Late-Life Home Purchases and the Limits of the Housing Ladder

Alexandra Pripadcheva*

October 8, 2026

Abstract

This paper examines the life-cycle dynamics of mortgage-financed home purchases in the United States and documents a systematic discrepancy between observed housing transactions and the predictions of canonical life-cycle models. Using individual-level mortgage originations from the Home Mortgage Disclosure Act (HMDA) for 2018-2025, I show that older households remain active participants in the housing market and, conditional on purchasing a home, acquire substantially more expensive properties than younger households.

An age-period-cohort analysis reveals a robust, increasing age profile of property values that persists across alternative identification strategies, geographic markets, borrower characteristics. These findings contrast with the conventional housing ladder hypothesis, which links housing acquisition primarily to wealth accumulation early in life and predicts declining housing expenditure toward retirement. I evaluate this discrepancy using a calibrated quantitative life-cycle model. Despite matching aggregate moments reasonably well, the model fails to reproduce the observed distribution of transactions across age groups, predicting negligible housing purchase activity during retirement.

Finally, I examine whether intergenerational wealth-transfer taxation may help explain these patterns. High estate- and inheritance-tax states exhibit a significantly lower share of older borrowers and a higher share of middle-aged borrowers among observed mortgage-financed purchases, while no statistically significant association emerges between transfer-tax status and property values among older purchasers. Although these results do not establish a causal mechanism, they are consistent with the possibility that fiscal incentives influence housing decisions beyond the savings accumulation channel emphasized in standard models. The findings highlight the importance of evaluating life-cycle housing models against transaction-level evidence, in addition to aggregate ownership and wealth distributions.

Keywords: life-cycle model, wealth accumulation, homeownership

JEL Codes: D15, G11, R21

*Georgetown University, ap1863@georgetown.edu

1 Introduction

Housing is the largest asset held by most households and occupies a significant role in theories of household finance and wealth accumulation. In canonical life-cycle models, households gradually accumulate savings during the early stages of life, use these savings to overcome borrowing constraints and purchase their first home, upgrade to larger or more valuable properties as income and wealth increase, and eventually size down their housing later in life to extract the wealth and to support consumption. This "housing ladder" mechanism, rooted in the life-cycle consumption smoothing of Modigliani and Brumberg (1954), underlies much of the modern quantitative literature on housing demand, mortgage borrowing, and household portfolios. Whether in models emphasizing borrowing constraints, transaction costs, or housing adjustment frictions, the timing of homeownership is fundamentally linked to the gradual accumulation of wealth.

This mechanism carries a clear empirical implication. If savings accumulation is the primary trigger for homeownership, home purchases should be concentrated among relatively young households, when accumulated wealth first becomes sufficient to finance a down payment. Although households may subsequently trade up as their incomes grow, the transactions should occur early in life. This prediction has become sufficiently embedded in both academic research and housing policy that it is rarely tested directly.

This paper documents evidence that challenges a central implication of the housing ladder hypothesis. Using individual-level mortgage originations from the Home Mortgage Disclosure Act (HMDA), covering the U.S. mortgage market from 2018 to 2025, I show that home purchases remain prevalent among older households and, importantly, that the value of housing acquired conditional on purchase increases systematically with borrower age. Older borrowers invest more in owner-occupied housing than any other age group: conditional on obtaining a mortgage to purchase a home, the expected value of the purchased property rises with age.

This age gradient is remarkably stable across years, geographic markets, borrower charac-

teristics, and periods of major economic change. The evidence therefore stands in contrast to the standard life-cycle pattern implied by the housing ladder hypothesis, under which housing expenditure is concentrated among younger and middle-aged households and declines progressively with age as households downsize or reduce their housing consumption. Instead, the data reveal substantial and in terms of purchase value, increasing engagement with the owner-occupied housing market throughout adulthood. In particular, older households do not merely continue to purchase homes at non-negligible rates; conditional on purchasing, they acquire substantially more expensive properties than their younger and middle-aged counterparts.

The first contribution of the paper is therefore empirical. While previous studies have examined age profiles of homeownership rates or housing tenure, this paper documents the age distribution of home purchase decisions directly using national individual-level mortgage originations. The distinction is important. Homeownership rates reflect the cumulative outcome of past purchase and sale decisions, whereas mortgage originations identify the housing transactions themselves. The evidence demonstrates that one of the central empirical predictions of the standard housing ladder is not fully supported in recent U.S. data.

The second contribution is to evaluate this discrepancy using a standard quantitative life-cycle model as a diagnostic benchmark. Rather than proposing a new structural framework, I deliberately give the canonical model its best opportunity to succeed by calibrating it to match conventional aggregate moments, including the age profile of homeownership, the distribution of housing wealth across income groups, aggregate wealth accumulation, and the concentration of wealth in the upper tail of the distribution. Despite matching these standard targets, the model predicts a sharply declining age distribution of home purchases that is inconsistent with the small but persistent participation of older households observed in the data. This failure is not a consequence of parameterization but follows directly from the mechanism that drives homeownership entry in the benchmark model. Because savings accumulation is the only channel through which households relax borrowing constraints,

the model necessarily predicts that home purchases occur predominantly during the earlier stages of the life cycle. The empirical evidence therefore points to a limitation of the standard mechanism rather than to a calibration error.

Finally, the paper provides empirical and quantitative evidence on whether the discrepancy can be reconciled by modifying key components of the benchmark life-cycle framework. In particular, I document evidence that mortgage contracts can substantially reduce households' tax liabilities and may facilitate tax-advantaged transfers of wealth from parents to their children. Incorporating a comparable fiscal incentive into the standard life-cycle model might generate a mechanism capable of reproducing the representation of retired households and their excessive spendings on housing assets among home purchasers observed in the data.

This paper contributes to three principal strands of the literature. First, it advances the literature on life-cycle housing demand by providing a direct empirical test of the central predictions underlying canonical housing ladder models (Ortalo-Magné and Rady, 2006; Banks et al., 2016). Within this framework, the primary investment margin concerns how households optimize the size and timing of housing transitions in response to income and price dynamics (Clark et al., 2003; Morrow-Jones and Wenning, 2005; Banks et al., 2016; Sinai and Souleles, 2005), borrowing constraints (Flavin and Yamashita, 2011), adjustment costs (Henderson and Ioannides, 1983), and tenure risk (Sinai and Souleles, 2003). I provide evidence that fiscal stimulus can introduce substantial investment incentives later in life, without necessitating modifications to these core components of the standard life-cycle model.

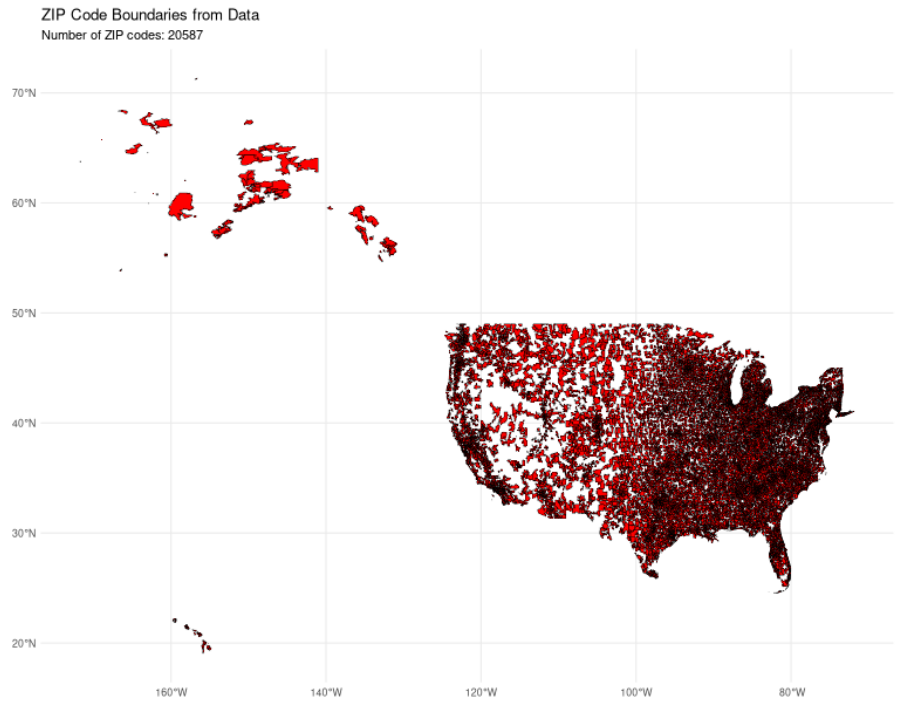
Second, this study contributes to the literature examining the reluctance of older households to decumulate accumulated wealth (De Nardi et al., 2010; De Nardi and Yang, 2014; De Nardi et al., 2016; Lockwood, 2018), particularly housing equity (Venti and Wise, 2001; Nakajima and Telyukova, 2019; French et al., 2022; Aydilek, 2025), as well as the literature documenting the accumulation of debt during the later stages of life (Lusardi et al., 2018,

2020). I present robust empirical evidence that seniors do not merely age in place; they also actively upsize their housing consumption using leverage, demonstrating a pronounced willingness to incur significant mortgage debt to acquire larger properties. And while a recent paper by Collins et al. (2018) establishes the rising prevalence of outstanding mortgage debt among older households, their stock-based measures cannot readily distinguish between the passive retention of earlier debt, refinancing, and new property acquisitions. Hence, their main conclusion is that the rising mortgage debts among seniors is due to earlier debt and refinancing. Using loan-level HMDA records, this paper isolates the flow of mortgage financed new home purchases. I show that older Americans are not merely aging in place with existing debt; some remain active participants in the housing and mortgage market and, conditional on purchasing, acquire higher value properties even within narrow geographic areas. Furthermore, I demonstrate that a calibrated life-cycle housing model failing sharply to capture these late-life transaction flows. This demonstrates that matching housing stocks is insufficient for validating life-cycle models and pointing to missing late-life triggers for housing adjustment and leverage. I rationalize this novel pattern through a tax-based mechanism previously unidentified in the literature.

Third, this paper offers a novel resolution to the reverse mortgage puzzle (Davidoff, 2015; Davidoff et al., 2017; Nakajima and Telyukova, 2017). My findings demonstrate that older households rather seek traditional forward mortgages and are ready to invest large amounts of savings into housing. Consequently, the findings suggest that the persistently low demand for reverse mortgages is not driven exclusively by bequest motives, precautionary savings, or product complexity. Instead, reverse mortgages fundamentally contradict the actual debt and housing market objectives pursued by senior households.

The remainder of the paper is structured as follows. Section 2 describes the data. Section 3 documents the empirical facts. Section 4 uses a standard life-cycle model as a diagnostic benchmark to evaluate the empirical evidence. Section 5 examines competing explanations for the observed patterns. Section 6 concludes.

Figure 1: Core HMDA Dataset by ZIP-code



2 Data

The core dataset is constructed from individual mortgage records in the Home Mortgage Disclosure Act (HMDA) database. I use HMDA observations for 2018-2025. The dataset is based on mortgage applications that were approved and resulted in an originated loan. I restrict the sample to home purchase loans and fixed-rate mortgages. I exclude observations with non-positive reported income and remove the top 0.25 percent of the income distribution. The resulting data contain information on household income, property value, loan amount, loan-to-value ratio, mortgage term and interest rate, debt-to-income ratio, mortgage costs, borrower characteristics, and selected characteristics of the census tract in which the property is located. Based on census tracts I match the relevant ZIP-code. Figure 1 shows the geographic distribution of the resulting mortgage observations by ZIP-code. Detailed data construction and matching procedures are described in Appendix C.

It is also worth to note that before 2018, the HMDA did not report the market value of properties associated with mortgage applications. While loan amounts were available, they

do not fully reflect property values due to the absent information about down payments. From 2018 onward the dataset includes reported property values, which are essential for this research, as well as loan-to-value ratios, debt-to-income ratios and some additional characteristics.

3 Data Facts

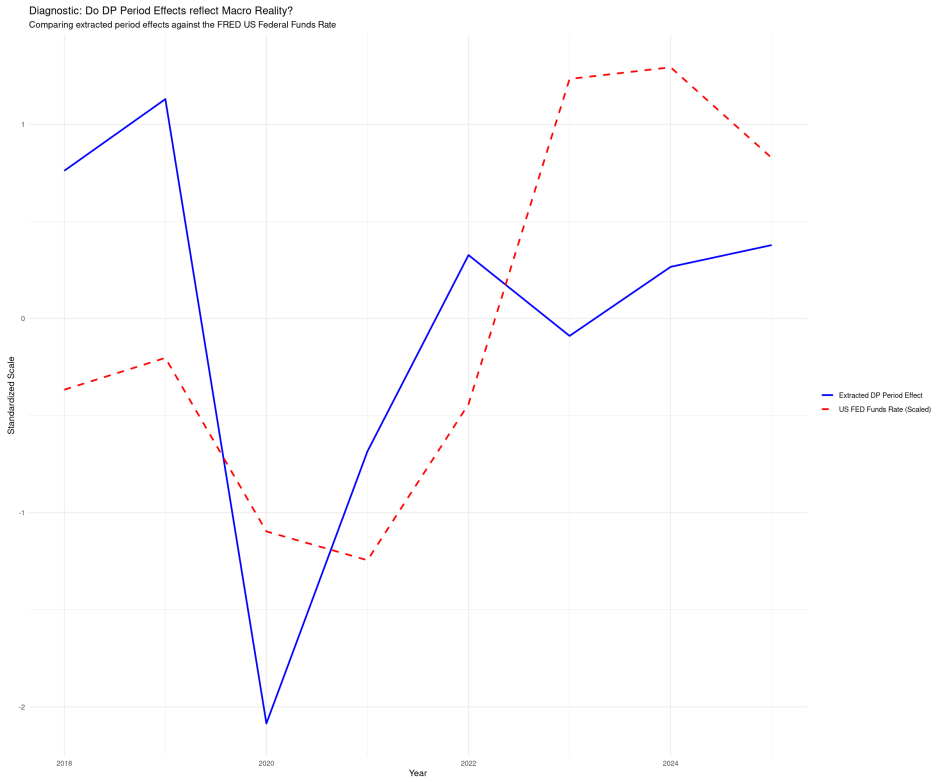
3.1 Age Effects in House Purchase

The raw evidence from HMDA data for 2018-2025 reveals a declining age distribution of home purchases, broadly consistent with the conventional life-cycle pattern. Nevertheless, a notable feature of the data is the persistent presence of older households, defined here as borrowers aged 55 and above, in the mortgage and owner-occupied housing markets. Moreover, conditional on purchasing a home, these households make substantially large housing expenditures, despite having lower income and more debts than other age groups. The corresponding figures and tables are reported in Appendix A.

This persistent presence of senior households among house buyers reveals that raw cross-sectional data may contradict standard underlying life-cycle dynamics. To isolate the true life-cycle profile of housing consumption, I decompose these observations into distinct age, period, and cohort effects.

The fundamental challenge in isolating the effect of age on housing consumption is the perfect linear dependence between age, period (year), and cohort (the birth year), as age equals current year minus birth year. Because of this exact collinearity, a regression cannot independently estimate unconstrained dummy variables for all three dimensions. To resolve this, I employ the normalization technique as in Deaton and Paxson (1994). This approach assumes that macroeconomic period effects are orthogonal to long-term deterministic trends and sum to zero over time. By projecting out the linear time trend, I construct a set of purged year regressors. This restriction breaks the collinearity, forcing any time trends into

Figure 2: Diagnosis of the Identification Assumptions



the age and cohort coefficients while allowing the purged year dummies to capture annual macroeconomic shocks. I validate this assumption by mapping the extracted period effects against the annual average US Federal Funds Rate [1]. The Figure 2 demonstrates that the constrained period coefficients capture the fluctuations accurately enough, considering the short span of years.

As an additional validation of the normalization Deaton and Paxson (1994), I re-estimate the age-period-cohort model by replacing the projected period dummies with the annual average US Federal Funds Rate. Because the Federal Funds Rate is not perfectly collinear with age and cohort, it breaks the exact age-period-cohort linear dependency without imposing the zero-sum or orthogonal restriction.

Using this age-period-cohort framework, I estimate life-cycle profiles of housing demand along two distinct margins:

[1] FRED, FEDFUNDS at annual frequency

- the **extensive margin**: the probability of purchasing a home at a given age (home-ownership entry); and
- the **intensive margin**: the property value conditional on purchase, capturing the scale of housing consumption.

3.1.1 The Extensive Margin (Propensity to Buy)

To determine how the likelihood of purchasing a home changes over the life cycle, I aggregate the mortgage data by age group and year, merging it with Census population estimates for the corresponding age bins and years. The dependent variable is the natural logarithm of the purchase propensity (the share of the population in a given age bin that purchases a home in a given year). To account for varying demographic cell sizes, I estimate a Weighted Least Squares (WLS) model, weighted by the population of each age-year group:

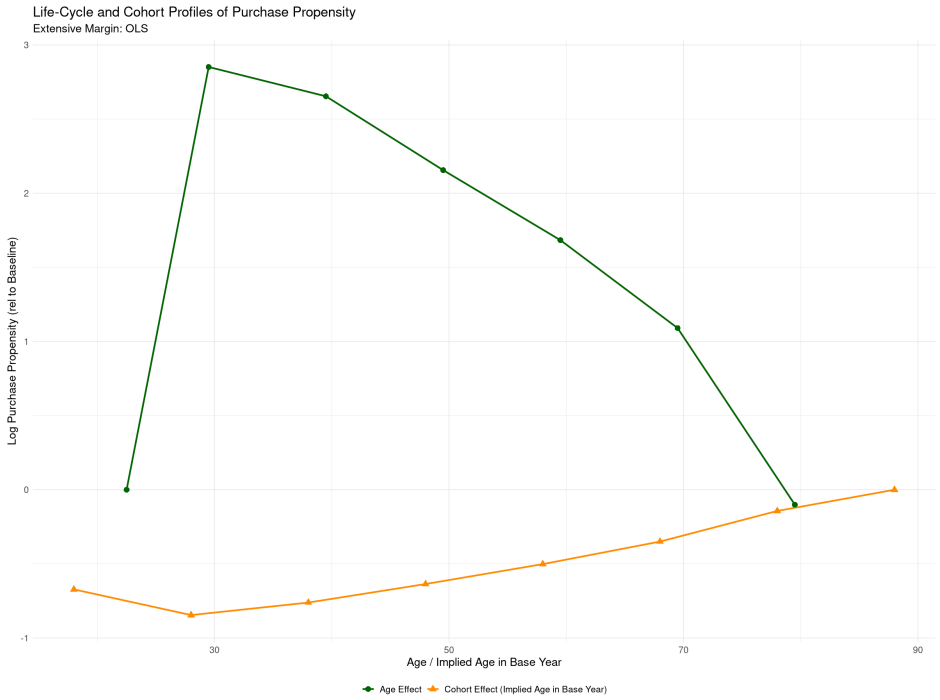
$$\ln(\text{BuyerShare}_{AgeBin,t}) = \beta_0 + \alpha\text{AgeBin} + \gamma\text{Cohort}_{AgeBin,t} + \sum_{k=1}^{T-2} \delta_k \text{PurgedYear}_{k,t} + \epsilon_{AgeBin,t},$$

where αAgeBin denotes age effects for seven categorical age bins (Under 25, 25-34, 35-44, 45-54, 55-64, 65-74, and Over 74); $\gamma\text{Cohort}_{AgeBin,t}$ represents birth cohort effects grouped into 10-year decades based on the age midpoints; $\text{PurgedYear}_{k,t}$ are Deaton-Paxson purged year regressors, which capture non-linear time trends. Finally, standard errors are computed using a heteroskedasticity-robust covariance matrix.

The regression results are presented in Table 1, and Figure 3 displays the estimated age and cohort effects. Although the inverted-U shape in the probability across the age groups is consistent with expectations, the upward cohort trend persists among older age groups, including "55-64", "65-74", and "Over 74", which comprise retired and senior households.

As a robust alternative to the log-linear specification, I employ a Poisson Pseudo-Maximum Likelihood (PPML) estimator. This approach handles any zero-value observation in the buyer share, which potentially can be observed. Moreover, if we allow for heteroskedasticity,

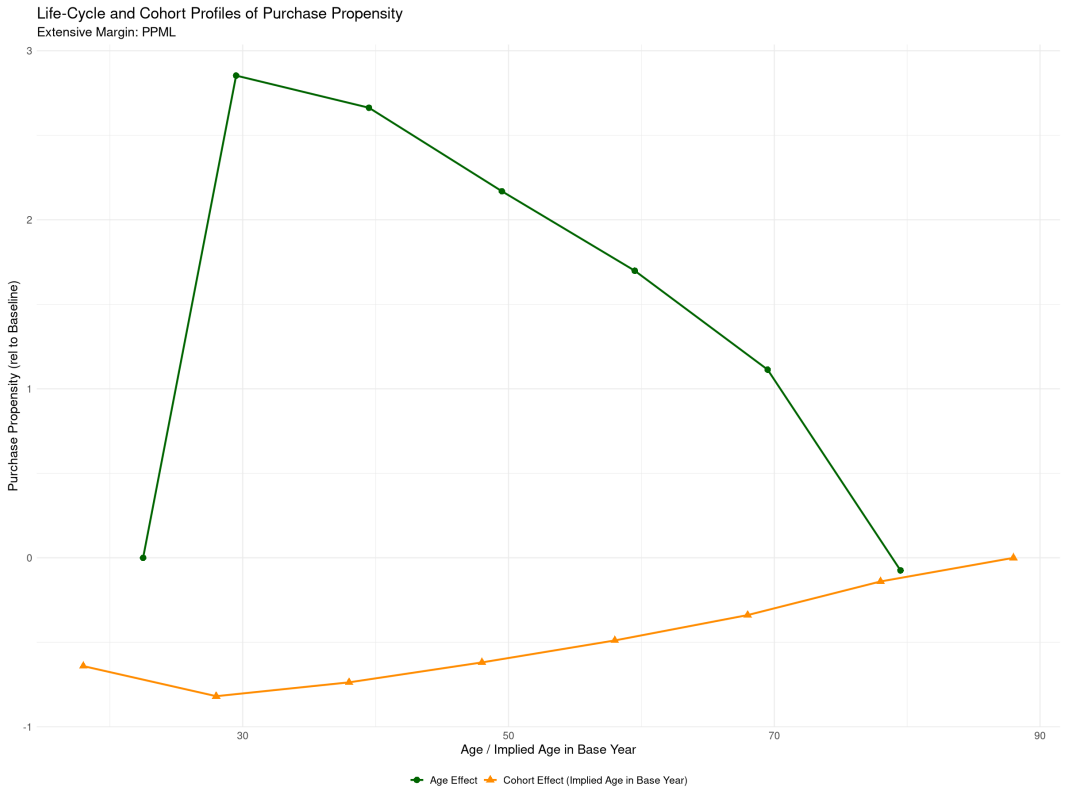
Figure 3: Age and Cohort Effects for Extensive Margin (OLS Model)



then OLS estimator might deliver biased and inconsistent estimation because of the logarithmic transaformation. To prevent this, I provide PPLM model as a robustness check. The results are reported in Table 1 and remain broadly consistent with those obtained from the baseline log-linear model. Figure 4 presents the corresponding estimated age and cohort effects. Notably, the unexpected upward trend in the estimated cohort effect among older households persists.

As a further robustness check, I re-estimate the PPML model by replacing the Deaton-Paxson purged period dummies with the annual average US Federal Funds Rate. The estimated age and cohort profiles, reported in Table 2 and depicted in Figure 5. The results are not significantly altered: now the upward trend of the cohort switches to the age effect. Thus, senior hoseholds are expected to have a probability to buy a house as a positive value - or at least, higher than the base age group. Hence, these results virtually identical to the baseline specification: they both qualitatively and quantitatively confirm that the results are not driven by the particular normalization used to identify period effects. However, they

Figure 4: Age and Cohort Effects for Extensive Margin (PPLM Model)



do attribute non-zero probability of senior households to buy a house to different life-cycle variables.

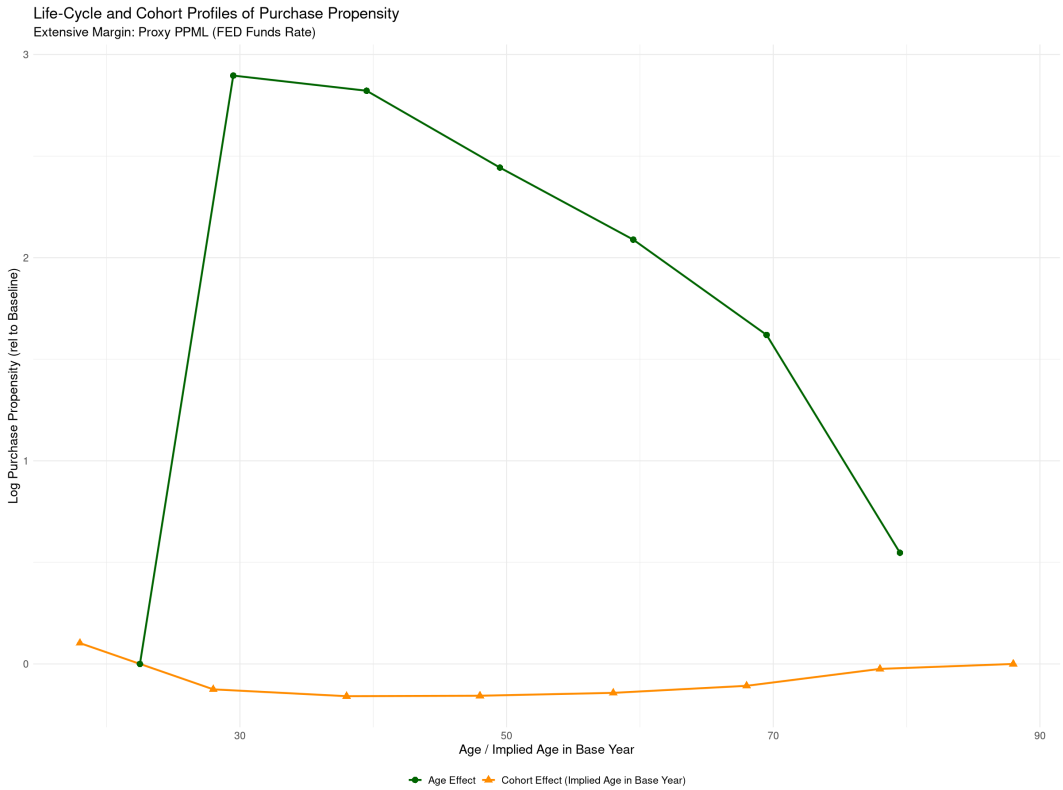
3.1.2 The Intensive Margin (Property Value)

Conditional on purchasing a home, I also examine how the value of the property varies across the lifecycle. I estimate fixed-effects log-linear model at the individual level:

$$\ln(\text{PropertyValue}_{i,c,t}) = \alpha \text{AgeBin}_i + \gamma \text{Cohort}_i + \sum_{k=1}^{T-2} \delta_k \text{PurgedYear}_{k,t} + \mu_c + \epsilon_{i,c,t},$$

where αAgeBin_i represents a set of fixed effects for the seven applicant's age group; γCohort_i captures fixed effects for the applicant's birth cohort grouped into 10-year decades; $\text{PurgedYear}_{k,t}$ are Deaton-Paxson purged year regressors, which capture non-linear time trends; μ_c absorbs unobserved, time-invariant geographic characteristics at the county level, with the error term

Figure 5: Age and Cohort Effects for Extensive Margin (PPLM Model)



$(\epsilon_{i,c,t})$ clustered by county. The results are reported in Table 1 and Figure 6 presents the estimated age and cohort effects. Although the cohort effect follows the expected pattern, the age effect exhibits a strictly increasing trajectory across the age groups. This finding is strikingly at odds with the standard life-cycle hypothesis, which predicts that the corresponding effect should flatten or decline after approximately age 60-65 rather than continue to increase.

As a further robustness check for the intensive margin, I re-estimate the age-period-cohort model substituting the Deaton-Paxson purged period dummies with the annual average US Federal Funds Rate. The estimated age and cohort profiles, reported in Table 2 and plotted in Figure 7, are virtually indistinguishable from the baseline specification and the qualitative patterns are preserved.

The primary finding is a robust, increasing age effect on housing expenditures, conditional on buying a new house. Two additional results worth to emphasise. First, appendices A.2-

Figure 6: Age and Cohort Effects for Intensive Margin

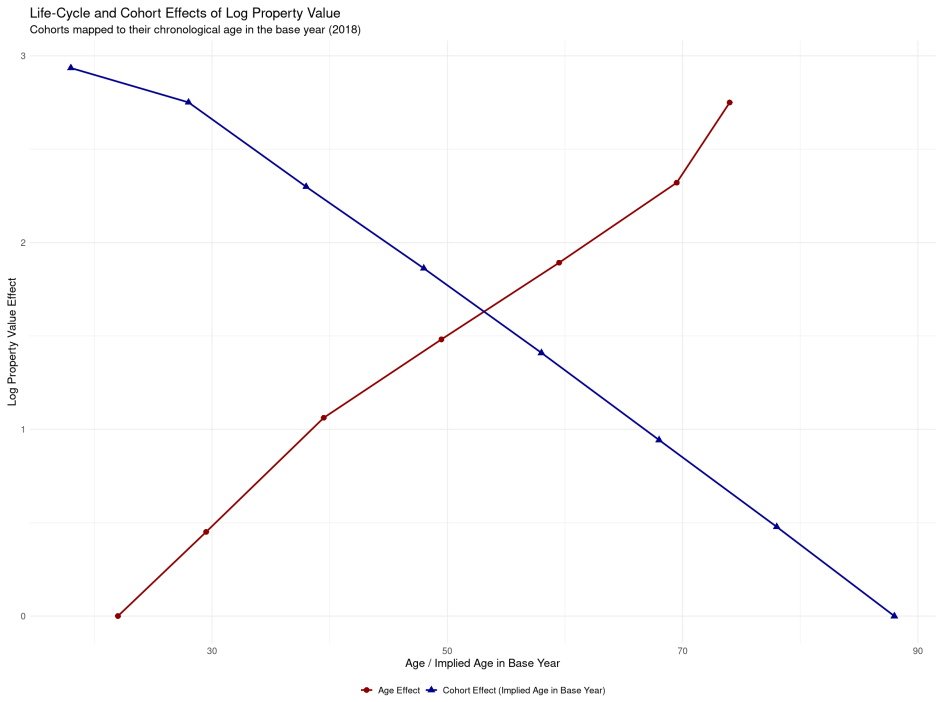


Figure 7: Age and Cohort Effects for Intensive Margin

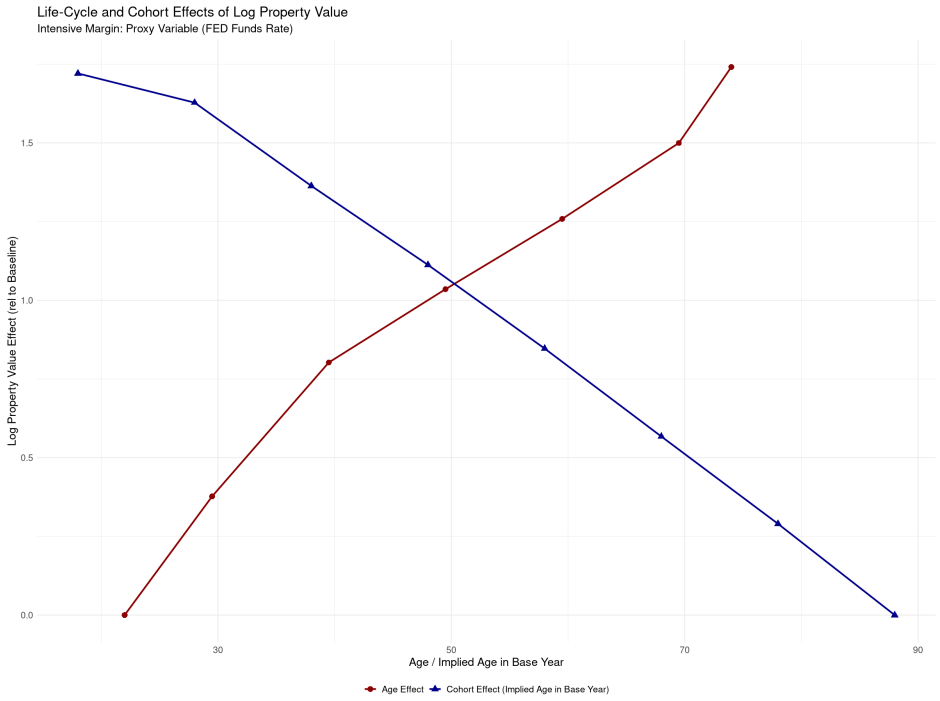


Table 1: Age-Period-Cohort Effects via Deaton-Paxson Normalization Across Intensive and Extensive Margins

Dependent Variables:	Log property value	Log buyer share of age bin	Buyer share of age bin
	OLS	OLS	Poisson
Age Effects (Base: < 25)			
Age 25-34	0.4508***	2.852***	2.853***
	(0.0030)	(0.0206)	(0.0028)
Age 35-44	1.062***	2.654***	2.663***
	(0.0084)	(0.0494)	(0.0068)
Age 45-54	1.482***	2.156***	2.169***
	(0.0153)	(0.0758)	(0.0113)
Age 55-64	1.892***	1.684***	1.699***
	(0.0215)	(0.1021)	(0.0159)
Age 65-74	2.321***	1.090***	1.113***
	(0.0261)	(0.1318)	(0.0206)
Over 74	2.750***	-0.1009	-0.0744***
	(0.0305)	(0.1607)	(0.0255)
Cohort Effects (Base: 1930)			
Cohort 1940	0.4778***	-0.1430***	-0.1398***
	(0.0062)	(0.0350)	(0.0085)
Cohort 1950	0.9428***	-0.3496***	-0.3386***
	(0.0115)	(0.0652)	(0.0122)
Cohort 1960	1.409***	-0.5016***	-0.4887***
	(0.0170)	(0.0909)	(0.0161)
Cohort 1970	1.862***	-0.6360***	-0.6187***
	(0.0231)	(0.1173)	(0.0203)
Cohort 1980	2.299***	-0.7610***	-0.7366***
	(0.0290)	(0.1449)	(0.0246)
Cohort 1990	2.751***	-0.8455***	-0.8185***
	(0.0351)	(0.1722)	(0.0289)
Cohort 2000	2.935***	-0.6723***	-0.6405***
	(0.0367)	(0.1886)	(0.0308)
Constant		-6.692***	-6.719***
		(0.1743)	(0.0285)
County Code	Yes	No	No
Observations	10,727,207	49	49
R^2	0.43078	0.99971	
Pseudo R^2	0.27590	1.0000	0.15013

Note: *p<0.1; **p<0.05; ***p<0.01

Table 2: Age-Period-Cohort Effects: Proxy Variable Approach (Fed Funds Rate)

Dependent Variables:	Log property value OLS	Buyer share of an age bin Poisson
Age Effects (Base: < 25)		
Age 25-34	0.3772***	2.896***
	(0.0030)	(0.0420)
Age 35-44	0.8025***	2.822***
	(0.0058)	(0.0555)
Age 45-54	1.035***	2.443***
	(0.0101)	(0.0663)
Age 55-64	1.259***	2.089***
	(0.0142)	(0.0756)
Age 65-74	1.500***	1.620***
	(0.0167)	(0.0841)
Over 74	1.741***	0.5472***
	(0.0188)	(0.0910)
Cohort Effects (Base: 1930)		
Cohort 1940	0.2899***	-0.0243
	(0.0046)	(0.0459)
Cohort 1950	0.5673***	-0.1076*
	(0.0077)	(0.0573)
Cohort 1960	0.8467***	-0.1423**
	(0.0109)	(0.0671)
Cohort 1970	1.112***	-0.1565**
	(0.0147)	(0.0776)
Cohort 1980	1.363***	-0.1587*
	(0.0184)	(0.0861)
Cohort 1990	1.628***	-0.1251
	(0.0221)	(0.0898)
Cohort 2000	1.721***	0.1026
	(0.0226)	(0.1066)
Proxy Variable		
Fed Funds Rate	0.0427***	-0.0964***
	(0.0007)	(0.0045)
Constant		-7.233***
		(0.0954)
County Code	Yes	No
Observations	10,727,207	49
R^2	0.42644	
Pseudo R^2	0.27219	0.15003

Note: *p<0.1; **p<0.05; ***p<0.01

A.3 investigate two additional mechanisms that could potentially account for the observed age patterns: heterogeneous borrowing constraints and differences in local housing market conditions. Neither mechanism, however, accounts for the estimated age gradient. The age effects remain statistically and economically stable after controlling for heterogeneity in credit constraints and geographic location, suggesting that the observed relationship between borrower age and housing expenditure is not driven by differences in access to credit or by the spatial distribution of housing purchases.

Second, I replicate the age-period-cohort estimation on the PSID (1968-2023) using a two-way fixed-effects panel model with household and year fixed effects. For housing assets held (self-reported house values at interview), no increasing age profile emerges. Crucially, the PSID measures housing wealth, not leveraged purchase decisions: it captures the stock of housing assets at prevailing values, not the flow of new mortgage-financed acquisitions. The absence of an increasing age gradient in housing wealth alongside its pronounced presence in mortgage borrowing underscores the central role of leverage in generating the observed life-cycle pattern.

Formal estimation details for both results are reported in Appendices A.3 and A.4.

4 Diagnosing the Standard Model: What Optimal Life-Cycle Behavior Predicts

4.1 Lifecycle Model with Housing

To interpret the two empirical facts documented in Section 3, I use a standard discrete time finite life-cycle model of household consumption and housing choice as a diagnostic benchmark. The goal of this section is not to propose a new theory of housing demand, but to ask a prior question: can the canonical framework, carefully implemented and calibrated to match well-established aggregate features of the U.S. economy, account for the facts in

the data? In most parts I follow Kaplan and Violante (2014), but allow liquid assets to be negative, treating them as net wealth rather than a specific liquid asset type. In addition, the model is cast in partial equilibrium. House prices and rental rates are taken as exogenous, consistent with the model's diagnostic role: I ask whether optimal household behavior given observed market conditions can replicate the empirical facts, not whether those market conditions themselves are endogenously determined.

Each household i lives for T periods. Each one also receives an individual labor income y_t^i. Labor income of an individual i at age t is given by $y_t^i = w \cdot \kappa_t \cdot \alpha_i \cdot z \cdot e$, where each component captures a distinct source of variation. w is the aggregate wage rate, normalized to 1; κ_t is a deterministic age-efficiency profile; α_i is a fixed, permanent productivity type, and there are 5 of them in the economy; z is a persistent idiosyncratic productivity shock; e is a transitory i.i.d. productivity shock. In every period t, households can allocate their wealth between two types of assets: (1) a risk-free savings asset a_{t+1}^i and (2) a risky housing asset h_{t+1}^i. The risk-free asset a_{t+1}^i is purchased at a normalized price, similar to the consumption good, and yields a uniform real rate of return $(1 + r_t^a)$ across the nation in the next period. The house service unit h_{t+1}^i can be purchased each month at a p_t^h price. In this way, the house purchase will return $1 + z_{t+1} = \frac{p_{t+1}^h + rent_{t+1}}{p_t^h}$. Each household also chooses each period how much of consumption good c_t^i to consume, and how much of shelter services s_t^i to consume.

The household's preferences follow a standard additively separable expected utility framework with a discount factor β. The objective function is given by: $\max_{\{c_t^i, s_t^i, a_{t+1}^i, h_{t+1}^i\}_{t=0}^T} \mathbb{E} \sum_{t=0}^T \beta^t U(c_t^i, s_t^i)$, where c_t^i represents consumed units of consumption good, s_t^i represents consumed units of shelter services and h_{t+1}^i is the chosen housing services, which transforms into shelter services next period. That is, there is a linear technology, which transforms owned housing services h_t^i into shelter services s_t^i in each period t. Although homeowners and renters share the same objective function and have the same universal borrowing limit b, their budget constraints differ according to the state variable h_t^i. If no h_t^i is owned, then shelter services should be rented for a price which constitutes a fixed portion ψ of the income. That is, renters who

have $h_t^i = 0$ will have to pay the rent for consumed shelter services. If a household owns more h_t^i than it demands shelter services s_t^i, then the difference is rented out and becomes an earned rent. If any household decides to change the value of the housing services owned, it will pay housing transaction costs htc_t, which are fixed share ϕ of a purchased (or sold) house value.

Problem 1 *Household i's problem:*

$$\max_{\{c_t^i, s_t^i, a_{t+1}^i, h_{t+1}^i\}_{t=0}^T} \mathbb{E} \sum_{t=0}^T \beta^t U(c_t^i, s_t^i),$$
$$s.t.\ c_t^i + rent_t s_t^i + p_t^h h_{t+1}^i + a_{t+1}^i = y_t^i + rent_t h_t^i + p_t^h h_t^i - htc_t + (1 + r_t^a) a_t^i,$$
$$a_{t+1}^i \geq -b * p_t^h h_{t+1}^i.$$

$$rent_t = \psi^i y_t^i$$

$$htc_t = \begin{cases} \phi(p_t^h h_{t+1} + p_t^h h_t) & \text{if } h_{t+1} \neq h_t \\ 0 & \text{otherwise} \end{cases}$$

Assumption 1 *Assume the following functional form of the utility function:* $U(c_t^i, s_t^i) = \frac{\left((c_t^i)^{1-\sigma_h} (s_t^i)^{\sigma_h}\right)^{1-\sigma}}{1-\sigma}$.

See the discussion of the partial eqilibrium in Appendix B.

4.2 Parametrization and Calibration

Before asking what the model predicts about transaction behavior, I calibrate it to match standard aggregate features of the U.S. housing market and the U.S. economy. This step is important for the diagnostic exercise: if the model were miscalibrated, its failure to match the empirical facts could reflect poor parameterization rather than a fundamental limitation of the framework. By showing that the model matches the aggregate targets well, I ensure that the subsequent failures are informative about the model's structural assumptions rather

than its parameter values. The calibration targets are the age profile of homeownership rates, the distribution of housing wealth across income quintiles, the aggregate wealth-to-income ratio, and the top decile wealth share.

4.2.1 Parameters: Income Process

Households live for 81 periods, corresponding to ages 20 through 100. They work until age 68 and receive labor income during this phase. After retirement, they receive pension, modeled as a constant fraction of pre-retirement earnings. Survival probabilities are defined using age-dependent mortality rates derived from life tables of CDC[2]. Conditional survival probabilities s_j govern the likelihood of transitioning between periods, with death occurring with certainty in the final period.

Labor income follows a deterministic life-cycle profile governed by age-dependent productivity κ_j. It is constructed to match a hump-shaped earnings profile: a steep increase early in life, a plateau during prime working years, and a decline approaching retirement. Full retirement occurs at $T_r = 68$, after which income is constant at $\kappa_t = 0.65$, capturing pension income. So, after retirement at 68, income drops to a constant level equal to approximately 65% of base earnings. During the period of life between the year 20 and 68 individuals also receive idiosyncratic permanent and transitory income shocks each period. Permanent component is AR(1) process with 0.95 persistence parameter and 0.1 innovation parameter (Meghir and Pistaferri, 2004; Storesletten et al., 2004). Transitory component is iid process with mean zero and standard deviation 0.15 (Blundell et al., 2008). In addition, agents differ by a fixed productivity type α_i, drawn at birth from a discrete distribution: exp(1.5),exp(0.5),exp(0),exp(-0.5),exp(-1.2), with equal probabilities. The fixed productivity multipliers and shares of the population are defined to match the data of the Congressional Budget Office[3].

[2]National Vital Statistics Report, volume 74, number 6, July 2025

[3]The Distribution of Household Income, 2019 (Data and Supplemental Information)

4.2.2 Parameters: Preferences

The discount factor is set to $\beta = 0.96$. Preferences exhibit constant relative risk aversion with coefficient $\sigma = 3$. Utility is defined over consumption and housing services, with an intratemporal elasticity of substitution governed by $\sigma_h = 0.45$.

4.2.3 Parameters: Housing Market

Housing transactions are subject to proportional adjustment costs. The transaction cost parameter is set to $\phi = 0.05$, applied symmetrically to purchases and sales.

Households face a collateral constraint that limits borrowing to a fraction $b = 0.8$ of the housing value, effectively imposing a minimum down payment requirement of 20%.

Rental prices vary across income groups and the values of the income shares are taken from the "Assessing the Severity of Rent Burden on Low-Income Families" (Board of Governors of the Federal Reserve System, 2017). The rent-to-income ratio range from 10% for high-income households to 56% for low-income households, reflecting the larger burden of rental costs for poorer families.

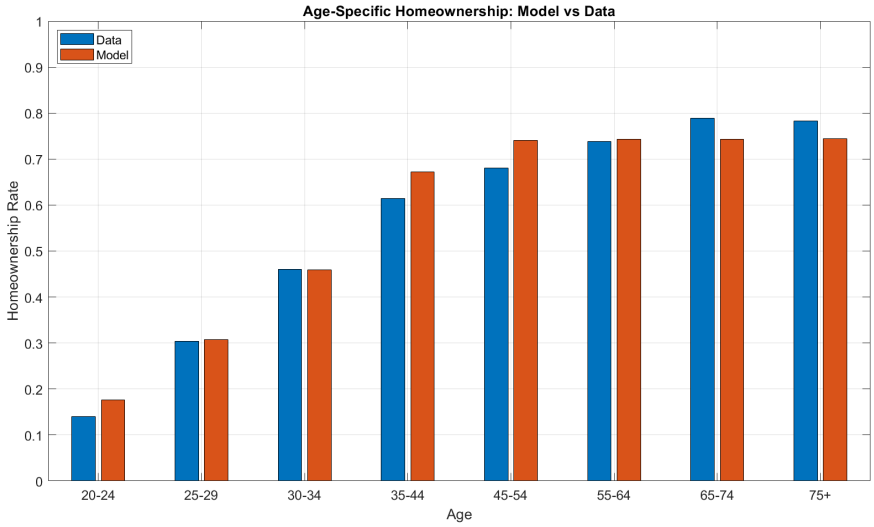
House prices follow a stochastic AR(1) process with persistence $\rho_{ph} = 0.9$ and innovation standard deviation $\sigma_\epsilon = 0.04$. The process is discretized using the Farmer-Toda method, and the resulting grid is normalized to have near unit mean.

4.2.4 Parameters: Financial Market

The risk-free interest rate is set to $r = 0.04$ to approximately match the historical yield on U.S. Treasury securities. Mortgage borrowing is priced at the same rate, implying that mortgage debt is effectively risk-free in the model.

Households can hold both liquid assets and housing wealth. Liquid asset holdings are allowed to be negative, subject to the collateral constraint tied to housing.

Figure 8: Homeownership by Age



4.2.5 Targeted External Moments

With these parameters, I simulate the model with 1,000 households per income type over 1,000 periods (exluding first 200 burned-in periods) and assess whether it matches the following statistics of the U.S. economy.

First, I match the age-specific homeownership rate. The statistics is taken from the American Housing Survey 2023. The match is presented in Table 3. While the model slightly underestimates homeownership, it follows the main pattern across ages (Figure 8).

Table 3: Model Simulation, Homeownership by Age

Age Group	Data	Model
Total	0.6519	0.5756
Under 25	0.1400	0.1758
25-29	0.3036	0.3076
30-34	0.4601	0.4593
35-44	0.6149	0.6718
45-54	0.6809	0.7408
55-64	0.7384	0.7425
65-74	0.7888	0.7429
75+	0.7827	0.7439

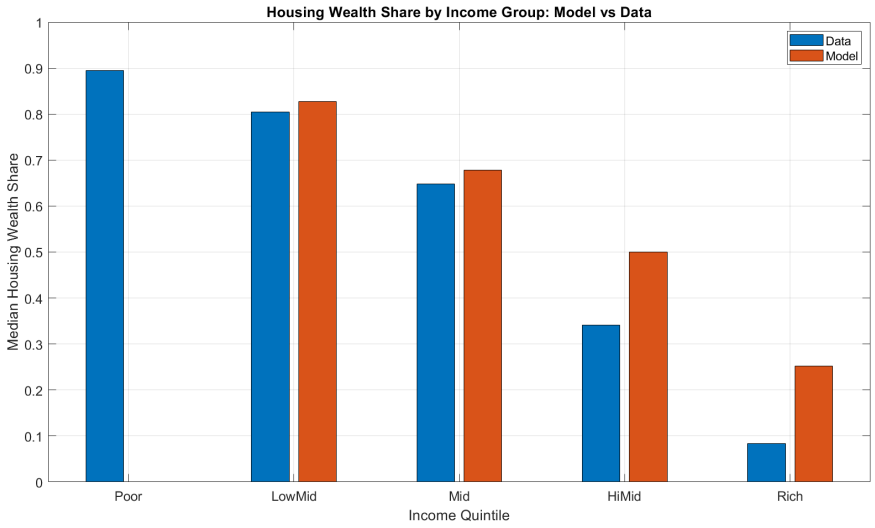
Second, I target the median housing wealth share of homeowners by income quintile,

computed from the Survey of Consumer Finances (Board of Governors of the Federal Reserve System, 2023). The housing wealth share is defined as the ratio of house value to total net worth. This share declines steeply with income: from 0.895 for the lowest quintile down to 0.083 for the highest. This gradient reflects the fact that lower-income homeowners hold almost all of their wealth in housing, while high-income households are more diversified into financial assets. The model targets this gradient, and the match is presented in Table 4 and Figure 9. Two systematic deviations are worth addressing. First, at the top of the income distribution, the model cannot replicate the extremely low housing share of 0.083 for the richest quintile. This occurs because the model assigns a single income level to each type, whereas the data's top quintile aggregates a wide range of incomes, including ultra-high earners whose financial wealth dwarfs their housing holdings. These extreme observations pull the quintile's mean and median housing share sharply downward in the data. Since the model abstracts from within income heterogeneity and from the right tail of the earnings distribution, it cannot generate sufficient financial wealth concentration. Second, at the bottom of the distribution, the poorest quintile presents zero homeowning observations in the model. This is an intentional feature rather than a failure: the model's poorest income type faces a binding borrowing constraint throughout most of the lifecycle, and absent any intergenerational housing transfers (households enter the lifecycle with zero housing assets and cannot inherit property from dying generations) these agents never accumulate sufficient wealth to meet the down payment requirement. In real life, the few low-income homeowners who do own tend to have acquired their home through inheritance or long tenure, and their net worth outside the house is negligible, producing the very high observed share of 0.895. Despite these two deviations, the model successfully reproduces the central quantitative feature of the data: a monotonically declining housing wealth share across the income distribution. The housing share for the richest type falls substantially relative to the middle types, declining from 0.83 to 0.25, while middle-income homeowners still hold the bulk of their wealth in housing.

Table 4: Model Simulation, Housing Wealth Share of Homeowners by Quintile

Quintile	Data	Model
Poor	0.895	NA
Lowmid	0.804	0.8278
Mid	0.648	0.6786
Himid	0.341	0.4998
Rich	0.083	0.2517

Figure 9: Housing Wealth Share by Quintile



Finally, I target two aggregate wealth distribution moments. The first is the aggregate wealth-to-income ratio, constructed as the ratio of averaged across the years median family net worth to averaged across the years median family income. I average across the SCF survey waves from 1989 to 2022 (Board of Governors of the Federal Reserve System, 2023) to abstract from cyclical fluctuations and get a stable economy-wide target. The data value is 2.0699 and the model delivers 2.1469. The second moment is the top 10% wealth share, taken from "The State of U.S. Household Wealth" report as of the fourth quarter of 2024 (Federal Reserve Bank of St. Louis, 2025). The data reports a share of 67.2%, while the model delivers 60.27%. This deviation is a known and expected limitation of the model's income process. The model assigns a single permanent income level to each of five household types, which by construction cannot capture the extreme right tail of the wealth distribution.

Importantly, the model does reproduce the broad concentration of wealth: nearly two-thirds of aggregate net worth accrues to the top decile in both the model and the data, and the qualitative structure of wealth inequality is preserved.

Having confirmed that the model provides a credible enough representation of the aggregate US housing market, I now subject it to a more demanding test: whether it can account for the transaction-level patterns which motivate this paper. As the next section shows, it cannot and the structure of its failure is the paper's central contribution.

5 Model Prediction and Its Failure

Having established that the benchmark model matches standard aggregate targets, I now ask whether it can reproduce the empirical facts documented in Section 3. The answer is no. The model's prediction of a sharply skewed age distribution of purchases with no purchases happening during the retirement against the data's persistent presence of senior households reveals that the standard savings-accumulation trigger for homeownership entry is insufficient to account for observed market participation across the life cycle. This failure points to the specific gap in the standard framework: is it missing triggers for housing demand and leverage during the retirement period.

5.1 Skewed Distribution of Purchases/Mortgages

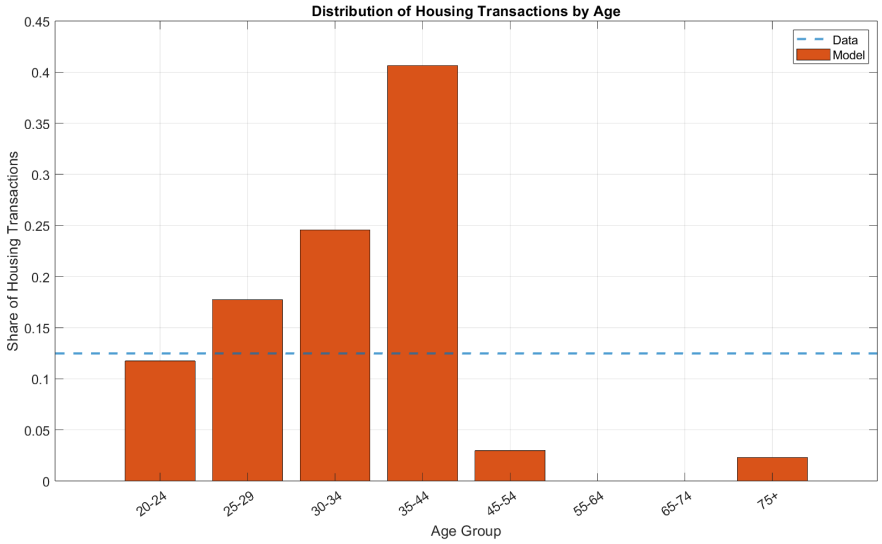
To characterize the age profile of housing market activity in the model, I define a purchase as any positive change in housing assets, excluding transitions from a positive stock to zero in order to isolate acquisitions from liquidations. This construction yields a simulated analogue to the transaction-level records in the HMDA database, where each observation corresponds to a discrete purchase event. I further decompose transactions into three economically distinct categories: entry into homeownership, upgrading an existing owned property, and downsizing. Aggregate results by age group are reported in Table 5; the full age distribution

is illustrated against a uniform line in Figure 10, and the transaction-type decomposition is shown in Figure 11.

The simulated distribution is sharply left-skewed, with purchasing activity heavily concentrated between ages 25 and 44, which together account for approximately 83 percent of all transactions. This pattern is not arbitrary as it reflects the joint operation of two life-cycle forces. The down payment constraint binds hardest in early working life, when asset accumulation is limited, and releases precisely as households reach the steepest phase of the earnings profile. Once the constraint becomes slack, households have both the means and the motive to enter ownership: the remaining working horizon is long enough that the illiquid housing asset can be held to maturity, and the expected stream of housing services justifies the portfolio choice. The result is a purchasing activity that spikes at the age of 30-44 and falls sharply thereafter, as shown in Table 5. The decomposition by transaction type sharpens this picture. Entry into homeownership is front-loaded, peaking in the 25-44 ages and accounting for the bulk of early age activity. Upgrading follows with a slight lag, as households who entered at the minimum feasible house size adjust upward as income and wealth grow. Downsizing is almost entirely absent before age 35 and remains negligible through middle age, before emerging as the dominant transaction type among households aged 75 and above, who account for nearly 90 percent of all downsizing events. This liquidation pattern is consistent with the consumption-smoothing motive: as only pension exits and mortality risk rises, housing wealth is progressively converted into liquid assets to finance consumption and get utility from it before the end of the life-cycle.

Taken together, these patterns confirm that the model's purchase distribution is driven by the interaction between the down payment constraint and life-cycle income dynamics. The concentration of activity in early-to-middle working life, and almost complete absence of transactions among households aged 55-74 stands in contrast to the empirically observed presence, and more importantly, excess intensive investments into housing. This reflects a missing trigger for purchases and missing incentive to purchase at highest prices at the

Figure 10: Age Distribution, All Transactions



retirement ages, to which the model assigns absence of activity. Hence, while we can document that the model cannot match the extensive margin, we cannot say anything about the intensive margin and how the model would treat it.

Table 5: Age Distribution of Housing Transactions by Type

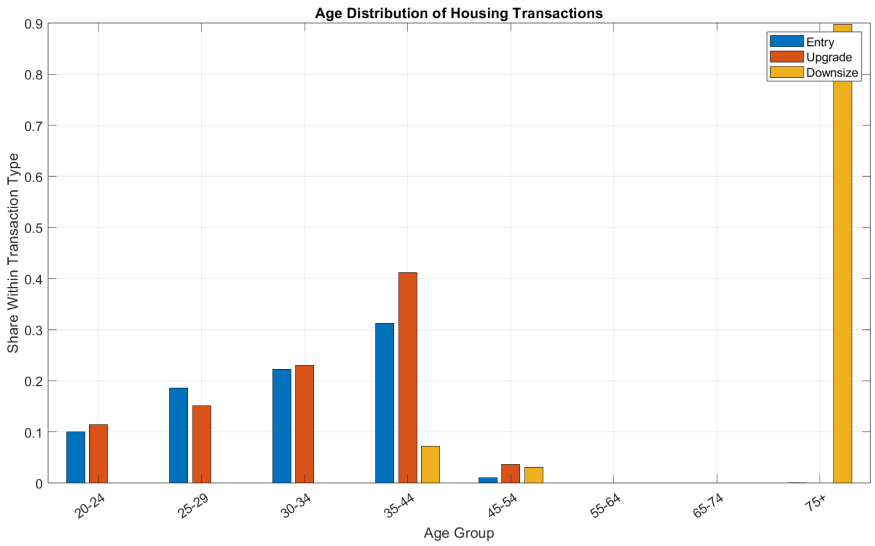
Age Group	All Transactions	Entry	Upgrade	Downsize
20-24	0.1175	0.0999	0.1140	0.0000
25-29	0.1774	0.1861	0.1513	0.0000
30-34	0.2458	0.2228	0.2304	0.0000
35-44	0.4066	0.3123	0.4116	0.0718
45-54	0.0298	0.0107	0.0364	0.0307
55-64	0.0001	0.0000	0.0001	0.0000
65-74	0.0000	0.0000	0.0000	0.0000
75+	0.0229	0.0005	0.0000	0.8975

6 Empirical Evidence for the Fiscal Mechanism

6.1 Estate and Inheritance Taxes

One potential explanation for this pattern is the role of bequest motives and the tax treatment of intergenerational real estate transfers. Although older households may participate

Figure 11: Age Distribution by Transaction Type



less frequently in the mortgage and owner-occupied housing markets, those who do purchase may acquire more valuable properties in anticipation of transferring wealth to their heirs. Under U.S. tax law, inherited real estate generally receives a step-up in basis to its fair market value at the date of the decedent's death, potentially eliminating capital gains tax liability on appreciation accrued during the decedent's lifetime. However, the transfer may still be subject to federal estate taxes or applicable state-level estate or inheritance taxes. This preferential treatment of unrealized capital gains may therefore strengthen the incentive to hold housing as a long-term store of wealth intended for intergenerational transfer. However, there are many more than just capital gains taxes.

To further examine whether taxes associated with the bequest motive influence housing demand, I compare the housing purchases of older and younger households across states with high versus low estate and inheritance taxes. I classify HI, IL, KY, MD, MA, MN, NE, NJ, NY, OR, PA, RI, VT, WA, and DC as high-death-tax jurisdictions because they impose either a top marginal estate-tax rate of at least 16% or a top inheritance-tax rate of at least 15% on at least one class of beneficiaries; all remaining states constitute the low-tax comparison group. The purpose of this classification is to test whether older households with

a bequest motive respond to the tax treatment of intergenerational transfers by allocating more wealth to expensive owner-occupied housing. Under this mechanism, high death-tax exposure may reduce older households' incentive to hold wealth in housing and therefore shift their purchases from housing towards other types of assets. The empirical test accordingly asks whether the housing expenditure and age composition of buyers differ systematically between high- and low-death-tax states, with younger households serving as a benchmark group because bequest-related tax considerations should be less salient for them. State-level tax rules provide plausibly exogenous policy variation relative to individual household characteristics, although the identifying assumption requires that, absent these tax differences, the relative housing choices of older and younger households would follow comparable trends across the two groups of states.

This specification identifies the coefficients in a linear probability model in which purchasing households are classified into three age groups: young (below 35), middle-aged (35-65), and older (above 65). Indicator variables for these categories are regressed on a high or low tax dummy, a set of controls (income, occupancy type, lien status, applicant race, applicant sex and county and year fixed effects), and an interaction term between the tax indicator and income. The estimating equation is given below, with standard errors clustered at the state level:

$$Y_{it}^{(g)} = \beta_0 + \beta_1 \,\text{HighTransferTax}_t + \beta_2 \,\text{Income}_{it} + \beta_3 \,(\text{HighTransferTax}_{ct} \times \text{Income}_{ict}) + \mathbf{X}_{it}'\boldsymbol{\gamma} + \delta_t + \varepsilon_{it},$$

where $Y_{ict}^{(g)}$ is an indicator for buyer i, year t, belonging to age group $g \in \{\text{Young, Middle, Old}\}$, HighTransferTax_t is an indicator for high transfer tax areas, Income_{it} is borrower income, $\mathbf{X}_{it}$ includes occupancy type, lien status, race, and sex controls, δ_t are year fixed effects, ε_{it} is the error term.

The estimates reported in Table 6 yield coefficients with the expected signs and are statistically significant, consistent with the hypothesis that real estate transfer taxes influence housing market participation across age groups. Higher estate or inheritance taxes do not

Table 6: Linear Probability Models of Homebuyer Age Categories

	Dependent variable:			
	Young	Middle	Old	log(Property Value)
	(1)	(2)	(3)	(4)
High Transfer Tax	-0.003	0.043^{***}	-0.040^{***}	0.036
	(0.005)	(0.011)	(0.013)	(0.086)
Income	-0.0002^{***}	0.0003^{***}	-0.0001^{***}	0.003^{***}
	(3.01×10^{-5})	(2.08×10^{-5})	(1.58×10^{-5})	(0.0001)
Occupancy: Second Residence	-0.027^{***}	-0.244^{***}	0.271^{***}	-0.169^{***}
	(0.002)	(0.008)	(0.009)	(0.023)
Occupancy: Investment Property	-0.028^{***}	-0.075^{***}	0.103^{***}	-0.347^{***}
	(0.002)	(0.012)	(0.011)	(0.045)
Subordinate Lien	0.039^{***}	0.029^{***}	-0.067^{***}	0.039^{**}
	(0.005)	(0.008)	(0.009)	(0.019)
Race: American Indian	0.013^{***}	0.037^{***}	-0.051^{***}	-0.134^{***}
	(0.003)	(0.004)	(0.007)	(0.029)
Race: Asian	-0.022^{***}	0.147^{***}	-0.125^{***}	0.273^{***}
	(0.002)	(0.008)	(0.009)	(0.062)
Race: African American	-0.030^{***}	0.048^{***}	-0.017^{**}	-0.185^{***}
	(0.003)	(0.006)	(0.007)	(0.028)
Race: Pacific Islander	-0.011^{***}	0.069^{***}	-0.058^{***}	0.085^{*}
	(0.003)	(0.006)	(0.007)	(0.048)
Female	-0.003^{***}	-0.012^{***}	0.015^{***}	-0.086^{***}
	(0.0010)	(0.002)	(0.002)	(0.007)
High Transfer Tax × Income	1.75×10^{-6}	-1.05×10^{-5}	8.76×10^{-6}	2.52×10^{-5}
	(4.25×10^{-5})	(3.47×10^{-5})	(2.33×10^{-5})	(0.0002)
Observations	9,299,280	9,299,280	9,299,280	1,971,138
R^2	0.01943	0.03402	0.03604	0.33157
Year fixed effects	Yes	Yes	Yes	Yes

Note: *p<0.1; **p<0.05; ***p<0.01

influence younger households' decisions to enter the housing market, while older households may be less likely to purchase properties in jurisdictions with higher bequest-related tax burdens. However, when the analysis is restricted to older households and examines the value of properties purchased, no statistically significant relationship emerges between property values and real estate transfer taxes. One possible explanation is strategic geographic sorting: older households may anticipate the tax implications of intergenerational wealth transfers and avoid purchasing properties in high-tax jurisdictions. Such endogenous location choice could attenuate the estimated relationship between transfer taxes and property values within the sample of older homebuyers, making it difficult to identify the effect. Thus, the absence of a statistically significant relationship among older purchasers does not necessarily rule out a role for transfer-tax considerations; rather, it may reflect selection into the housing market and across jurisdictions.

7 Conclusion

This paper documents an empirical regularity that challenges a central implication of the standard housing ladder hypothesis and canonical life-cycle models of housing demand. Although home purchases are concentrated among younger and middle-aged households, older households remain persistently active in the mortgage financed housing market. Moreover, conditional on purchasing a property, older borrowers acquire substantially more expensive homes than their younger counterparts. These transaction-level patterns reveal a dimension of life-cycle housing behavior that aggregate moments alone may fail to capture.

The quantitative life-cycle model developed in this paper reproduces several aggregate features of the U.S. housing market reasonably well, including the age profile of homeownership rates, the distribution of housing wealth across income quintiles, the aggregate wealth-to-income ratio, and the concentration of wealth at the top of the distribution. However, it fails to replicate the observed age distribution of housing transactions. Simulated

purchases are disproportionately concentrated among younger and middle-aged households, with essentially no participation among older households. This discrepancy reflects the model's reliance on accumulated savings and the relaxation of down-payment constraints as the primary mechanism governing entry into homeownership. When the borrowing constraint ceases to bind for households following similar income trajectories, purchases become concentrated within a relatively narrow age range, generating a pronounced age profile that is inconsistent with the more persistent participation of older households observed in the data.

To investigate whether intergenerational wealth-transfer incentives may help explain this discrepancy, the empirical analysis examines how the composition of homebuyers varies with state-level estate and inheritance tax regimes. Linear probability model estimates reveal that, conditional on borrower characteristics and other controls, obtaining a residence in a high transfer tax state is associated with a lower probability of a mortgage-financed purchase being made by an older household. The estimated association for younger households is statistically insignificant, as it should be. By contrast, among older homebuyers, the estimated relationship between intergenerational transfer tax status and log property value is statistically insignificant. The results therefore suggest that estate and inheritance tax regimes are associated primarily with the age composition of observed housing transactions, rather than with the value of properties acquired by older purchasers. Although these patterns are consistent with the hypothesis that intergenerational transfer considerations influence housing market participation, they do not establish a causal mechanism, and the estimates do not independently identify bequest motives.

Taken together, these findings suggest that models of life-cycle housing demand may place excessive emphasis on savings accumulation and the relaxation of borrowing constraints as determinants of the timing of home purchases. Housing transactions reflect not only households' financial capacity to enter the market but potentially also incentives associated with the intergenerational transfer of wealth and related fiscal incentives. Incorporating

such considerations may help explain why older households continue to participate in the housing market and why their purchases are disproportionately concentrated in higher-value properties. More broadly, the results highlight the importance of evaluating life-cycle housing models against transaction-level age patterns, in addition to aggregate ownership and wealth distributions.

References

Aydilek, A. (2025). The role of health shocks after age 70 on housing and wealth profiles. *PLoS One*, 20(1):e0312349.

Banks, J., Blundell, R., Oldfield, Z., and Smith, J. P. (2016). House price volatility and the housing ladder. In Wise, D. A., editor, *Insights in the Economics of Aging*, pages 87–119. National Bureau of Economic Research, Cambridge, MA. NBER Chapter 13647.

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

Board of Governors of the Federal Reserve System (2017). Assessing the severity of rent burden on low-income families. FEDS Notes.

Board of Governors of the Federal Reserve System (2023). Survey of consumer finances (scf). The 2022 Survey of Consumer Finances is the most recent survey conducted.

Clark, W. A. V., Deurloo, M. C., and Dieleman, F. M. (2003). Housing careers in the united states, 1968-93: Modelling the sequencing of housing states. *Urban Studies*, 40(1):143–160.

Collins, J. M., Hembre, E., and Urban, C. (2018). Exploring the rise of mortgage borrowing among older americans. Working Paper CRR WP 2018-3, Center for Retirement Research at Boston College, Chestnut Hill, MA.

Davidoff, T. (2015). Can "high costs" justify weak demand for the home equity conversion mortgage? *The Review of Financial Studies*, 28(8):2364–2398.

Davidoff, T., Gerhard, P., and Post, T. (2017). Reverse mortgages: What homeowners (don't) know and how it matters. *Journal of Economic Behavior Organization*, 133:151–171.

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

De Nardi, M., French, E., and Jones, J. B. (2016). Savings after retirement: A survey. *Annual Review of Economics*, 8:177–204.

De Nardi, M. and Yang, F. (2014). Bequests and heterogeneity in retirement wealth. *European Economic Review*, 72:182–196.

Deaton and Paxson (1994). *Saving, Growth, and Aging in Taiwan*. University of Chicago Press, Chicago.

Federal Reserve Bank of St. Louis (2025). The state of u.s. household wealth. Open Vault, published June 23, 2025.

Flavin, M. and Yamashita, T. (2011). Owner-occupied housing: Life-cycle implications for the household portfolio. *American Economic Review*, 101(3):609–14.

French, E., Jones, J. B., and McGee, R. (2022). Savings after retirement. Technical Report 22-12, Federal Reserve Bank of Richmond.

Henderson, J. V. and Ioannides, Y. M. (1983). A model of housing tenure choice. *The American Economic Review*, 73(1):98–113.

Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.

Lockwood, L. M. (2018). Incidental bequests and the choice to self-insure late-life risks. *The American Economic Review*, 108(9):2513–2550.

Lusardi, A., Mitchell, O. S., and Oggero, N. (2018). The changing face of debt and financial fragility at older ages. *AEA Papers and Proceedings*, 108:407–11.

Lusardi, A., Mitchell, O. S., and Oggero, N. (2020). Understanding debt in the older population. Working Paper 28236, National Bureau of Economic Research.

Meghir, C. and Pistaferri, L. (2004). Income variance dynamics and heterogeneity. *Econometrica*, 72(1):1–32.

Modigliani, F. and Brumberg, R. H. (1954). Utility analysis and the consumption function: An interpretation of cross-section data. In Kurihara, K. K., editor, *Post-Keynesian Economics*. Rutgers University Press, New Brunswick, NJ.

Morrow-Jones, H. A. and Wenning, M. V. (2005). The housing ladder, the housing life-cycle and the housing life-course: Upward and downward movement among repeat home-buyers in a us metropolitan housing market. *Urban Studies*, 42(10):1739–1754.

Nakajima, M. and Telyukova, I. (2019). Home equity in retirement. Working Paper 19-50, Federal Reserve Bank of Philadelphia.

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

Ortalo-Magné, F. and Rady, S. (2006). Housing market dynamics: On the contribution of income shocks and credit constraints*. *The Review of Economic Studies*, 73(2):459–485.

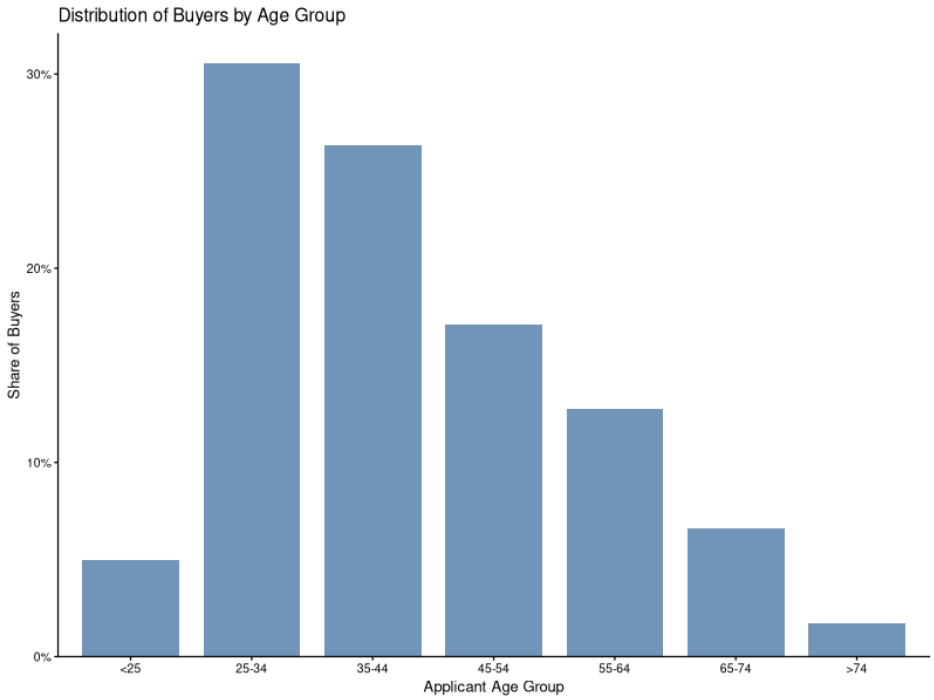
Sinai, T. and Souleles, N. S. (2003). Owner-occupied housing as a hedge against rent risk. Working Paper 9462, National Bureau of Economic Research.

Sinai, T. and Souleles, N. S. (2005). Owner-occupied housing as a hedge against rent risk*. *The Quarterly Journal of Economics*, 120(2):763–789.

Storesletten, K., Telmer, C. I., and Yaron, A. (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, 112(3):695–717.

Venti, S. F. and Wise, D. A. (2001). Aging and housing equity: Another look. Working Paper 8608, National Bureau of Economic Research.

Figure 12: Age Distribution of Conventional Mortgages



A Appendix

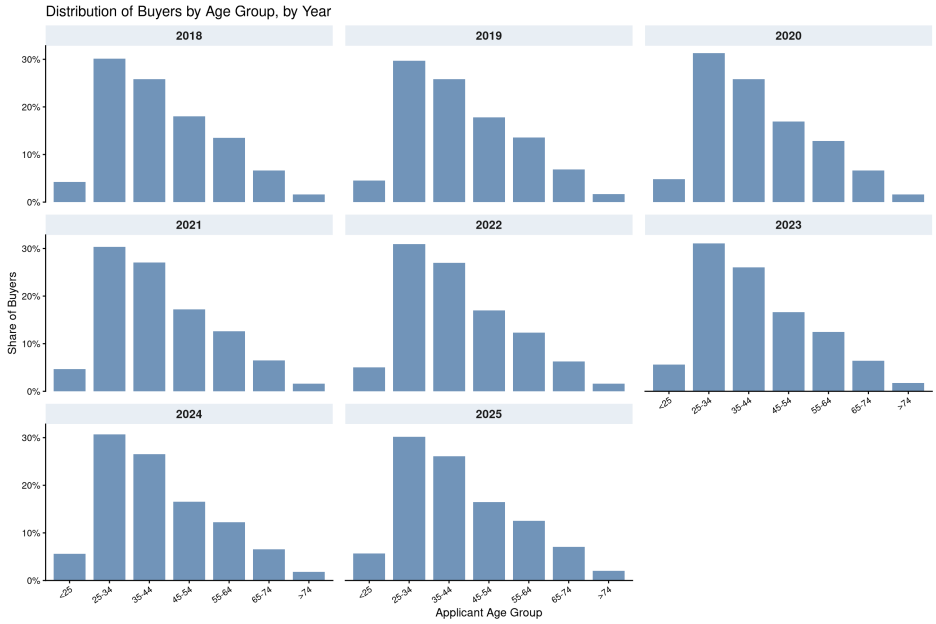
A.1 Normal Distribution Across Age

On Figure 12, I present the distribution of house purchases across age groups for the whole sample 2018-2025. The distribution appears approximately normal across the age bins defined in the HMDA data. This pattern remains consistent when the data disaggregated by year (Figure 13). Restricting the sample to FHA-financed loans yields a similar age distribution in percentage terms, although at a lower level. Furthermore, using the full sample of conventional mortgages, I examine the age distribution under several restrictions: owner-occupied properties only, owner-occupied first-lien properties, all purchases excluding second homes, and all purchases excluding investment properties. In each case, the resulting distributions remain approximately uniform. The corresponding figures are reported below.

To ensure that the observed distribution is not driven by second-home purchases, investment properties, or other transactions that deviate from the life-cycle hypothesis, I present a series of figures that explicitly exclude these categories. See Figures 14-17.

I also present a series of figures (see Figures 18-20) illustrating the penetration rate of homebuyers by age group, their relative representation compared to the national age distribution, and the overall population composition. The age distribution of the U.S. population is drawn from the American Community Survey (ACS) for the years 2018-2024. The normally distributed penetration rate shown in Figure 18 is the arithmetic complement of the narmally distributed transaction share: because younger cohorts are larger in the population and the adult population remains mostly flat (Figure 20), a normal distribution of transactions implies similar distribution across the age groups.

Figure 13: Age Distribution of Conventional Mortgages by Year



A.2 Role of the Borrowing Constraints: Loan-To-Value Ratio and Down Payment

Credit availability shapes the ability to purchase homes by governing when the down payment constraint becomes slack. In the baseline model, this constraint relaxes endogenously as households accumulate savings along a common income process, generating a pronounced spike in purchases at middle age, which is one of the model's central shortcomings. The HMDA data, which contains detailed information on loan applications and outcomes, provides a rich empirical basis for further assessing the role of borrowing constraints.

If heterogeneity in borrowing constraints were driving the rising age effects, younger buyers who successfully purchase despite limited savings would be expected to exhibit systematically higher LTV ratios, debt-to-income ratios and lower down payments than older buyers. Table 7 shows no such significant pattern: older buyers take slightly less loan as the share of the property value, but so they have less income and more debts which they have to serve. Hence, 5-10% of down payments serve as a trade-off for their tightened budget constraint.

Figure 21 displays the age distribution of borrowers across LTV ratio bins. The distribution of ages has spikes for yound and middle ages, but not for the seniors. This rules out credit-constraint proxies (such as high LTV) as an explanation for the large amount of investment into housing assets of older households.

A.3 Geo-Spatial Sorting

Next, I examine whether local market conditions can account for the behavior of the senior households. For example it can be the case that the observed distribution reflects heterogeneity in housing markets across areas with differing price levels. Younger households may

Figure 14: Owner occupied

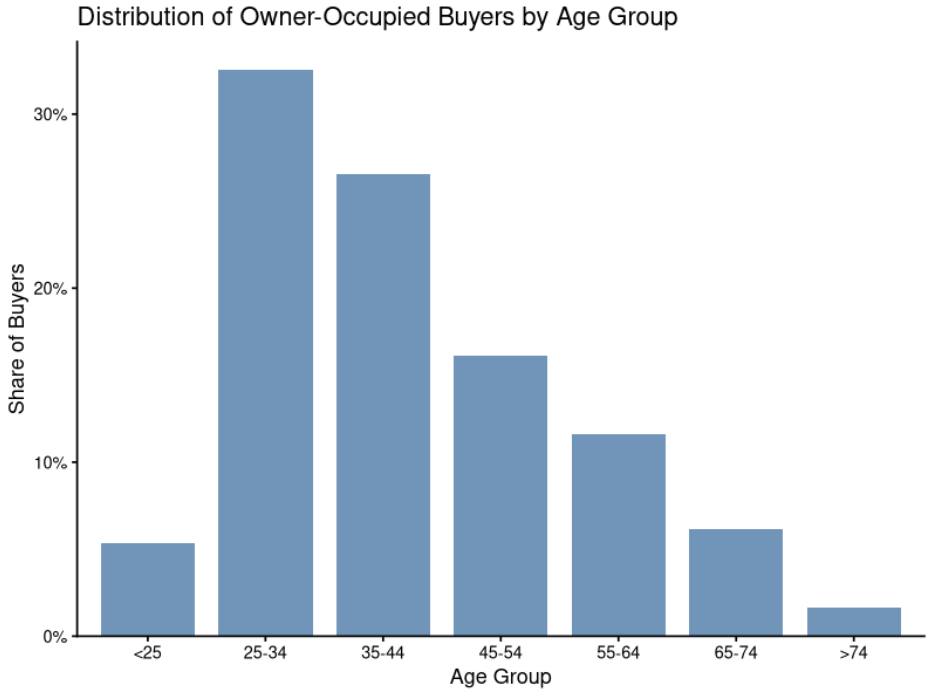


Figure 15: Owner first lien

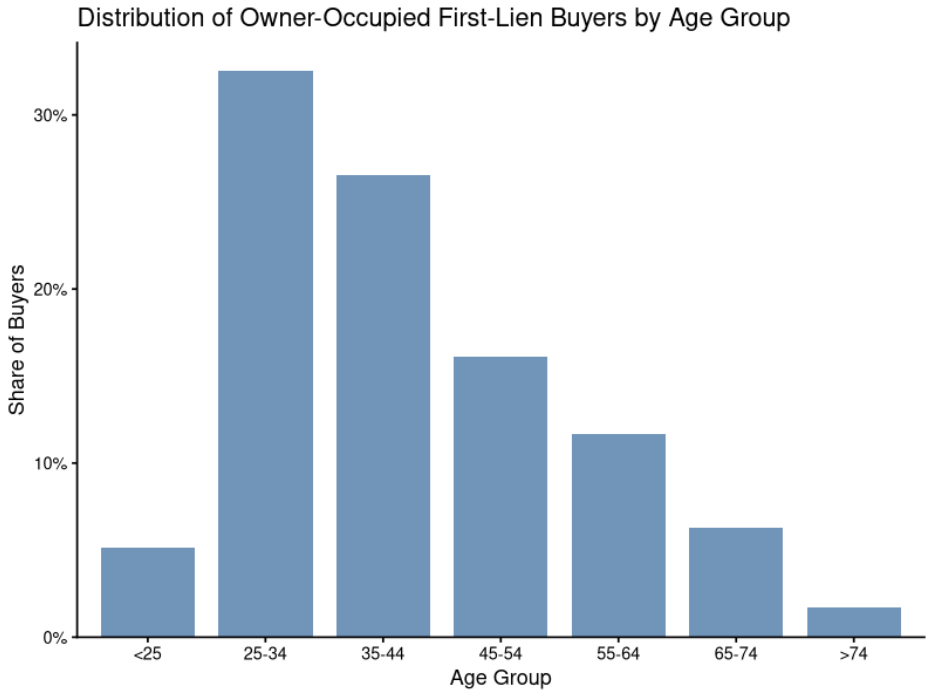


Figure 16: No second home

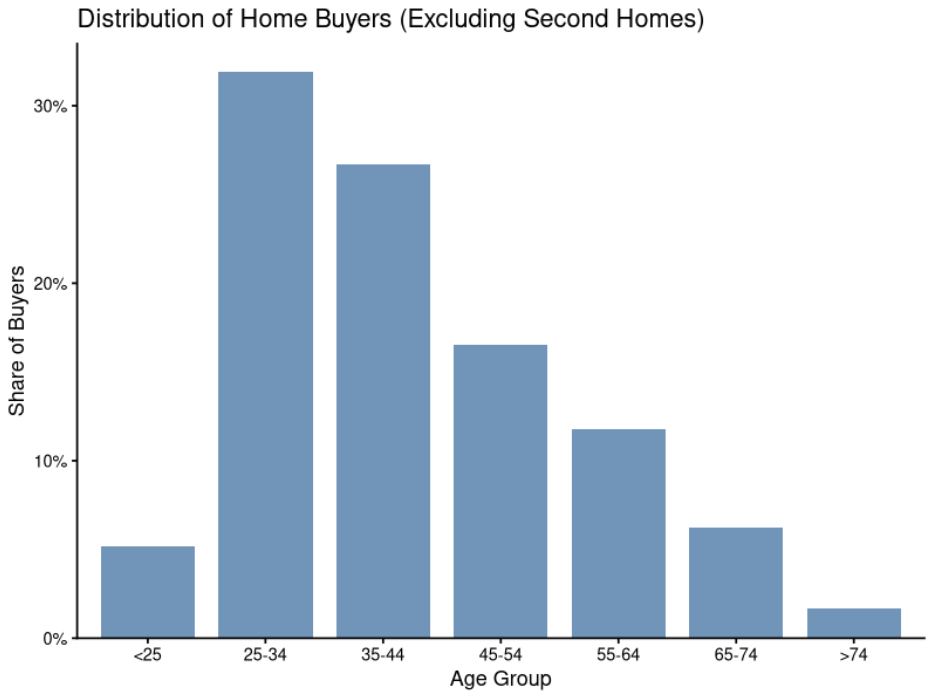


Figure 17: No investment

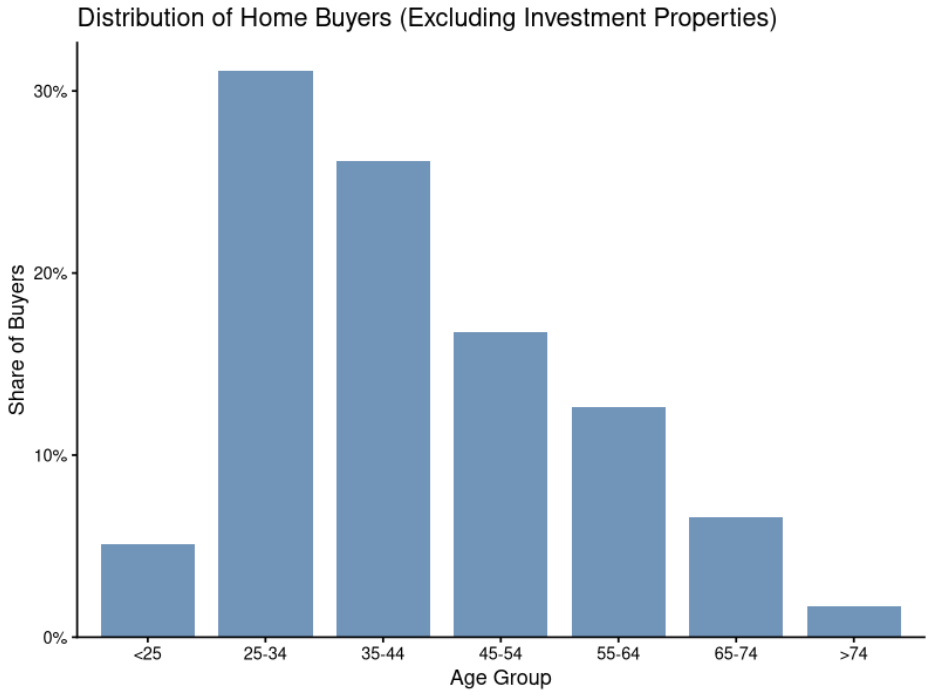


Figure 18: Buyer Penetration Rate

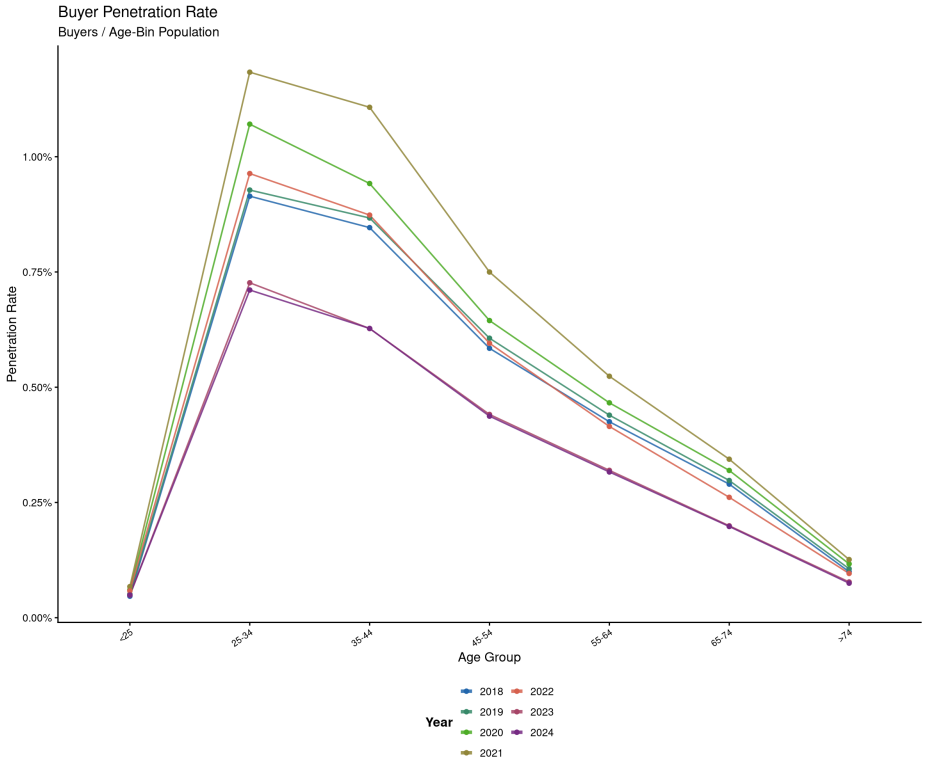


Figure 19: Relative House Buyer Representation

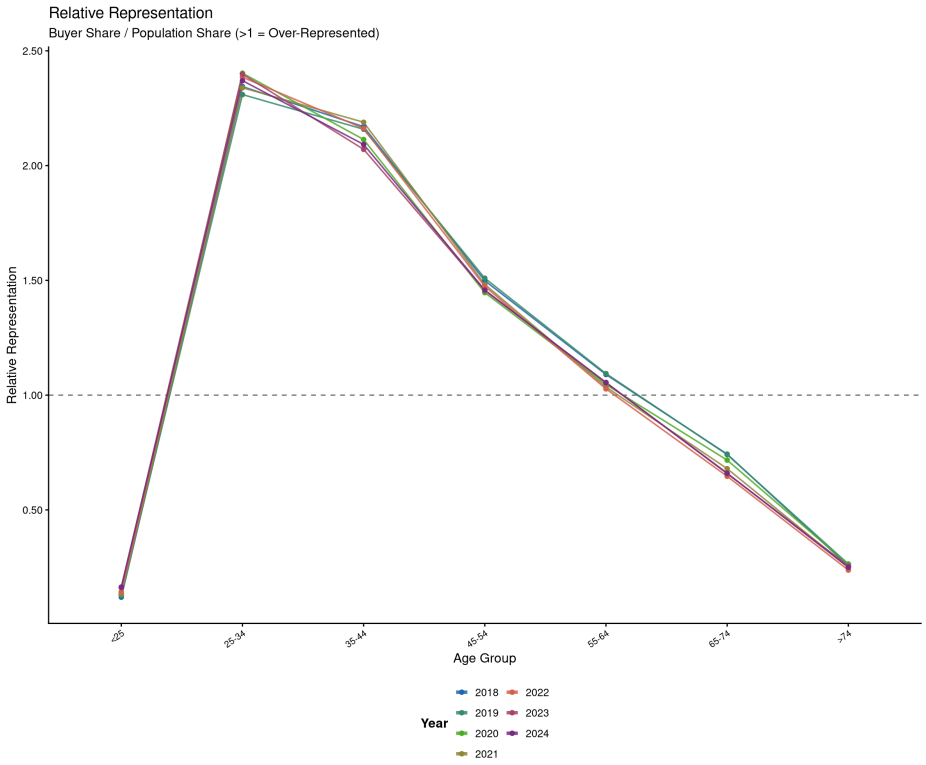


Figure 20: US Population Composition (Census)

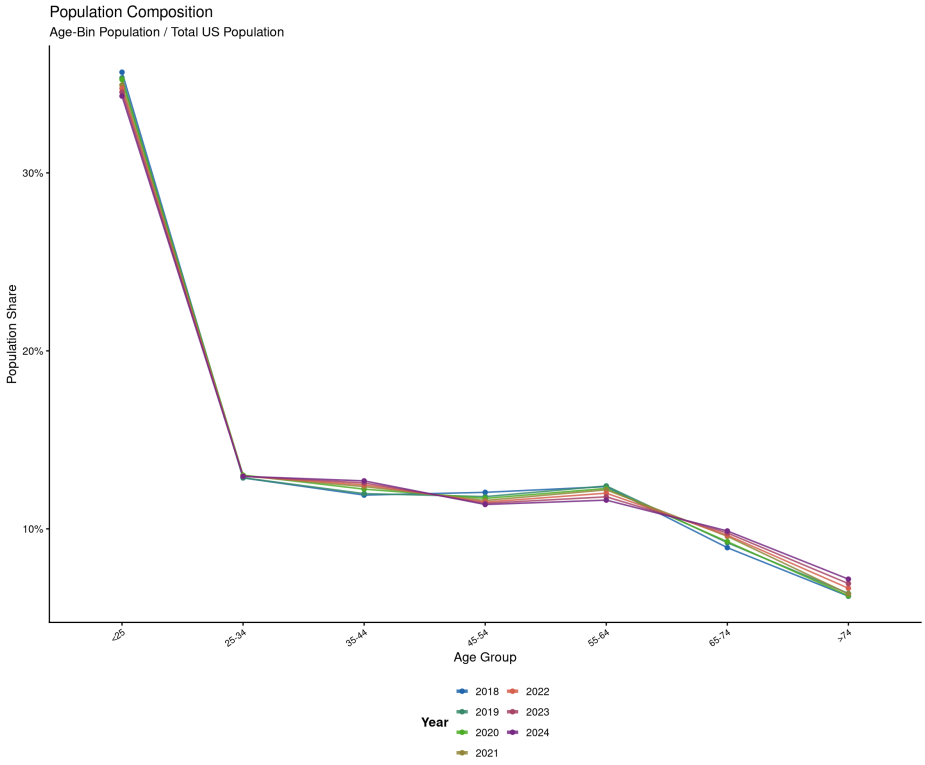
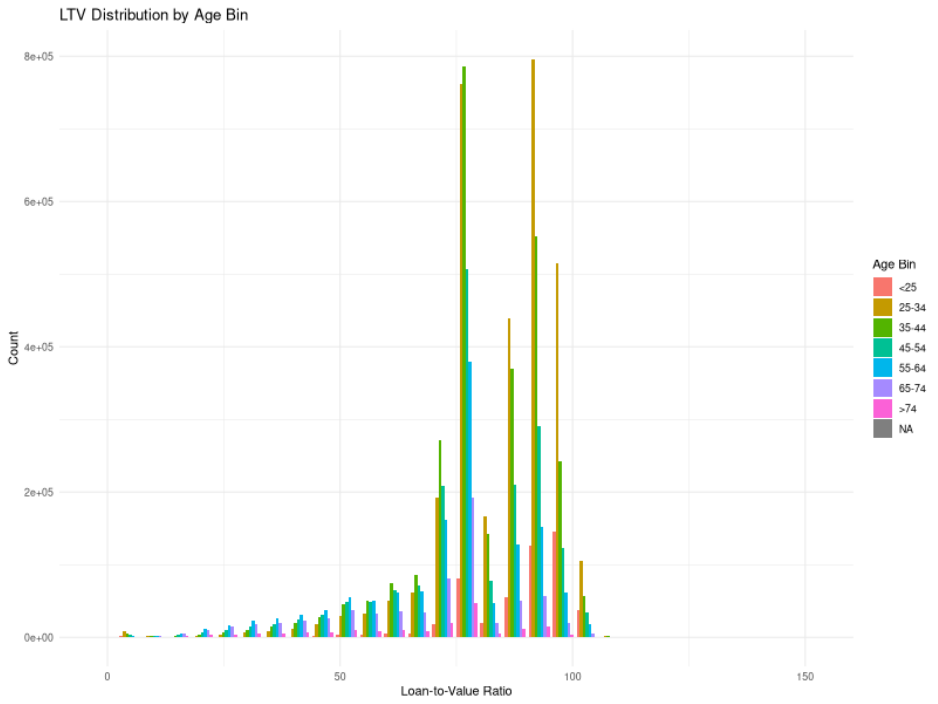


Table 7: Average Financial Constraints by Age Bin

Age Bin	Median LTV	Mean DTI	Mean Income	Mean Property Value	Median Property Value
Under 25	95.00	36.27	80.5	243,970.1	215,000.0
Age 25-34	90.00	35.74	119.3	393,569.2	325,000.0
Age 35-44	80.00	36.00	156.2	494,193.2	405,000.0
Age 45-54	80.00	35.90	163.9	473,135.1	385,000.0
Age 55-64	80.00	35.71	150.0	438,209.6	355,000.0
Age 65-74	75.19	37.15	117.8	412,467.0	335,000.0
Over 74	75.00	38.11	103.5	401,139.8	325,000.0

Figure 21: Loan-to-Value Ratios and Age Distributions



sort into lower-cost areas, while older households may concentrate in higher-cost locations. As a result, purchase prices across age groups may be different due to this spatial allocation combined with variation in house prices. However, Figure 22 shows that the age distribution for seniors (above 55 years old) remains largely unchanged even when counties are partitioned into high-cost and low-cost groups. I classify a county as low-cost if its average house price falls below the median across all U.S. counties.

To test this spatial concern further, I first estimate the intensive margin model (log property value on age, cohort, and purged period effects) under four progressively restrictive geographic fixed-effect specifications: no spatial controls; state fixed effects (within-state comparison); county fixed effects (within-county comparison); and census-tract fixed effects (within-neighborhood comparison). The census-tract specification effectively holds neighborhood amenities, local supply conditions, and even street housing services constant. Contrary to the spatial sorting hypothesis, the estimated age gap between young and retired buyers persists as geographic controls become more granular. The old-young disparity persists even at the census-tract level, precisely where unobserved locational quality differences are most fully absorbed. This pattern implies that aggregate estimates capture the true life-cycle effect precisely enough and that seniors' demand for expensive properties is satisfied even at the census-tract level. Results are reported in Table 8 and the age effects coefficients across specifications are plotted in Figure 23.

A parallel concern for the extensive margin is that, foe example, younger households' lower home purchase rates might reflect spatial sorting: young adults concentrating in markets with fundamentally lower purchase propensities, rather than life-cycle factors. To test this, I construct age-specific home purchase rates at four geographic granularities (national,

Figure 22: Distribution of Age for High Cost and Low Cost Counties

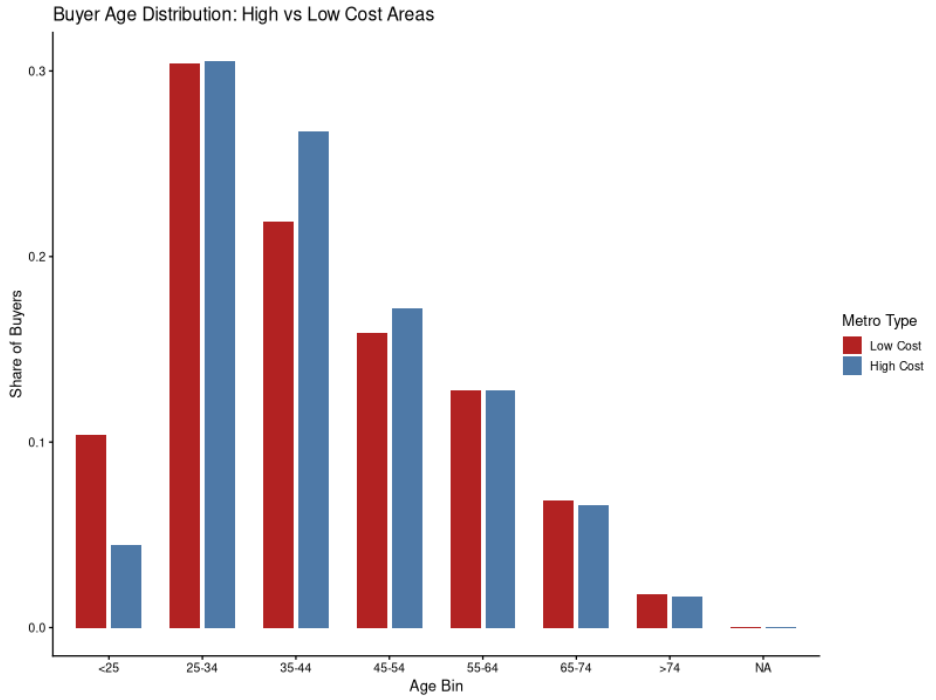


Figure 23: Age Effects for Intensive Margin with Varying Spatial FE

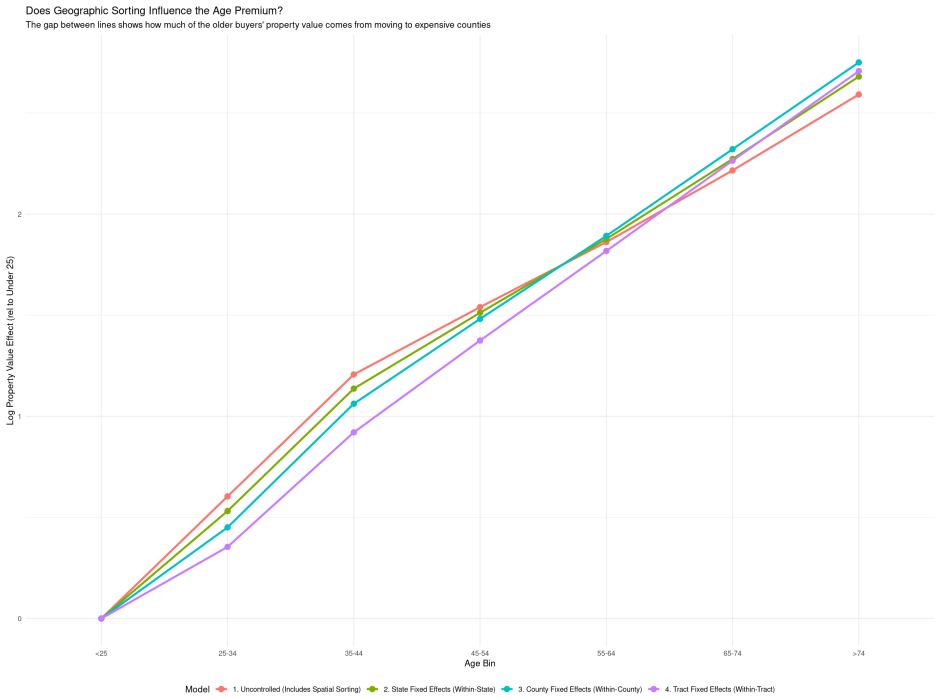
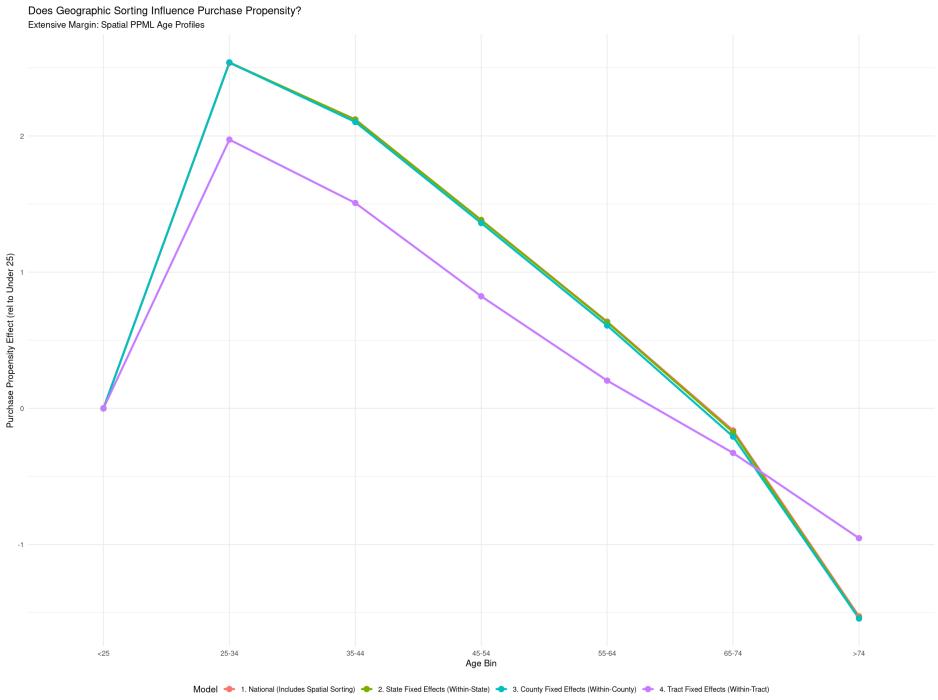


Table 8: Intensive Margin APC Effects Across Geographic Aggregations (Log Property Value)

Dependent Variable:	Log property value			
	National	State FE	County FE	Tract FE
	OLS	OLS	OLS	OLS
Age Effects (Base: < 25)				
Age 25-34	0.6039***	0.5316***	0.4508***	0.3543***
	(0.0012)	(0.0010)	(0.0009)	(0.0007)
Age 35-44	1.207***	1.137***	1.062***	0.9208***
	(0.0015)	(0.0014)	(0.0012)	(0.0010)
Age 45-54	1.540***	1.512***	1.482***	1.375***
	(0.0019)	(0.0017)	(0.0015)	(0.0013)
Age 55-64	1.861***	1.877***	1.892***	1.817***
	(0.0025)	(0.0022)	(0.0019)	(0.0017)
Age 65-74	2.216***	2.272***	2.321***	2.264***
	(0.0030)	(0.0027)	(0.0024)	(0.0022)
Over 74	2.591***	2.680***	2.750***	2.707***
	(0.0040)	(0.0036)	(0.0032)	(0.0028)
Cohort Effects (Base: 1930)				
Cohort 1940	0.4239***	0.4604***	0.4778***	0.4812***
	(0.0035)	(0.0031)	(0.0028)	(0.0022)
Cohort 1950	0.8307***	0.9039***	0.9428***	0.9579***
	(0.0040)	(0.0036)	(0.0032)	(0.0026)
Cohort 1960	1.241***	1.349***	1.409***	1.437***
	(0.0043)	(0.0038)	(0.0034)	(0.0028)
Cohort 1970	1.639***	1.779***	1.862***	1.908***
	(0.0046)	(0.0041)	(0.0036)	(0.0031)
Cohort 1980	2.025***	2.193***	2.299***	2.367***
	(0.0048)	(0.0043)	(0.0038)	(0.0033)
Cohort 1990	2.425***	2.623***	2.751***	2.835***
	(0.0051)	(0.0046)	(0.0040)	(0.0035)
Cohort 2000	2.561***	2.783***	2.935***	3.029***
	(0.0053)	(0.0048)	(0.0042)	(0.0037)
Constant	9.776***			
	(0.0052)			
Geographic FE	No	State	County	Tract
Observations	10,742,318	10,737,783	10,727,207	10,721,881
R^2	0.08682	0.26915	0.43078	0.63723

Note: *p<0.1; **p<0.05; ***p<0.01

Figure 24: Age Effects for Extensive Margin (PPLM Model) at Varying Spatial Level



state, county, census tract) by combining 2022 ACS 5-year population counts by age bin with HMDA purchase counts by age bin, year, and geography. I then estimate PPML regressions of the buyer share (transactions over population in each age-year-geography cell) on age, cohort, and purged period effects,weighted by cell population, under progressively restrictive geographic fixed effects: state, county, and census tract. The tract specification conditions on local market conditions, effectively comparing purchase rates of different aged households within the same neighborhood. The age coefficients mostly follow the same pattern at the different levels of geographical granularity. The persistence of a pronounced age effect at the most granular spatial unit rules out spatial sorting as an explanation for the observed purchase rates. Estimation results are reported in Table 9 and the age effects coefficients across specifications are plotted in Figure 24.

A.4 PSID Panel Age-Cohort-Period Effects

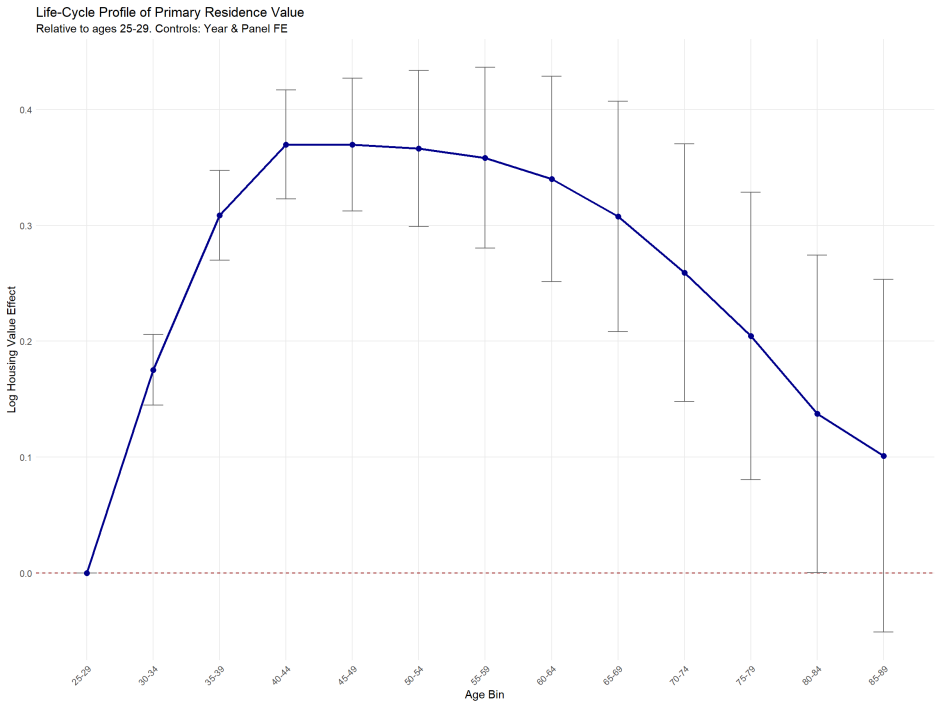
To contrast the mortgage origination results with a broader measure of housing demand, I replicate the age-period-cohort analysis on the Panel Study of Income Dynamics (PSID), 1968-2023. The sample comprises household-year observations with strictly positive, non-top-coded self-reported primary residence values and non-missing non-housing wealth, income, and family size. I estimate a two-way fixed-effects panel model of log self-reported house value on 5-year age group indicators, with household and survey year fixed effects and standard errors clustered at the household level. The estimated age profile reported in Table 10 and plotted in Figure 25 follows the canonical life-cycle pattern: housing wealth rises until retirement and declines thereafter.

Table 9: Extensive Margin APC Effects Across Geographic Aggregations (PPML)

Dependent Variable:	Buyer share of an age bin			
	National	State FE	County FE	Tract FE
	Poisson	Poisson	Poisson	Poisson
Age Effects (Base: < 25)				
Age 25-34	2.536***	2.540***	2.540***	1.973***
	(0.0029)	(0.0024)	(0.0024)	(0.0028)
Age 35-44	2.121***	2.120***	2.102***	1.508***
	(0.0035)	(0.0029)	(0.0029)	(0.0034)
Age 45-54	1.383***	1.381***	1.360***	0.8228***
	(0.0042)	(0.0035)	(0.0035)	(0.0041)
Age 55-64	0.6368***	0.6321***	0.6082***	0.2035***
	(0.0052)	(0.0042)	(0.0042)	(0.0050)
Age 65-74	-0.1624***	-0.1748***	-0.2077***	-0.3279***
	(0.0062)	(0.0051)	(0.0051)	(0.0060)
Over 74	-1.527***	-1.538***	-1.543***	-0.9535***
	(0.0081)	(0.0066)	(0.0066)	(0.0078)
Cohort Effects (Base: 1930)				
Cohort 1940	-0.3111***	-0.3117***	-0.3200***	-0.2419***
	(0.0068)	(0.0055)	(0.0056)	(0.0067)
Cohort 1950	-0.6911***	-0.6924***	-0.7052***	-0.5227***
	(0.0077)	(0.0063)	(0.0063)	(0.0076)
Cohort 1960	-1.120***	-1.122***	-1.138***	-0.8439***
	(0.0083)	(0.0068)	(0.0068)	(0.0082)
Cohort 1970	-1.534***	-1.537***	-1.556***	-1.188***
	(0.0088)	(0.0072)	(0.0072)	(0.0087)
Cohort 1980	-1.860***	-1.865***	-1.886***	-1.479***
	(0.0093)	(0.0076)	(0.0076)	(0.0092)
Cohort 1990	-2.185***	-2.190***	-2.213***	-1.769***
	(0.0098)	(0.0080)	(0.0081)	(0.0097)
Cohort 2000	-2.257***	-2.263***	-2.288***	-1.861***
	(0.0105)	(0.0086)	(0.0086)	(0.0103)
Constant	-5.156***			
	(0.0102)			
Geographic FE	No	State	County	Tract
Observations	56	2,856	151,673	2,362,096

Note: *p<0.1; **p<0.05; ***p<0.01

Figure 25: Age Effects for Intensive Margin in Panel Model



Critically, the PSID measures the stock of housing wealth (self-reported current value of the owned home), not the flow of leveraged mortgage-financed purchases captured by HMDA. The absence of an increasing age effect in housing wealth, alongside its pronounced presence in mortgage borrowing, implies that leverage is central to the observed phenomenon: older buyers not only purchase higher-value homes but do so with substantially high loan-to-value ratios. This divergence between the wealth and borrowing margins underscores that the senior demand surge is a credit-market phenomenon as much as a housing-demand phenomenon.

B Appendix

B.1 Constrained Optimization Problem for a Household

Household i's problem:

$$\max_{\{c_t^i, s_t^i, a_{t+1}^i, h_{t+1}^i\}_{t=0}^T} \mathbb{E} \sum_{t=0}^T \beta^t U(c_t^i, s_t^i),$$

$$\text{s.t. } c_t^i + rent_t s_t^i + p_t^h h_{t+1}^i + a_{t+1}^i = y_t^i + rent_t h_t^i + p_t^h h_t^i - htc_t + (1 + r_t^a) a_t^i,$$

$$a_{t+1}^i \geq -b * p_t^h h_{t+1}^i.$$

$$rent_t = \psi p_t^h$$

$$htc_t = \begin{cases} \phi(p_t^h h_{t+1} + p_t^h h_t) & \text{if } h_{t+1} \neq h_t \\ 0 & \text{otherwise} \end{cases}$$

Table 10: Life-Cycle Age Effects on Primary Residence Value

Dependent Variable:	Log house value
Age Effects	
Age 30-34	0.1752***
	(0.0156)
Age 35-39	0.3086***
	(0.0198)
Age 40-44	0.3698***
	(0.0240)
Age 45-49	0.3695***
	(0.0292)
Age 50-54	0.3662***
	(0.0343)
Age 55-59	0.3583***
	(0.0398)
Age 60-64	0.3401***
	(0.0452)
Age 65-69	0.3077***
	(0.0508)
Age 70-74	0.2591***
	(0.0568)
Age 75-79	0.2044***
	(0.0632)
Age 80-84	0.1373**
	(0.0699)
Age 85-89	0.1012
	(0.0776)
Individual FE	Yes
Year FE	Yes
Observations	56,118
R^2	0.87141
Within R^2	0.03331

Note: *p<0.1; **p<0.05; ***p<0.01

b - collateral constraint, (1-down payment) for the mortgage.

If the solution exists, the constrained optimization problem can be rewritten in the Lagrangian formulation:

$L = \sum_{t=0}^{T} L(c_t^i, s_t^i, a_{t+1}^i, h_{t+1}^i; \mu_t^i, \lambda_t^i) = \mathbb{E}[\sum_{t=0}^{T}\{\beta^t U(c_t^i, s_t^i) + \beta^t \mu_t^i(y_t^i + rent_t h_t^i + p_t^h h_t^i - \phi(p_t^h h_{t+1} + p_t^h h_t) + (1 + r_t^a)a_t^i - c_t^i - rent_t s_t^i - p_t^h h_{t+1}^i - a_{t+1}^i) + \beta^t \lambda_t^i(a_{t+1}^i + b * p_t^h h_{t+1}^i)\}]$

Kuhn-Tucker conditions for houseowner:

(1) $\frac{\partial L}{\partial c_t^i} = \frac{\partial U(c_t^i, s_t^i)}{\partial c_t^i} - \mu_t^i = 0$

(2) $\frac{\partial L}{\partial s_t^i} = \frac{\partial U(c_t^i, s_t^i)}{\partial s_t^i} - \mu_t^i rent_t = 0$

(3) $\frac{\partial L}{\partial a_{t+1}^i} = -\mu_t^i + \beta\mathbb{E}[\mu_{t+1}^i(1 + r_{t+1}^a)] + \lambda_t^i = 0$

(4) $\frac{\partial L}{\partial h_{t+1}^i} = -\mu_t^i(\phi + 1)p_t^h + \lambda_t^i b p_t^h + \beta\mathbb{E}[\mu_{t+1}^i(rent_{t+1} + p_{t+1}^h(1 - \phi))] = 0$

(5) $\mu_t^i(y_t^i + rent_t h_t^i + p_t^h h_t^i - htc_t + (1 + r_t^a)a_t^i - c_t^i - rent_t s_t^i - p_t^h h_{t+1}^i - a_{t+1}^i) = 0$

(6) $\lambda_t^i(a_{t+1}^i + b * p_t^h h_{t+1}^i) = 0$

We consider the simple case of always binding budget constraint, which implies that $\mu_t^i > 0$. From (1) get $\frac{\partial U(c_t^i, s_t^i)}{\partial c_t^i} = \mu_t^i$ and substitute into (2), (3), (4) to get

(7) $\frac{1}{rent_t}\frac{\partial U(c_t^i, s_t^i)}{\partial s_t^i} = \frac{\partial U(c_t^i, s_t^i)}{\partial c_t^i}$,

(8) $-\frac{\partial U(c_t^i, s_t^i)}{\partial c_t^i} + \beta\mathbb{E}[\frac{\partial U(c_{t+1}^i, s_{t+1}^i)}{\partial c_{t+1}^i}(1 + r_{t+1}^a)] + \lambda_t^i = 0$,

(9) $-\frac{\partial U(c_t^i, s_t^i)}{\partial c_t^i}(\phi + 1)p_t^h + \lambda_t^i b p_t^h + \beta\mathbb{E}[\frac{\partial U(c_{t+1}^i, s_{t+1}^i)}{\partial c_{t+1}^i}(rent_{t+1} + p_{t+1}^h(1 - \phi))] = 0$

Next, from (8) express $\lambda_t^i = \frac{\partial U(c_t^i, s_t^i)}{\partial c_t^i} - \beta\mathbb{E}[\frac{\partial U(c_{t+1}^i, s_{t+1}^i)}{\partial c_{t+1}^i}(1 + r_{t+1}^a)]$ and substitute into (9):

$-\frac{\partial U(c_t^i, s_t^i)}{\partial c_t^i}(\phi + 1)p_t^h + (\frac{\partial U(c_t^i, s_t^i)}{\partial c_t^i} - \beta\mathbb{E}[\frac{\partial U(c_{t+1}^i, s_{t+1}^i)}{\partial c_{t+1}^i}(1 + r_{t+1}^a)])bp_t^h + \beta\mathbb{E}[\frac{\partial U(c_{t+1}^i, s_{t+1}^i)}{\partial c_{t+1}^i}(rent_{t+1} + p_{t+1}^h(1 - \phi))] = 0$. We get $-\frac{\partial U(c_t^i, s_t^i)}{\partial c_t^i}(\phi + 1)p_t^h + \frac{\partial U(c_t^i, s_t^i)}{\partial c_t^i}bp_t^h - \beta\mathbb{E}[\frac{\partial U(c_{t+1}^i, s_{t+1}^i)}{\partial c_{t+1}^i}(1 + r_{t+1}^a)]bp_t^h + \beta\mathbb{E}[\frac{\partial U(c_{t+1}^i, s_{t+1}^i)}{\partial c_{t+1}^i}(rent_{t+1} + p_{t+1}^h(1 - \phi))] = 0$, which by rearranging becomes $\beta\mathbb{E}[\frac{\partial U(c_{t+1}^i, s_{t+1}^i)}{\partial c_{t+1}^i}(\frac{rent_{t+1} + p_{t+1}^h(1-\phi)}{p_t^h} - (1 + r_{t+1}^a)b)] = \frac{\partial U(c_t^i, s_t^i)}{\partial c_t^i}(\phi + 1 - b)$.

Recall that now we can also use equation (7), which relates consumption, housing services purchased one period before and shelter services. Substituting (7) into the previous equation, we end up with the Euler equation of the main interest for any period t:

$\beta\mathbb{E}\left[\frac{1}{\psi p_{t+1}^h}\frac{\partial U(c_{t+1}^i, s_{t+1}^i)}{\partial s_{t+1}^i}\left(\frac{rent_{t+1} + p_{t+1}^h(1-\phi)}{p_t^h} - (1 + r_{t+1}^a)b\right)\right] = \frac{\partial U(c_t^i, s_t^i)}{\partial c_t^i}(\phi + 1 - b)$ for $t \in [0; T]$.

Considering the linear transformation of h_t^i into s_t^i, which provides $\frac{\partial s_{t+1}^i}{\partial h_{t+1}^i} = 1$, for homeowners the following Euler equation will also be true

$\beta\mathbb{E}\left[\frac{\partial U(c_{t+1}^i, s_{t+1}^i)}{\partial s_{t+1}^i}\frac{\partial s_{t+1}^i}{\psi p_{t+1}^h \partial h_{t+1}^i}\left(\frac{rent_{t+1} + p_{t+1}^h(1-\phi)}{p_t^h} - (1 + r_{t+1}^a)b\right)\right] = \frac{\partial U(c_t^i, s_t^i)}{\partial c_t^i}(\phi + 1 - b)$ for $t \in [0; T]$.

C Appendix

C.1 Data Subset Selection

I construct the empirical sample from the HMDA mortgage level data for 2018-2025. I retain observations for which the mortgage application was approved and resulted in an originated loan, restrict the sample to home-purchase loans, and require a fixed-rate mortgage, defined by an introductory-rate period equal to zero. I exclude observations with non-positive reported income and remove the top 0.25 percent of the income distribution. The resulting sample contains 9,784,213 mortgage observations before the subsequent merges with geographic and market level data. The selection criteria are designed to produce a relatively homogeneous sample of purchase mortgages for which household income, property value, mortgage terms, and debt information are available. The final constructed dataset contains no missing, zero, or negative reported property values; property values range from 5,000 to 3.235 million dollars.

C.2 Assignment of ZIP-codes

HMDA does not provide the ZIP-code of the mortgaged property. Instead, it reports the census tract identifier. I therefore link HMDA observations to ZIP-codes using the HUD USPS ZIP-Code Crosswalk Files. The crosswalk can assign a census tract to multiple ZIP-codes. To obtain a single ZIP-code for each tract-year observation, I select the ZIP-code with the largest residential share of the tract whenever residential shares are available. If no positive residential share is available for a tract, I instead use the ZIP-code with the largest total tract share. Ties are resolved deterministically by ZIP-code order. I use the 2010-vintage crosswalk for HMDA years 2018-2021, the 2023-vintage crosswalk for 2022-2023, and the 2025-vintage crosswalk for 2024-2025. This procedure results in a one-to-one tract-to-ZIP assignment.

C.3 Assignment of Mortgage Decision Month

HMDA does not directly report the month in which the mortgage decision was made. I infer the decision month using the relationship between the mortgage interest rate, the reported rate spread, and the Average Prime Offer Rate (APOR). For each mortgage observation, I calculate the implied APOR as $\text{APOR}_i = \text{InterestRate}_i - \text{RateSpread}_i$. I then match this implied APOR to the weekly fixed-rate APOR series published by the Federal Financial Institutions Examination Council (FFIEC) and the Consumer Financial Protection Bureau (CFPB). The match is performed within calendar year and uses the nearest available weekly APOR. The date of the matched weekly APOR determines the month assigned to the mortgage observation. This procedure exploits the fact that the HMDA rate spread is defined relative to the applicable APOR and therefore provides information about the timing of the mortgage rate offer. The resulting month is used as the time index for matching the mortgage observation to monthly housing, rental, and financial data.