambda\mu\overline{C} - \Delta p)/\theta$, which is positive under the condition of part (a). At a purchase, $\mu\overline{C}_b + (1-\mu)R_s' \leq (1 + \zeta_b - \delta_h)\Delta p'$, since

$\overline{C}_b \le (1+\zeta_b-\delta_h)\Delta p'$ and $R'_s < \Delta p'$. So the derivative is at most $[\Lambda(1+\zeta_b-\delta_h)-1]\Delta p'/\theta'$, which is negative under the condition of part (b). As $\Lambda \le \varrho$, that condition holds when $\varrho < 1/(1+\zeta_b-\delta_h) = 0.981$.

The gap. With $h' = h$, $\theta(c_j^* - c_b^*) = (\overline{C} - \overline{C}_b)/\widetilde{H}$ and $\overline{C} - \overline{C}_b = C(h) - C_b + \Delta e > 0$, since $C_b \le C(h)$. And $1/\widetilde{H} = 1 - \tilde{\phi}\Lambda\mu$ falls in $\tilde{\phi}$ and μ. A household with $c_b^* \le c < c_j^*$ would not renovate in place by Proposition 1, and would take an efficient home at a purchase by Proposition 2. □

D. Computation

D.1. Networks. Each age has its own value network and policy network, 120 networks for the 60 ages. Both take the same 42 inputs: the state, functions of it, the seven parameters in ϑ, and the discount factor β. Each parameter is scaled to $[0,1]$ over its search range, the bequest scale υ on a log scale.

The value network has three hidden layers of 512 units with the SiLU activation. The activation is smooth because the maximization at the previous age reads the slope of the value in liquid wealth. The network does not fit the value directly. When the bequest motive is strong, the value of a household is dominated by the bequest term, and the accuracy the choices need is a small fraction of it. The network therefore fits the deviation from an analytic baseline,

$$\widehat{V}_j(\mathbf{x}_j, \vartheta) = A_j(\beta)\, g(\tilde{q}_j) + \bar{v}_j + \varsigma_j\, N_j(\mathbf{x}_j, \vartheta), \qquad A_j(\beta) = \sum_{k\ge 0} \beta^k \Big(\prod_{i=0}^{k-1} \phi_{j+i}\Big)(1-\phi_{j+k}),$$

where g is the bequest utility (3.3), $\tilde{q}_j$ is the household's net worth at the start of the period (cash on hand plus home equity), $A_j(\beta)$ is the discounted probability of dying at each future age, N_j is the network, and $\bar{v}_j$ and ς_j standardize the values at age j.

The policy network has two hidden layers of 128 units and two to four outputs for each alternative. For every alternative it predicts the logarithm of consumption and the probability that the household consumes the most the alternative allows: it saves nothing and, if it buys, borrows only what the purchase requires. When this probability exceeds 0.9, the household consumes that maximum. An owner who stays also has a payment share, which places the mortgage payment between the smallest payment the contract allows and the largest compatible with the predicted consumption. A buyer also has the size share a_h and the mortgage share a_m, described below. A renter's home size is the optimal size given total spending, in closed form, because it does not affect the future. Liquid saving is the residual in every alternative.

These outputs are not the choices themselves. A network outputs unrestricted real numbers, but the choices must satisfy the budget and borrowing constraints, so a fixed mapping turns the outputs into a feasible allocation. Take a renter who buys a home. The network predicts consumption c and the two shares $a_h, a_m \in [0,1]$. The first places the size of the new home in the feasible range,

$$h' = \underline{h}^{\text{own}} + a_h\big(\overline{h}^{\text{own}} - \underline{h}^{\text{own}}\big),$$

and the second places the mortgage between the smallest balance that finances the purchase together

with c and the largest balance that the loan-to-value and payment limits allow. Liquid saving is what is left, so the budget holds by construction. If c is not affordable at h', the size is reduced until it is. An error in a_m then moves money between the mortgage and liquid saving and leaves consumption unchanged. A purchase involves large sums relative to a year's consumption, so this matters: if the network predicted saving instead, a small error in it would be a large error in consumption.

D.2. Training. At each age, the training sample has 20,480 points. It consists of 4,096 draws of the state and the parameters, each with four values of liquid wealth, and 512 groups of points with small mortgages. The parameters are drawn uniformly over the box, v log-uniformly. Ten percent of the draws are held out to measure the fit.

At each point the expectation in equation (3.16) is computed by quadrature, with five nodes for the permanent income shock and three for the transitory shock, and exactly over whether an efficient home stays efficient. For every feasible alternative, the maximization in equation (3.17) combines a grid in the choices refined by golden-section search, a grid in consumption, and candidates on the boundary of the feasible set, and keeps the best. The target of the value network is the logsum (3.18). By the envelope condition, its derivative with respect to liquid wealth is $(1+r_b)$ times the probability-weighted marginal utility of consumption across alternatives.

The value network minimizes the relative error in the value plus 0.01 times the error in this derivative, in 24,000 steps of the Adam optimizer, with a learning rate that decays from 10^{-4} to 10^{-6}. The policy network fits the maximizing choices in 8,000 steps, with a learning rate that decays from 10^{-3} to 10^{-5}. In its loss, the choices of each alternative are weighted by $0.05 + 0.95$ times the probability that the household chooses it, so the fit concentrates on the alternatives that households choose. Infeasible alternatives have zero weight. Both networks start from the trained networks of age $j + 1$.

The networks have no explicit penalty on their weights, such as weight decay or dropout. Instead, three devices keep them from fitting noise in the sample. Every 2,000 steps for the value network and every 1,000 steps for the policy network, the loss is computed on the held-out points, and the weights with the lowest held-out loss are kept. The norm of the gradient is clipped at 5. And each step uses a random minibatch of 2,048 points. Table D3 collects these settings.

At ages 82 and above the problem is stationary. The networks of this age are fitted to targets computed from their previous iteration, at least 6 and at most 12 times. The iteration stops early when the choices change by less than a tolerance in two consecutive iterations. With the 59 younger ages, training has 71 stages. They took 9.7 hours on one GPU, 8.8 of them in computing the targets.

D.3. Accuracy. On the held-out points at age 23, the value network's mean relative error is 0.3 percent and its 95th percentile is 1.0 percent. We also check the policies at the estimate with Euler-equation errors, expressed as a share of consumption, at 128 states at each of the ages 23, 40, 64, 65, 78, and 82+, among the states in which neither the borrowing limit nor the consumption floor binds. The median error is 0.9 percent at 23 and at 82+, 1.2 percent at 78, 3.3 percent at 40, and 2.9 and 4.7 percent at 64 and 65. The 95th percentiles range from 3.6 percent at 78 to 13.6 percent at 65. The largest error, 54.6 percent, is at age 40. The errors are largest around retirement, where income falls.

Table D3: Network architecture and training

	Value network	Policy network
Architecture		
Inputs	42	42
Hidden layers	3×512 units	2×128 units
Activation	SiLU	SiLU
Outputs	1, the deviation from the bequest baseline	57, two to four for each alternative
Sample at each age		
Points	20,480	20,480
Held out	10 percent of the draws	10 percent of the draws
Loss		
Fit	squared error relative to $1 + \lvert V \rvert$	Huber ($\delta = 0.25$) in log consumption; cross-entropy in shares and probabilities
Derivative	$0.01\times$ squared relative error in $\partial V / \partial b$	–
Weight of a point	equal	$0.05 + 0.95\times$ choice probability of the alternative
Optimization		
Optimizer	Adam	Adam
Steps	24,000	8,000
Learning rate, cosine decay	10^{-4} to 10^{-6}	10^{-3} to 10^{-5}
Minibatch	2,048 points	2,048 points
Initial weights	networks of age $j + 1$	networks of age $j + 1$
Regularization		
Weight decay, dropout	none	none
Gradient-norm clipping	5	5
Early stopping	best held-out loss, checked every 2,000 steps	best held-out loss, checked every 1,000 steps

Notes: Each of the 60 ages has one network of each kind. The sample has 4,096 draws of the state and the parameters, each at four values of liquid wealth, and 512 draws of owners, each with a small mortgage and with none, at the same four values. Points that share a draw are held out together.

D.4. Search. The search has three stages, all with the same draws of the shocks. First, it simulates 96 points of a scrambled Sobol sequence over the box with 2,048 households each. Three of the points are replaced by estimates from earlier runs. It keeps the best 12. Second, it refines them locally over three rounds, drawing 16, 12, and 8 normal perturbations of each candidate with standard deviations of 10, 5, and 2 percent of the width of each range, with 4,096 households, and keeping the best 8 after each round. Third, it simulates the best 10 candidates with 32,768 households, and the best of them is the estimate. The search evaluates 486 candidates. The moments in Panel B of Table 5 are simulated at the estimate with 65,536 households and new draws.

D.5. Standard errors. The derivative G in equation (4.2) is computed by central differences at ±3 percent of the width of each range, with 65,536 households and the same draws at every point, on the scale on which the moments enter the criterion: levels for shares, logs for ratios. The covariance Ω of the data moments comes from a bootstrap of the surveys that resamples clusters, on the same scale.

The simulation noise is $\Sigma_S = r\, D\Omega D$, with $D = \mathrm{diag}(1/\sqrt{S_k})$. Here S_k is the number of simulated households in the age group of moment k, at the 4,096 households of the search, divided by the number of observations in the data. A simulated household counts once however many ages it contributes, so $D\Omega D$ overstates the noise. The factor $r = 0.069$ corrects for this. It is the ratio of the simulation noise measured at four parameter vectors, as the difference between two independent sets of draws with 65,536 households, to the noise $D\Omega D$ predicts for that number of households.

E. A Longer-Lasting Payoff

The retention of efficiency ϱ is imprecisely estimated and lies close to the lower end of its search range (Section 4.5). At the estimate, $\varrho = 0.84$, an efficient home keeps its rating for about six years on average. This appendix repeats the three calculations of Section 5 at $\varrho = 0.96$, at which a home keeps its rating for 25 years. The other parameters stay at the estimates of Table 5, and households re-optimize. The networks of Appendix D take ϱ as an input and were trained over its search range, $[0.80, 0.99]$, so the model need not be solved again. The calculations ask how the decomposition changes when efficiency lasts longer, not whether the model still fits the data: at $\varrho = 0.96$ the efficient share is 43.1, 45.9, and 36.8 percent of owners at 25–34, 35–64, and 65–75, about twice the shares in the data.

E.1. The payoff horizon. Table E4 computes the horizon of Section 5.1 at both values of ϱ. Column (1), with discounting and survival only, does not depend on ϱ. Once the loss of efficiency is added, the horizon at 75 is 83 percent of the horizon at 25 at the estimate (column (2)), but 55 percent at $\varrho = 0.96$ (column (5)): when efficiency lasts, survival starts to matter. Moving undoes most of this. With the model's probabilities of moving, the horizon at $\varrho = 0.96$ is 6.2 years at 25 and 5.4 at 75 (column (6)); with those in Understanding Society it is 6.7 and 7.0 years (column (7)). Most owners move long before a rating that lasts 25 years wears off, so the horizon stays short and close to flat in age.

E.2. The accounting. Table E5 repeats Table 7 at $\varrho = 0.96$. Buying an efficient home is still the main way into the stock and the only one that falls steeply with age: 8.8 per 100 owners at 25–34, 4.4 at 35–64, and 1.7 at 65–75, while renovating in place adds 1.0 to 1.1 at every age (Panel A). The loss of efficiency is now a smaller outflow, 1.5 to 2.1 per 100 owners against 2.3 to 4.6 at the estimate, so older owners keep more of what they bought when young, and the stock falls less with age than the inflow. In Panel B purchases fall from 15.0 per 100 owners at 25–34 to 3.8 at 65–75, close to the 15.0 and 4.5 at the estimate. Unlike at the estimate, the share of those purchases that are of efficient homes also falls with age, from 58 percent to 45 percent. Panel C attributes the fall in the inflow. Purchases account for 90 percent of the fall from 25–34 and 89 percent of the fall from 35–64, against 99 and 103 percent at the estimate. The share of efficient purchases, the margin on which the horizon acts, accounts for 17 and 13 percent. Rebuilt from these margins, the efficient share at 65–75 would be 72.3 percent instead of 36.8 with purchases held at their rate at 25–34, and 41.6 percent with the share of efficient purchases held at its rate at 25–34.

E.3. Closing the window. Table E6 removes efficient homes from the market at both values of ϱ. A buyer can still buy an inefficient home and renovate it at once, an owner can still renovate in place, and households re-optimize. At $\varrho = 0.96$ the efficient share falls to 13.8, 10.3, and 12.4 percent at 25–34, 35–64, and 65–75, and the gap between 25–34 and 65–75 shrinks from 6.2 points to 1.4 (column (4)). At the estimate it shrinks from 9.6 points to 1.3. Households that cannot buy efficiency renovate more: at 25–34, renovation in place rises from 1.0 to 1.7 per 100 owners a year, and renovation at purchase from 0.3 to 0.6. Yet the gap that remains without the window is about the same at both values of ϱ. A longer-lasting payoff adds little to the part of the gradient that does not run through purchases.

The three calculations point the same way as at the estimate. A longer-lasting payoff gives the horizon a role, through the share of buyers who choose an efficient home, but the window still accounts for most of the gradient.

Table E4: The payoff horizon when efficiency lasts longer

		R0 ($\varrho = 0.84$)			$\varrho = 0.96$		
Age at	Discounting	+ decay	+ moving		+ decay	+ moving	
renovation	and survival		model	observed		model	observed
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
25	26.9	5.3	3.4	3.8	14.1	6.2	6.7
35	24.8	5.3	3.4	4.2	13.7	6.3	8.6
45	22.0	5.3	3.0	4.7	13.1	5.7	9.8
55	18.6	5.2	2.8	4.7	12.0	5.3	9.6
65	14.6	5.0	3.1	4.6	10.3	6.1	8.8
75	10.0	4.4	2.9	4.2	7.8	5.4	7.0
Age 75 relative to age 25	37%	83%	86%	112%	55%	86%	105%

Notes: The horizon H_j of equation (3.19), in years, as in Table 6. Columns (1)–(4) are those of Table 6, at the estimates. Columns (5)–(7) set $\varrho = 0.96$ and keep the other parameters at the estimates; the model's probability that an owner of an A–C home moves, in column (6), is from the simulation at $\varrho = 0.96$. Column (1) does not depend on ϱ. The last row is the horizon at 75 as a percentage of the horizon at 25.

Table E5: Flows into and out of efficient ownership when efficiency lasts longer

	Age		
	25–34 (1)	35–64 (2)	65–75 (3)
A. Flows into and out of A–C ownership, per 100 owners a year			
A–C owners (stock)	43.07	45.88	36.85
Inflows			
Bought an A–C home	8.76	4.43	1.73
Renovated at purchase	0.26	0.38	0.26
Renovated in place	0.99	1.13	1.04
Total	10.02	5.93	3.02
Outflows			
Lost A–C status	2.05	1.93	1.50
Sold, now renting	2.03	1.94	0.93
Sold, bought a D–G home	1.08	1.55	1.09
Died	0.03	0.17	0.63
Total	5.19	5.60	4.14
Net inflow	4.83	0.33	−1.12
B. Margins of the inflow			
Purchases by renters and D–G owners, per 100 owners	15.04	8.61	3.80
Percent of those purchases			
A–C home	58.3	51.4	45.3
Renovated at purchase	1.7	4.4	6.8
In-place renovation, percent of D–G owners	1.7	2.1	1.6

C. Decomposition of the fall in the inflow to 65–75

	From 25–34	From 35–64
Fall in the inflow, per 100 owners	6.99	2.91
Percent of the fall due to		
Purchases	90	89
A–C share of purchases	17	13
Share renovated at purchase	−7	−5
D–G owners	−2	−6
In-place renovation rate	1	9

Notes: Simulated at $\varrho = 0.96$, the other parameters at the estimates of Table 5, 65,536 households, who re-optimize. Rows, panels, and the decomposition as in Table 7.

Table E6: The efficient share of owners without efficient homes for sale

	A–C share of owners (%)			Gap, 25–34	Ratio, 25–34
	25–34	35–64	65–75	less 65–75	to 65–75
	(1)	(2)	(3)	(4)	(5)
$\varrho = 0.84$ (estimate)[a]					
A–C homes for sale	22.4	17.8	12.8	9.6	1.75
No A–C homes for sale	6.9	5.5	5.6	1.3	1.24
$\varrho = 0.96$					
A–C homes for sale	43.1	45.9	36.8	6.2	1.17
No A–C homes for sale	13.8	10.3	12.4	1.4	1.11
EHS (target)	21.0	16.9	9.7	11.3	2.16

Notes: Shares, gap, and ratio as in Table 8, other parameters at the estimates of Table 5. No A–C homes for sale: a buyer cannot buy an efficient home, but can still buy an inefficient home and renovate it at once, and an owner can still renovate in place; households re-optimize. EHS: the targets of Table 5.

[a] The rows at $\varrho = 0.84$ are those of Table 8, simulated on a GPU; the rows at $\varrho = 0.96$ are simulated on a CPU. At the estimates the CPU gives shares of 22.5, 17.7, and 12.7 percent; the two differ by simulation noise only.

F. Additional Results

Table F7 gives the numbers behind Figure 6, with the efficient share at 35–64, which the figure leaves out, and purchases at 65–75. The share at 35–64 lies between the shares at 25–34 and 65–75 in every row but one: at $\varrho = 0.96$ it is the highest of the three, 45.9 percent. Purchases at 65–75 rise with comfort, from 5.5 per 100 owners a year without it to 6.6 with comfort doubled, but stay at 6.0 to 6.1 as ϱ rises.

Table F7: The efficient share of owners when the payoff changes

	A–C share of owners (%)			Gap, 25–34	Ratio, 25–34	Purchases,
	25–34	35–64	65–75	less 65–75	to 65–75	65–75[a]
	(1)	(2)	(3)	(4)	(5)	(6)
At the estimates	22.4	17.8	12.8	9.6	1.75	6.1
Comfort value of A–C ($\eta = 0.061$)						
None	17.2	13.3	9.7	7.4	1.76	5.5
Doubled	27.9	23.1	16.6	11.3	1.68	6.6
A–C status decays ($\varrho = 0.84$)[b]						
$\varrho = 0.88$	28.0	23.5	16.7	11.4	1.68	6.0
$\varrho = 0.92$	34.7	32.0	23.3	11.5	1.49	6.0
$\varrho = 0.96$	43.1	45.9	36.8	6.2	1.17	6.1
EHS (target)	21.0	16.9	9.7	11.3	2.16	

Notes: The runs of Figure 6. Each row changes one element of the model, keeps the other parameters at the estimates of Table 5 (values in parentheses), and lets households re-optimize. Shares are A–C owners as a percentage of owners in each age group, from 65,536 simulated households. The gap is the share at 25–34 less the share at 65–75, in percentage points; the ratio is the share at 25–34 over the share at 65–75. EHS: the targets of Table 5.

[a] All home purchases, by renters and owners, per 100 owners aged 65–75 a year.

[b] The rows for ϱ are simulated on a CPU and the others on a GPU. The CPU simulation at the estimates gives shares of 22.5, 17.7, and 12.7 percent, a gap of 9.9, and a ratio of 1.78; the two differ by simulation noise only.