From Flipping to Holding: Dynamic Strategies of Housing Investors*

Carlos Garriga,† Pedro Gete,‡ Athena Tsouderou,§ and Franco Zecchetto¶

April 2026

Abstract

We develop a dynamic portfolio model in which housing investors' holding periods are endogenously determined by resale, rental, and financing decisions. Credit conditions, capital gains, rental yields, and returns on alternative assets shape investors' decisions to flip properties or hold them as long-term rentals. This, in turn, determines market turnover. Leverage-driven strategies are fragile to price reversals, while rent-oriented strategies are more resilient. Using transaction-level U.S. deed records, we document a sharp shift in investor behavior across two housing booms, from mortgage-financed flipping toward longer-horizon, cash-based, rent-oriented strategies, consistent with the model's predictions and underlying economic mechanisms.

Keywords: Housing investment, holding periods, buy-and-hold, flipping, real estate investors, leverage, portfolio choice, house prices, rental yields.

---

*We thank Jose Maria Barrero, Elena Carletti, Alejandro Casado, João Cocco, Morris Davis, Anthony DeFusco, Alessandro Gavazza, Andra Ghent, Marco Giacoletti, Adam Guren, Jonathan Halket, Jean Helwege, Finn Kydland, Matthijs Korevaar, Lauren Lanahan, Jack Liebersohn, Matias Lamas, Peter Maxted, Alexander Michaelides, Dmitriy Muravyev, Charles Nathanson, Rafael Repullo, Heather Royer, Daniel Ruf, Lawrence Schmidt, Martin Schneider, David Sraer, Eva Steiner, Michael Weber, Michael Wittry, Dacheng Xiu, and several seminar and conference participants. The views expressed herein do not necessarily reflect those of the Federal Reserve Bank of St. Louis or the Federal Reserve System. Research reported in this paper was partially funded by MCIN /AEI /10.13039/501100011033 / FEDER, UE Grant No. PID2021-125359NB-I00.

†Federal Reserve Bank of St. Louis. Email: carlos.garriga@stls.frb.org.

‡IE University. Email: pedro.gete@ie.edu.

§University of Miami. Email: atsouderou@miami.edu.

¶ITAM Business School. Email: franco.zecchetto@itam.mx.

1 Introduction

Housing investors play a central role in U.S. single-family housing markets, yet their investment horizons (how long they hold properties before selling) remain poorly understood. Holding periods are a first-order determinant of housing market turnover and resale activity (e.g., DeFusco, Nathanson, and Zwick, 2022; Bayer et al., 2020), with broader implications for investor composition and housing market dynamics. Yet most existing work treats them as fixed or secondary. This paper argues that holding periods are an endogenous outcome of investors' forward-looking portfolio decisions and a key channel through which macro-financial conditions reshape housing investment strategies.

We develop a dynamic portfolio model in which housing investors optimally choose whether to resell properties quickly or operate them as long-term rental assets. Investors decide how much housing to hold and when to sell. They also choose whether to rent or keep properties vacant to preserve resale options, whether to sell vacant units, and how to finance housing purchases using mortgage credit or by reallocating wealth away from financial assets. Holding periods emerge endogenously from these choices and respond to expected capital gains, rental yields, credit conditions, and returns on alternative investments.

The model centers on two margins. The first is a selling margin, which governs when vacant units are sold immediately rather than retained. The second is a forward-looking vacancy-rental margin, which determines whether unsold units are kept vacant to preserve the option to sell next period or are rented to earn cash-flow returns. Together, these margins determine turnover and the distribution of holding periods across housing units. High expected capital gains lower resale cutoffs and shorten holding periods, while high rental yields tilt portfolios toward renting and longer horizons by raising the opportunity cost of selling. Loose credit lowers the effective cost of replacing sold properties, amplifying short-term resale incentives. Declines in safe returns induce portfolio reallocation toward housing. This strengthens rent-oriented strategies even in the presence of rising house prices.

This framework delivers sharp and testable predictions for holding periods. When mortgage credit is loose and rental yields are low, holding periods shorten and short-term resale becomes highly sensitive to expected capital gains. When returns on alternative assets are low and housing is operated primarily for rental income, holding periods lengthen, resale decisions become less responsive to capital gains and more sensitive to rental yields. More generally, the sensitivity of short-term resale to capital gains and rental returns depends on the macro-financial environment. Consequently, flipping and buy-and-hold strategies do not simply reflect fixed investor types, but also emerge endogenously from these forces.

Our analysis complements the framework of Glaeser and Nathanson (2017), which generates house price momentum and reversals through extrapolation of past price growth, while treating holding periods as exogenous. By contrast, we endogenize resale behavior, allowing it to respond to macro-financial conditions. This shift moves the focus away from price expectations alone toward dynamic selling incentives and highlights how capital gains, credit conditions, rental returns, and outside asset yields shape investment horizons.

We bring the model's predictions to the data using transaction-level deed records from CoreLogic covering the universe of U.S. single-family home sales from 2000 to 2017. The data reveal a striking regime change across the two major U.S. housing booms in our sample. Despite similarly large and persistent increases in real house prices, investor behavior differed sharply across periods. During the pre-Global Financial Crisis (GFC) boom, investors relied heavily on mortgage credit (e.g., Albanesi, DeGiorgi, and Nosal, 2022), resold properties quickly, and exhibited strong short-term resale responses to capital gains. By contrast, in the post-GFC boom, investors held properties longer, relied predominantly on cash financing, and shifted toward rent-oriented strategies, in line with broader search-for-yield forces (e.g., Daniel, Garlappi, and Xiao, 2021). Consistent with the model's predictions, short-term resale became far less sensitive to capital gains and increasingly sensitive to rental yields.

Empirically, our analysis focuses on retail housing investors, who account for the majority of investor activity in our data. While institutional investors have grown in importance in recent years (e.g., Gurun et al., 2023), they represent a small share of housing transactions in our data and may differ along important dimensions such as constraints and objectives. Accordingly, the framework is best interpreted as capturing retail investor behavior.

Calibrated versions of the model reproduce the observed shift in short-term resale sensitivities across periods. In a pre-GFC housing boom characterized by loose credit and a rapidly rising price-to-rent ratio, the model generates leverage-driven expansion, high turnover, and a steep relationship between expected appreciation and short-term resale. In a post-GFC housing boom with tighter credit, higher rental yields, and low returns on safe assets, the model instead produces long-term, rent-oriented investment behavior and muted resale responses to capital gains, even when house price growth is held fixed.

The model also highlights important differences in the resilience of investment strategies. Counterfactual boom-bust experiments show that leverage-driven, short-term strategies are inherently fragile: when house prices fall and credit conditions tighten, resale activity collapses. By contrast, long-term, rent-oriented strategies are substantially more resilient, as rental cash flows continue to support housing positions and limit sales even under sharp price reversals. These differences imply that shifts in investor strategies have first-order consequences not only for housing booms, but also for the dynamics of busts.

Taken together, the model and evidence point to the emergence of a new investment regime following the GFC, in which housing increasingly serves as an asset delivering cash-flow returns rather than short-term speculative gains. This shift has important implications for the transmission of monetary policy: when investor purchases are largely cash-financed, policy rates affect housing demand not only through portfolio reallocation, but also by altering incentives to hold properties for longer periods and reducing market turnover. More broadly, by treating holding periods as a key endogenous margin, the paper helps bridge models of housing booms driven by price expectations with the observed shift toward long-term, rent-oriented investment behavior in the post-GFC era.

The remainder of the paper proceeds as follows. Section 2 discusses the related literature. Section 3 presents the model, and Section 4 derives its key mechanisms and predictions. Section 5 documents empirical evidence consistent with the model's mechanisms using transaction-level deed records. Section 6 presents the model calibration and compares model-implied and empirical short-term resale behavior across macro-financial environments. Section 7 presents housing boom and bust experiments, and Section 8 concludes. Appendices A and B provide additional details on the model and data.

2 Related Literature

A large literature studies the role of investors in U.S. housing markets during the boom of the early 2000s, emphasizing short-term speculation, rapid resale activity, and heavy reliance on mortgage credit. Much of the expansion and subsequent collapse in housing market activity during this period has been linked to the rise and fall of speculative investors. DeFusco, Nathanson, and Zwick (2022) document that fluctuations in housing transaction volume were largely driven by speculative activity. Bayer et al. (2020) show that short-term investors responded strongly to past house price appreciation, while Bhutta (2015) documents a sharp expansion of mortgage debt among real estate investors.

Related work highlights the role of lax credit conditions in sustaining these investment strategies (Albanesi, DeGiorgi, and Nosal, 2022), the growth and geographic concentration of short-term speculation and house flipping during the pre-GFC period (Haughwout et al., 2011; Garcia, 2022), and belief-driven trading based on expectations of continued price growth (Piazzesi and Schneider, 2009). While this literature provides important insights into investor behavior during the pre-GFC period, it leaves open how investor strategies evolve when credit conditions, rental incentives, and returns on alternative assets change.

A related strand of the literature interprets pre-GFC investor behavior in terms of behavioral biases. Chinco and Mayer (2016) show that out-of-town second-home buyers behaved like misinformed speculators, while Gao, Sockin, and Xiong (2020) link speculative activity to extrapolative expectations based on past price growth. Glaeser and Nathanson (2017) emphasize belief distortions and underreaction to fundamentals, and Loewenstein and Willen (2023) attribute part of the 2000s house price boom to exuberant expectations. More broadly, Andersen et al. (2022) study how reference dependence and loss aversion shape housing sale decisions. While this work provides important insights into pre-GFC investor behavior, it leaves open the extent to which post-GFC changes in investment strategies reflect endogenous responses to macro-financial conditions, beyond shifts in beliefs and biases.

Recent work emphasizes the role of macro-financial forces in shaping post-GFC housing investment behavior. A growing literature documents search-for-yield and portfolio reallocation in low interest rate environments (Daniel, Garlappi, and Xiao, 2021), with investors shifting toward higher-yielding assets when safe returns decline. Using historical data from 18th-century Amsterdam, Korevaar (2023) shows that investors reached for yield in response to low safe returns. Consistent with this view, housing increasingly functions as a cash-flow asset when safe yields fall (Demers and Eisfeldt, 2022; Eichholtz et al., 2021).

Mabille (2023) and Gupta, Hansman, and Mabille (2025) highlight the importance of credit conditions and financial constraints for housing outcomes and housing allocation. Building on this perspective, we show that these forces also determine investors' holding periods and the equilibrium balance between flipping and buy-and-hold strategies.

Existing evidence also shows that investor activity has contributed to higher house prices and rents in the post-GFC period, both through direct demand effects (Allen et al., 2018) and through the expansion of investor ownership (Mills, Molloy, and Zarutskie, 2019; Lambie-Hanson, Li, and Slonkosky, 2022). In this context, Garriga, Gete, and Tsouderou (2023) show that monetary policy affects housing demand through portfolio reallocation rather than solely through borrowing costs, while Boddin et al. (2023) provide evidence of a housing portfolio channel of QE transmission. Ghent (2012) studies how adjustment frictions shape monetary transmission through housing markets. Related work also documents reaching-for-income behavior among specific household groups (e.g., Gargano and Giacoletti, 2024). Our focus, instead, is on how macro-financial conditions shape investment horizons, resale behavior, and financing choices for housing investors across housing cycles.[1]

Relative to this literature, our contribution is to provide a structural framework that links house price growth, rental incentives, credit conditions, and outside returns to the emergence of short-term resale and long-term rent-oriented strategies. Rather than taking holding periods as given, the model makes them an endogenous outcome of forward-looking portfolio decisions, and we assess its predictions using transaction-level U.S. deed records. This approach allows us to interpret post-GFC shifts in investor behavior as responses to changing macro-financial conditions, not solely as shifts in beliefs. More broadly, our framework relates to dynamic models of housing finance and mortgage choice (Campbell and Cocco, 2015), while focusing on holding periods and resale behavior as key outcomes.

3 A Model of Housing Investors

This section presents a dynamic model in which housing investors choose whether to flip properties, that is, resell them quickly, or to rent and hold them. A key feature of the model is that holding periods are fully endogenous to the macro-financial environment: changes in expected house price growth, rental returns, credit availability, and returns on alternative assets shift the balance between short-term resale and long-term rental activity.

Investors purchase housing, decide whether to finance via mortgage credit or by reallocating wealth away from bonds and other liquid assets, choose whether to rent units or

[1] A related literature studies the rise of large institutional landlords in single-family rental markets and their effects on rents and local welfare (e.g., Gurun et al., 2023).

keep them vacant, and decide which vacant units to sell. Heterogeneity in rental and resale benefits generates interior rental and resale shares, creating meaningful, forward-looking trade-offs between preserving the option to sell and earning rental income.

3.1 Investor Preferences

Consider an investor who makes dynamic housing investment decisions over time. Time is discrete and denoted by t. The investor has preferences over consumption c_t, represented by the per-period utility function $u(c_t)$. Preferences are time-separable, and future utility is discounted at rate β per period. The investor's expected lifetime utility at time t is

$$\mathbb{E}_t \left[\sum_{k=0}^{\infty} \beta^k u(c_{t+k}) \right], \tag{1}$$

where the expectation is taken over the stochastic processes for house prices, rental prices, credit conditions, and bond returns. These processes are described below.

3.2 Housing Investment

The investor enters period t with a stock of housing units H_{t-1}, and a share ϕ_{t-1} of those units rented out. The investor collects rental income $\phi_{t-1}q_t$ per unit on these rented houses, where q_t is the rental price at time t. The investor then chooses the share $\tilde{\rho}_t$ of vacant units to sell, generating income $(1-\phi_{t-1})\tilde{\rho}_t p_t$ per unit sold, where p_t is the house price at time t. The unsold portion is retained and becomes part of end-of-period holdings H_t. The investor also incurs maintenance costs that offset the physical depreciation of the housing stock, $\delta p_t H_{t-1}$.

Let $\rho_t = (1-\phi_{t-1})\tilde{\rho}_t$ be the share of previous-period holdings H_{t-1} sold at time t. After selling, the investor chooses housing purchases H_t^* at cost $p_t H_t^*$, subject to convex adjustment costs $\psi(H_t^*, H_{t-1})$ that capture transaction and renovation frictions. We assume

$\psi_{H^*} > 0$, $\psi_H < 0$, and $\psi_{H^*H^*} > 0$. There is a minimum purchase size $H_t^* \geq h_{\min}$.

Total housing holdings then evolve according to:

$$H_t = \underbrace{H_t^*}_{\text{new houses}} + \underbrace{(1-\rho_t)H_{t-1}}_{\text{old houses}}. \tag{2}$$

Finally, the investor chooses the share ϕ_t of total holdings H_t to allocate to rental use.

3.3 Heterogeneity Across Housing Units

Heterogeneity is introduced at the level of individual housing units within the investor's portfolio. Each unit receives idiosyncratic shocks that affect its rental and resale values.

Specifically, every unit i rented at time $t-1$ yields a payoff $(1+\omega_{i,t-1})q_t$ at the beginning of period t, where the idiosyncratic rental shock $\omega_{i,t-1}$ is drawn i.i.d. from a distribution Γ_ω at time $t-1$. The investor therefore knows each unit's rental suitability when deciding whether to rent it out or leave it vacant. Likewise, each unit i kept vacant at time $t-1$ yields a payoff $(1+\kappa_{i,t})p_t$ if sold at time t, where the idiosyncratic resale shock $\kappa_{i,t}$ is drawn i.i.d. from a distribution Γ_κ at time t. Thus, the investor knows the resale valuation of each vacant unit when making the selling decision. The two shocks are mutually independent. The investor applies cutoff rules: units with $\omega_{i,t} \geq \bar{\omega}_t$ are rented at time t (and the rest are kept vacant for potential sale next period), and vacant units with $\kappa_{i,t} \geq \bar{\kappa}_t$ are sold at time t (and the rest are retained). The cutoffs $\bar{\omega}_t$ and $\bar{\kappa}_t$ are determined endogenously.[2]

The share of units rented, ϕ_t, is determined by the rental cutoff $\bar{\omega}_t$ through $\phi_t = 1 - \Gamma_\omega(\bar{\omega}_t)$, with the associated surplus $S_\omega(\phi_t) = \int_{\Gamma_\omega^{-1}(1-\phi_t)}^{\infty} \omega \, d\Gamma_\omega$. Analogously, given the selling cutoff $\bar{\kappa}_t$, the share of vacant units sold is $\tilde{\rho}_t = 1 - \Gamma_\kappa(\bar{\kappa}_t)$, with corresponding sales

[2]Heterogeneity in rental and selling benefits may reflect property- or location-specific factors such as differences in management or maintenance costs, geographic proximity affecting monitoring, neighborhood amenities, local regulations, transaction and marketing frictions, and variation in local market tightness.

surplus $S_\kappa(\tilde{\rho}_t) = \int_{\Gamma_\kappa^{-1}(1-\tilde{\rho}_t)}^{\infty} \kappa \, d\Gamma_\kappa$.[3] These cutoff-based shares and surplus terms translate directly into revenue expressions: rental income per unit of housing H_{t-1} at time t is

$$\underbrace{\left(\int_{\bar{\omega}_{t-1}}^{\infty} (1+\omega)\, d\Gamma_\omega\right)}_{\text{rented houses}} q_t = \left(\phi_{t-1} + S_\omega(\phi_{t-1})\right) q_t, \tag{3}$$

and sales income per unit of housing H_{t-1} at time t is

$$\underbrace{\left(1 - \int_{\bar{\omega}_{t-1}}^{\infty} d\Gamma_\omega\right)}_{\text{vacant houses}} \underbrace{\left(\int_{\bar{\kappa}_t}^{\infty} (1+\kappa)\, d\Gamma_\kappa\right)}_{\text{houses sold}} p_t = (1-\phi_{t-1})\left(\tilde{\rho}_t + S_\kappa(\tilde{\rho}_t)\right) p_t. \tag{4}$$

Finally, the unit-level shocks shape the distribution of holding periods across the investor's portfolio: units transition between rental, vacancy, and sale according to their idiosyncratic valuations and prevailing macro-financial conditions.

3.4 Mortgage Financing

The investor can finance housing purchases with long-term, fixed-rate mortgage debt. Mortgages amortize geometrically: the outstanding balance decays at rate $(1-\nu)$, where ν denotes the share of principal repaid each period. For a mortgage of size M_t^* originated at time t, the sequence of payments (amortization plus interest) is $(\nu + r_t^*)M_t^*$, $(\nu + r_t^*)(1-\nu)M_t^*$, $(\nu + r_t^*)(1-\nu)^2 M_t^*$, and so on, where r_t^* is the interest rate at origination.

Mortgages are subject to a loan-to-value (LTV) limit θ_t at origination:

$$M_t^* \leq \theta_t p_t H_t^*. \tag{5}$$

When units are sold, the associated mortgages are fully prepaid at the time of sale.

[3] We assume that the distributions Γ_ω and Γ_κ are strictly increasing with support on the entire real line. This implies a one-to-one mapping between the cutoffs $(\bar{\omega}_t, \bar{\kappa}_t)$ and the corresponding interior rental and sales shares $(\phi_t, \tilde{\rho}_t)$, so the decision problem can be expressed equivalently in terms of cutoffs or shares.

The investor enters period t with a total outstanding mortgage balance M_{t-1} and interest payments X_{t-1}. Both objects must be tracked separately because mortgages carry fixed interest rates. After choosing the share of vacant houses to sell $\tilde{\rho}_t$ and housing purchases H_t^*, the outstanding mortgage balance evolves according to:

$$M_t = \underbrace{M_t^*}_{\text{new mortgages}} + \underbrace{(1-\rho_t)(1-\nu)M_{t-1}}_{\text{old mortgages}}. \tag{6}$$

Mortgage interest payments evolve according to:

$$X_t = \underbrace{M_t^* r_t^*}_{\text{new interest}} + \underbrace{(1-\rho_t)(1-\nu)X_{t-1}}_{\text{old interest}}. \tag{7}$$

3.5 Bond Investment

At the beginning of period t, the investor receives the face value of the maturing bond holdings, B_{t-1}. The investor then chooses a new position B_t, purchased at price $\frac{1}{1+r_t}$ per unit, where r_t is the real risk-free interest rate between periods t and $t+1$. Bond holdings are chosen jointly with housing purchases H_t^* as part of the investor's portfolio allocation.

Bond positions must be non-negative, $B_t \geq 0$, and this constraint may occasionally bind.

3.6 Stochastic Process for Prices, Credit Conditions, and Interest Rates

The investor faces exogenous, stationary Markov processes governing housing and rental prices, credit conditions, and bond returns. House prices p_t follow a log AR(1) process:

$$\log(p_t) = \rho_p \log(p_{t-1}) + (1-\rho_p)\bar{p} + \epsilon_{p,t}, \tag{8}$$

where $\epsilon_{p,t}$ are i.i.d. normally distributed shocks with mean zero and variance σ_p^2.

Rental prices q_t follow a similar log AR(1) process:

$$\log(q_t) = \rho_q \log(q_{t-1}) + (1 - \rho_q)\bar{q} + \epsilon_{q,t}, \tag{9}$$

where $\epsilon_{q,t}$ are i.i.d. normally distributed shocks with mean zero and variance σ_q^2.

Credit conditions are captured by a single index ξ_t, which jointly determines the LTV limit θ_t and the mortgage interest rate r_t^* for new originations. These variables move perfectly together, reflecting shifts between loose and tight credit regimes.

Finally, bond interest rates r_t follow an AR(1) process:

$$r_t = \rho_r r_{t-1} + (1 - \rho_r)\bar{r} + \epsilon_{r,t}, \tag{10}$$

where $\epsilon_{r,t}$ are i.i.d. normally distributed shocks with mean zero and variance σ_r^2.

All four shock processes are assumed to be mutually independent.

3.7 Investor's Dynamic Problem

We now formalize the investor's decision problem in recursive form. At the start of period t, the investor enters with the following endogenous state variables: previous housing holdings H_{t-1}, the share of those holdings rented out ϕ_{t-1}, outstanding mortgage balance M_{t-1}, mortgage interest payments X_{t-1}, and previous bond holdings B_{t-1}. These are summarized by z_{t-1}. The exogenous state variables are the current house price p_t, rental price q_t, credit conditions ξ_t, and bond interest rate r_t, which are collectively denoted by x_t.

The investor chooses consumption c_t, new housing purchases H_t^*, new mortgage borrowing M_t^*, the rental share of the new end-of-period stock ϕ_t, the share of previously vacant units to sell $\tilde{\rho}_t$, and new bond holdings B_t. The investor's recursive problem is given by:

$$V(z_{t-1}; x_t) = \max_{c_t, H_t^*, M_t^*, \phi_t, \tilde{\rho}_t, B_t} \{u(c_t) + \beta \mathbb{E}_t [V(z_t; x_{t+1})]\} \quad \text{s.t.} \tag{11}$$

$$\underbrace{c_t}_{\text{consumption}} + \underbrace{\delta p_t H_{t-1}}_{\text{maintenance}} + \underbrace{p_t H_t^*}_{\text{new houses}} + \underbrace{\psi(H_t^*, H_{t-1})}_{\text{adjustment cost}} + \underbrace{\rho_t (1-\nu) M_{t-1}}_{\text{prepayment from selling}} + \underbrace{\nu M_{t-1}}_{\text{principal}} + \underbrace{X_{t-1}}_{\text{interest}} + \underbrace{\frac{B_t}{1+r_t}}_{\text{new bonds}}$$

$$= \underbrace{y_t}_{\text{exog income}} + \left[\underbrace{\phi_{t-1} \left(1 + \frac{S_\omega(\phi_{t-1})}{\phi_{t-1}}\right) q_t}_{\text{rental income}} + \underbrace{(1-\phi_{t-1}) \tilde{\rho}_t \left(1 + \frac{S_\kappa(\tilde{\rho}_t)}{\tilde{\rho}_t}\right) p_t}_{\text{sales income}} \right] \underbrace{H_{t-1}}_{\text{old houses}} + \underbrace{M_t^*}_{\text{new mortgages}} + \underbrace{B_{t-1}}_{\text{old bonds}},$$

subject to the laws of motion for housing, mortgages, and interest payments, given by (2), (6), and (7), the LTV constraint at origination (5), and the constraints $H_t^* \geq h_{\min}$ and $B_t \geq 0$.

3.8 Short-Term Sales and Holding Periods

We track the holding period of housing units. Let s denote the number of periods a unit has been held since its purchase. Define $\Phi_t(s \leq T)$ as the share of the investor's portfolio at the end of period t consisting of units whose holding periods are less than or equal to T.

The short-term share evolves according to

$$\Phi_t(s \leq T) = \underbrace{\frac{H_t^*}{H_t}}_{\text{new houses}} + \underbrace{\frac{(1-\rho_t) H_{t-1}}{H_t} \Phi_{t-1}(s \leq T-1)}_{\text{old houses advancing one period}}, \tag{12}$$

where H_t^* are new purchases, H_t is the end-of-period stock, and ρ_t is the share of units sold at the start of period t. Because holding periods are measured at the end of period t, new purchases enter with a holding period of one, while units that remain in the portfolio simply age by one period. The short-term share does not enter the investor's dynamic problem; it is constructed only in simulations to compute holding-period statistics.

The investor starts period t with a short-term share $\Phi_{t-1}(s \leq T)$. The share of the previous period's stock that is sold at time t with a holding period not exceeding T is:

$$\text{Short-Term Sales}_t = \rho_t \, \Phi_{t-1}(s \leq T). \tag{13}$$

The parameter T determines what is classified as a short holding period.

3.9 Computational Approach

In practice, solving the full problem with both mortgages and bonds would imply an intractably large state space, as it combines the endogenous states $\{H_{t-1}, \phi_{t-1}, M_{t-1}, X_{t-1}, B_{t-1}\}$ with the exogenous states $\{p_t, q_t, \xi_t, r_t\}$. To make the analysis tractable and to isolate distinct financing channels within the model, we study two environments: (i) a *mortgage environment* without bonds ($B_t = 0$), which isolates mortgage financing; and (ii) a *bond environment* without mortgages ($M_t = X_t = 0$), which isolates portfolio rebalancing between housing and bonds. We solve each environment using projection-based numerical methods, with the details of the global nonlinear algorithm provided in Appendix A.

The key model predictions in Section 4 are not specific to this decomposition. These two environments are most informative in regimes where one margin is more relevant than the other. When credit conditions are loose and leverage constraints are not binding, mortgage financing is the primary driver of housing investment decisions, making the mortgage environment particularly informative. Conversely, when credit conditions are tight and investors rely on portfolio reallocation, variation in alternative asset returns governs housing investment decisions, making the bond environment more relevant.

4 Key Predictions of the Model

This section describes the key mechanisms through which the model determines investors' investment horizons and holding periods. Two margins are central: the selling margin, which governs when vacant units are sold, and the vacancy-rental margin, which determines whether units are kept vacant or rented, and thus whether the option to sell in the next period is preserved. These margins link investors' decisions to expected capital gains, rental returns, and credit and portfolio conditions. See Appendix A for derivations.

4.1 Selling Choice for Vacant Units

Vacant units $(1-\phi_{t-1})H_{t-1}$ are sold at the start of period t if their idiosyncratic sale shock $\kappa_{i,t}$ exceeds a cutoff $\bar{\kappa}_t$ (and retained otherwise). The model implies the following cutoff:

$$\bar{\kappa}_t = \frac{1}{p_t}\psi_{H^*}(H_t^*, H_{t-1}) - \mu_t^{\mathrm{LTV}}\theta_t - \frac{\mu_t^{H^*}}{p_t} + \frac{(1-\nu)M_{t-1}}{p_t H_{t-1}}\left(1 - \Omega_t^M - \Omega_t^X \bar{r}_{t-1}\right), \tag{14}$$

where μ_t^{LTV} and $\mu_t^{H^*}$ are the multipliers on the LTV limit and minimum purchase constraints, Ω_t^M and Ω_t^X denote the marginal continuation values of mortgage principal and interest obligations, while $\bar{r}_{t-1}$ is the average interest rate on outstanding mortgage debt. A lower cutoff $\bar{\kappa}_t$ increases the share of vacant units sold, since $\tilde{\rho}_t = 1 - \Gamma_\kappa(\bar{\kappa}_t)$.

Higher house prices p_t lower the cutoff and raise immediate resale activity among vacant units. A low marginal cost of purchasing new units or a loose LTV limit reduces the cutoff, making resale more attractive. A low mortgage prepayment burden $(1-\nu)M_{t-1}$ relative to portfolio value $p_t H_{t-1}$ increases the equity recovered upon sale, further lowering the cutoff. Finally, when the net continuation value of the existing mortgage, $1 - \Omega_t^M - \Omega_t^X \bar{r}_{t-1}$, is low (i.e., servicing the mortgage is costly), the incentive to sell rises.[4]

[4] A larger multiplier $\mu_t^{H^*}$ indicates that the minimum purchase constraint is binding, limiting the investor's ability to reduce housing exposure through purchases and making the sale of vacant units more attractive.

4.2 Vacancy-Rental Choice

End-of-period holdings H_t are rented at time t if their idiosyncratic rental shock $\omega_{i,t}$ exceeds a cutoff $\bar{\omega}_t$ (and kept vacant otherwise). The model implies the following cutoff:

$$\bar{\omega}_t = \frac{\mathbb{E}_t\left[\Lambda_{t+1}\left\{\left(S_\kappa(\tilde{\rho}_{t+1}) - \tilde{\rho}_{t+1}\bar{\kappa}_{t+1}\right)\frac{p_{t+1}}{p_t}\right\}\right]}{\mathbb{E}_t\left[\Lambda_{t+1}\left\{\frac{q_{t+1}}{p_t}\right\}\right]} - 1, \tag{15}$$

where Λ_{t+1} is the stochastic discount factor and $S_\kappa(\tilde{\rho}_{t+1})$ is the surplus from selling vacant units next period. A lower cutoff $\bar{\omega}_t$ increases the rental share, since $\phi_t = 1 - \Gamma_\omega(\bar{\omega}_t)$.

High expected rental yields tilt investors toward renting, while high expected capital gains strengthen the incentive to keep units vacant. The numerator captures the option value of vacancy: the expected capital gains foregone by renting today instead of preserving the possibility of selling at the start of period $t+1$. The term $S_\kappa(\tilde{\rho}_{t+1}) - \tilde{\rho}_{t+1}\bar{\kappa}_{t+1} > 0$ reflects the surplus from selling only high-κ units next period. The denominator captures the expected rental return $\frac{q_{t+1}}{p_t}$. Intuitively, when expected house price growth is high, vacancy becomes more valuable; when expected rental yields are high, the flow return dominates and investors rent more units.

The vacancy-rental margin is also connected to the selling margin. The rental cutoff $\bar{\omega}_t$ depends on the next-period selling cutoff $\bar{\kappa}_{t+1}$, which enters the future sales surplus. Intuitively, higher expected house price growth implies a lower selling cutoff next period, which raises the expected sales surplus and increases the option value of keeping units vacant today.[5] Thus, expected capital gains affect short-term resale behavior both through the current selling margin and the forward-looking vacancy-rental margin.

[5]Formally, the derivative of the surplus term $S_\kappa(\tilde{\rho}_{t+1}) - \tilde{\rho}_{t+1}\bar{\kappa}_{t+1}$ with respect to $\bar{\kappa}_{t+1}$ is $-\tilde{\rho}_{t+1} < 0$; hence a lower selling cutoff next period raises the surplus and the option value of keeping units vacant today.

4.3 Credit Conditions and Housing Purchases

Credit conditions play a central role in investors' willingness to acquire additional housing units. House prices, the LTV limit, and adjustment costs jointly determine the effective marginal cost of expanding the housing portfolio. This cost is given by

$$\mathcal{C}_t = \left(1 - \mu_t^{\mathrm{LTV}} \theta_t\right) p_t + \psi_{H^*}(H_t^*, H_{t-1}). \tag{16}$$

When credit is loose (a high LTV limit) and marginal adjustment costs are small, the marginal cost of acquiring a unit falls. As a result, the selling cutoff $\bar{\kappa}_t$ in (14) also declines, making investors more willing to sell and subsequently repurchase units. Intuitively, easier credit increases short-term resale activity and shortens holding periods.

4.4 Rebalancing Between Housing and Bonds

Investors allocate wealth between new housing purchases H_t^* and one-period risk-free bonds B_t. These assets are linked through the Euler conditions that govern their expected discounted returns. Let P_{t+1}^H denote the gross marginal payoff from holding a housing unit from t to $t+1$. This payoff incorporates the rental income, capital gains if the unit is sold, and the continuation value if it is retained (Appendix A provides the details).

The Euler equations for housing and bonds combine into the following condition:

$$\mathbb{E}_t \left[\Lambda_{t+1} \left\{ \frac{P_{t+1}^H}{\mathcal{C}_t} - (1 + r_t) \right\} \right] = \mu_t^B (1 + r_t) - \frac{\mu_t^{H^*}}{\mathcal{C}_t}, \tag{17}$$

where μ_t^B is the multiplier on the bond non-negativity constraint.

When the constraints are not binding (i.e., $\mu_t^{H^*} = \mu_t^B = 0$), a fall in the risk-free rate r_t requires investors to lower the marginal return to housing. An immediate adjustment occurs through higher demand for housing: as investors substitute away from bonds toward housing, purchases H_t^* increase, raising the marginal acquisition cost $\mathcal{C}_t$ and thereby reducing the housing return.[6] Hence, lower interest rates shift portfolios toward housing.

4.5 Summary of Model Implications

The model yields a unified set of predictions. First, the selling and vacancy-rental cutoffs imply that short-term sales are more likely when expected house price growth is high and less likely when expected rental returns are high. Second, new housing purchases expand when credit conditions are loose and when real yields on alternative assets fall, shifting portfolios toward housing. Together, these mechanisms determine holding periods, resale activity, and the allocation of the housing portfolio across selling, vacancy, and renting margins.

5 Macro Evidence Supporting the Model's Mechanisms

This section documents empirical patterns that align with the model's key mechanisms in Section 4. Although the pre- and post-GFC U.S. housing booms featured similar house price growth, the macro-financial environment and investor behavior differed sharply. We show that (i) rental yields and price-rent ratios evolved in ways consistent with shifts in the vacancy-rental margin, (ii) mortgage credit availability moved in opposite directions across booms, affecting the cost of expanding portfolios and the selling margin, and (iii) the capitalization and financing choices of housing investors changed substantially. Together, these patterns align with the model's prediction that holding periods adjust endogenously to rental returns, capital gains, credit conditions, and outside opportunities.

Our investor-level analysis uses CoreLogic deed records covering all single-family trans-

[6] There is an additional forward-looking effect: higher housing purchases today increase next period's housing stock, raising next period's marginal adjustment cost with respect to H_t and lowering the payoff P_{t+1}^H.

actions from 2000-2017, merged with assessor data on property characteristics. We focus on retail investors, who account for the majority of investor transactions in our data.[7] Deeds report buyer identities, mailing addresses, transaction prices, and lien information, allowing us to identify retail investors and distinguish between cash and credit-financed purchases. We define retail investors as individuals who purchase more than one property within the same MSA over a two-year window, as well as small legal entities, excluding leisure or vacation buyers and large institutional landlords. Buyer mailing addresses are linked to census-tract characteristics to proxy for investors' wealth. See Appendix B for details.

5.1 Macro Drivers and Model-Implied Incentives Across Booms

We define the pre-GFC boom as 2001-2007 and the post-GFC boom as 2012-2018. Panel (a) of Figure 1 shows that both periods featured similarly large and persistent increases in real house prices of roughly 40-45% over six years. However, despite comparable price appreciation, the macro-financial environment evolved very differently across the two booms.

Panel (b) shows that the price-rent ratio (the inverse of the rental yield) diverged sharply across the two booms. Before the GFC, it rose from roughly 15 to above 20 as house prices outpaced rents, implying declining rental yields. After the GFC, rents grew more in line with prices and the ratio stabilized around 16.6, generating substantially higher rental yields. In the model, these differences shift incentives toward operating properties rather than reselling them, a mechanism consistent with the evidence we present below.

Panel (c) documents the evolution of mortgage credit availability.[8] The pre-GFC boom was marked by rapid credit expansion, rising LTV ratios, and looser underwriting, allowing

[7] Before the GFC, most investor purchases in our data are made by individuals using personal names on deeds. After the GFC, a set of large institutional buyers ("Wall Street landlords") emerges (e.g., Gurun et al., 2023), but they account for less than 2% of annual transactions and are concentrated in a few selected locations. Their exclusion does not materially change the empirical patterns documented below.

[8] Measured using the Urban Institute's Housing Credit Availability Index (HCAI), which proxies the overall underwriting stance of mortgage lenders (higher values indicate looser credit standards).

investors to finance purchases with minimal equity and facilitating high turnover. In contrast, the post-GFC boom featured persistently tight credit and stricter documentation. Limited access to mortgage leverage increased reliance on cash-based financing (through liquid savings or portfolio reallocation toward real estate), a pattern consistent with the model's mechanism in which tighter credit raises the effective cost of expanding housing portfolios.

Panel (d) shows that real five-year Treasury yields fell sharply during the first half of the 2010s, reaching about -1.5% at their trough before stabilizing around zero. In contrast, pre-GFC real yields were consistently positive, ranging from roughly 1.0% to 2.5%.[9] This large decline in real rates increased the relative attractiveness of rental housing. In the model, lower outside returns induce portfolio rebalancing toward real estate.

Summarizing, the macro-financial environment differed sharply across booms. The pre-GFC boom combined flat rents and easy credit: conditions that, in the model, tilt incentives toward short-term, leverage-driven strategies. The post-GFC boom instead featured rising rents, tight credit, and low real rates: conditions that, in the model, favor longer-term rental operations. These contrasting environments map directly into the model's mechanisms and frame the investor-level evidence that follows in support of them.

5.2 Investor Capitalization and the Leverage Margin

High leverage and widespread use of mortgage credit were central features of the pre-GFC boom (Sufi and Taylor, 2022). Consistent with this pattern, 65% of investor purchases in the pre-GFC period were financed with a mortgage, versus only 23% in the post-GFC period (Table 1). Conditional LTV ratios also decline modestly, and the mortgage credit availability index corroborates substantially tighter underwriting standards in the post-GFC period.

Post-GFC investors also tended to reside in substantially wealthier neighborhoods, con-

[9] Real ten-year yields fall from a pre-GFC range of 1.5-3.0% to roughly -0.75-0.5% in the post-GFC period.

sistent with greater liquidity and a reduced need for leverage. We proxy investor wealth using the ratio of the median house value in the investor's census tract to the median value of the corresponding MSA; this measure increases from 0.92 in 2001-2004 to 1.50 in 2012-2015. This shift aligns with the rise in cash purchases: better-capitalized investors were more able to finance acquisitions directly rather than through high-LTV mortgages.[10] This pattern is also consistent with recent evidence on the growing importance of cash purchases in housing transactions (Han and Hong, 2024).

Taken together, these facts indicate that the post-GFC investment boom was driven predominantly by well-capitalized cash buyers. In the context of our model, this reflects a tightening of the leverage margin and greater reliance on internal liquidity, consistent with a rebalancing of portfolios toward real estate in a low-interest-rate environment.

5.3 Holding Periods and the Decline in Short-Term Resale Activity

Short-term resale strategies were a defining feature of the pre-GFC boom. Using deed-recorded purchase and sale dates, we measure holding periods for all investor purchases in the 2001-2004 and 2012-2015 cohorts.[11] Figure 2 shows that the share of investor purchases resold within six months to three years declines from about 25% in the pre-GFC boom to 20% in the post-GFC boom. In MSAs that experienced the largest boom-bust cycles in the 2000s (e.g., DeFusco, Nathanson, and Zwick, 2022), the decline is even sharper: for example, from roughly 34% to 22% in Las Vegas, NV. These patterns indicate a broad reduction in short-term resale activity, especially in markets previously prone to speculative trading.

To shed light on the underlying incentives, Figure 3 compares the distribution of county-level rental yields and price growth across the areas where investors buy. While price growth is similarly distributed across the two booms, rental yields shift markedly upward: the av-

[10]Appendix B shows that post-GFC investors were largely new entrants to the single-family market. The appendix also documents that these investors are drawn from higher-income and more educated areas.

[11]Transactions within six months are excluded to remove duplicate deed records; results are unaffected.

erage rises from 7.6% in 2001-2004 to 9.4% in 2012-2015, with post-GFC investors disproportionately purchasing in counties offering yields of 9-12%. This shift toward stronger cash-flow fundamentals is consistent with longer holding periods and reduced reliance on short-term capital gains. These patterns are consistent with the model's mechanism in which higher rental returns raise the value of operating properties relative to resale.

6 Calibration and Model Validation Using U.S. Deeds Data

In this section, we calibrate the model and validate its key mechanisms by comparing model-implied short-term resale behavior to its empirical counterpart. We consider the two environments introduced in Section 3 to isolate the macro-financial channels emphasized in Section 4. The mortgage environment captures how credit conditions, together with movements in prices and rents, shape incentives toward capital-gains-oriented, short-term resale behavior. The bond environment captures how changes in outside asset returns, along with prices and rents, shape incentives for rent-oriented, buy-and-hold behavior.

Once calibrated, we assess whether these mechanisms are supported by the data. The model predicts that short-term resale reflects a trade-off between capital gains and rental returns, and that the relative importance of these margins varies across macro-financial environments. Using transaction-level regressions from CoreLogic deed records, we show that the data are consistent with these predictions, lending support to the model.

6.1 Baseline Environment and Common Parameters

Each model period corresponds to one year. Preferences follow a CRRA utility function

$$u(c) = \frac{c^{1-\sigma} - 1}{1 - \sigma}, \tag{18}$$

with coefficient of relative risk aversion $\sigma = 2$, implying an intertemporal elasticity of substitution of 0.5. Exogenous income is normalized to $y = 1$ and assumed constant over time. The minimum purchase size is set to $h_{\min} = 0.5$, which disciplines the scale of acquisitions.

The housing depreciation rate is set to $\delta = 1.5\%$, consistent with Jeske, Krueger, and Mitman (2013). In the U.S., a typical mortgage is a 30-year fixed-rate loan with constant payments. At a 6% annual mortgage rate, the halfway point of principal repayment occurs at 20.99 years. In the model, the principal balance instead declines geometrically at rate $(1 - \nu)$, with $\nu = \frac{1}{30}$, which implies a half-life of $\frac{\ln(0.5)}{\ln(1-\nu)} = 20.45$ years.

Log house prices p_t and log rents q_t follow AR(1) processes in (8) and (9), with $(\rho_p, \sigma_p) = (0.85, 0.035)$ and $(\rho_q, \sigma_q) = (0.90, 0.02)$, based on AR(1) estimates of real house prices (Case-Shiller Index) and real rents (Rent of Primary Residence Index), both deflated by CPI All Items Less Shelter. The persistence of rents is set above the empirical estimate to match the peak price-rent ratio observed during the post-GFC boom in the model experiments. We normalize the unconditional mean of log house prices to $\bar{p} = 0$ and set log rents to $\bar{q} =$ -2.76, implying an initial price-rent ratio of 15 in the model experiments.

Credit conditions ξ_t follow a two-state Markov chain, with state n denoting normal conditions and state l denoting loose conditions. We set $\Pi_{nn} = \Pi_{ll} = 0.9$, implying that regime shifts occur, on average, once every ten years (Justiniano, Primiceri, and Tambalotti, 2019). Under normal conditions, the LTV limit and mortgage rate are $(\theta_n, r_n^*) = (80\%, 6\%)$, while in the loose regime they are $(\theta_l, r_l^*) = (100\%, 4\%)$. Bond returns r_t follow the AR(1) process in (10), with persistence $\rho_r = 0.80$, innovation standard deviation $\sigma_r = 0.01$, and unconditional mean $\bar{r} = 1.17\%$, based on AR(1) estimates of real returns on U.S. Treasury securities.

All AR(1) processes are discretized as 7-point Markov chains using the Rouwenhorst (1995) method. Further details are provided in Appendix A.

6.2 Calibration of the Mortgage and Bond Environments

For the quantitative analysis, we specify the housing adjustment cost as

$$\psi(H_t^*, H_{t-1}) = \frac{\zeta}{2}\left(\frac{H_t^*}{H_{t-1}}\right)^2 p_t H_{t-1}, \tag{19}$$

where ζ governs the strength of adjustment frictions in housing purchases. We assume logistic distributions for rental and resale heterogeneity, following Greenwald and Guren (2025), with location and scale parameters calibrated separately across environments.

The mortgage environment is calibrated so that the model's ergodic distribution is dominated by long-term investors, with most properties rented and about 15% of houses sold. We jointly choose the discount factor β, the adjustment cost parameter ζ, the rental-income location and dispersion parameters (μ_ω, s_ω), and the sales-income parameters (μ_κ, s_κ) to match the following moments: (i) the value of annual housing purchases equal to 6.5 times income, (ii) a housing value-to-mortgage ratio of 1.45, (iii) a 70% share of rented houses, (iv) a 50% share of vacant houses sold, (v) a rental yield (including the surplus) of 7.5%, and (vi) a capital gain of 4.0% (also including the surplus). Table 2 reports the parameter values.

The bond environment is calibrated so that its ergodic distribution features substantial short-term resale activity, with about 44% of properties sold. The endogenous parameters are chosen to match the following moments: (i) the value of bond holdings equal to 5.5 times income, (ii) a 50% housing share in investors' portfolios, (iii) a 45% share of rented houses, (iv) a 75% share of vacant houses sold, (v) a rental yield of 3.0%, and (vi) a capital gain of 4.0%. Table 3 reports the parameter values. Appendix A provides additional details.

6.3 Model-Based Sensitivities of Short-Term Sales

This subsection computes model-implied sensitivities of short-term resale to capital gains and rental yields, which we then confront with the empirical evidence.

We simulate the model under the optimal policy rules, allowing all exogenous shocks to evolve freely according to their respective stochastic processes. Environment-specific samples are constructed by conditioning on a pre- and post-GFC indicator. In the mortgage environment, we focus on periods with loose credit conditions (LTV limit $\theta_t = 100\%$ and mortgage rate $r_t^* = 4\%$), which capture macro-financial conditions characteristic of the pre-GFC period. In the bond environment, we focus on periods with low bond yields ($r_t < 5.25\%$), which capture macro-financial conditions characteristic of the post-GFC period.

For each simulated observation, we construct the short-term sale margin, defined as the share of properties resold within three model years, using equation (13). Using the simulated data, we estimate OLS regressions including separate lags of capital gains or rental yields, interactions with the post-GFC indicator, and controls for endogenous states. Separate regressions are run for capital gains and rental yields. To summarize the effects, we focus on the impact of a uniform one-percentage-point increase in capital gains or rental yields, aggregating the responses associated with current and lagged returns.[12]

6.4 Empirical Sensitivities of Short-Term Sales

We now assess whether the sensitivities of short-term resale implied by the calibrated model are borne out in the data. Using CoreLogic deed records, we examine how the probability that an investor makes a short-term sale responds to potential capital gains and rental yields. We estimate logit specifications similar to Gilbukh and Goldsmith-Pinkham (2024). Table 4 reports summary statistics for the key variables. Details on data construction, matching

[12] As implied by equations (12) and (13), short-term resale outcomes depend on both current and past returns.

procedures, and variable definitions are provided in Appendix B.

To measure investors' short-term sales responsiveness to capital gains, we estimate:

$$\log\left(\frac{\pi_{ijt}}{1-\pi_{ijt}}\right) = \alpha + \beta\,\text{PriceGrowth}_{jt} \times \text{Post}_t + \gamma\,\text{PriceGrowth}_{jt} + C_j + C_i + C_{ct} + u_{ijt}, \quad (20)$$

where the dependent variable is the log-odds of a short-term sale, with π_{ijt} denoting the probability that property j, purchased by investor i at time t, is sold within two years; we also consider a three-year horizon. The coefficients of interest are γ, which captures the sensitivity of short-term sales to local house price growth in the pre-GFC period, and β, which measures how this sensitivity changes in the post-GFC period.

The variable PriceGrowth_{jt} measures house price appreciation in the property's zip code during the year after purchase.[13] Using zip-code-level price indices allows us to measure potential capital gains for both sold and unsold properties, while avoiding reliance on realized transaction prices that may reflect idiosyncratic characteristics or renovation costs. We interpret this measure as the relevant potential capital gain for the investor's sale decision, with taxes, transaction fees, and management costs absorbed by MSA fixed effects. The indicator Post_t equals 0 in the pre-GFC boom (2001-2004) and 1 in the post-GFC boom (2012-2015).

The specification includes controls for property characteristics C_j, investor characteristics C_i, and county-year economic conditions C_{ct} to mitigate omitted variable bias. Property controls include purchase price, log size (square feet), number of rooms, and house age or effective age when available.[14] Investor controls include indicators for local investors (address in the same MSA as the property), foreign investors (address outside the property's MSA), and legal entities (purchases made through LLCs, LPs, trusts, or corporations). Local economic conditions capture county-year population growth, income growth, and the change in the unemployment rate, all measured in the year prior to sale.

[13] As robustness checks, we also compute house price appreciation using either two-year zip-code price growth or annual growth beginning one year after purchase; the main results are unaffected.

[14] Using only actual age or including both actual and effective age does not affect the results.

The specification also includes MSA and year fixed effects, or MSA×year fixed effects, to further control for demand factors, as well as month-of-purchase fixed effects to capture seasonality. Standard errors are robust and clustered at the zip-code level.[15]

Table 5 reports logit coefficient estimates for equation (20). Capital gains are positively and significantly associated with the probability of a short-term sale, while the interaction with Post_t is negative and significant, indicating a weaker sensitivity in the post-GFC period. Results are highly consistent across the two- and three-year resale horizons.

To measure investors' short-term sales responsiveness to rental yields, we estimate:

$$\log\left(\frac{\pi_{ijt}}{1-\pi_{ijt}}\right) = \alpha + \beta\,\text{RentalYield}_{jt} \times \text{Post}_t + \gamma\,\text{RentalYield}_{jt} + C_j + C_i + C_{ct} + u_{ijt}, \quad (21)$$

where the coefficients of interest are γ, which captures the pre-GFC sensitivity of short-term sales to rental yields, and β, which measures the post-GFC change in that sensitivity.

The variable RentalYield_{jt} is measured as the average rent-to-price ratio in the property's MSA during the year after purchase.[16] Controls C_j, C_i, and C_{ct} are defined as in the capital-gains specification. We additionally include state and year fixed effects, or state×year fixed effects, to absorb state-specific taxes, transaction fees, management costs, and demand factors. Standard errors are robust and clustered at the MSA level.

Table 6 reports logit coefficient estimates for equation (21). Rental yields are negatively and significantly associated with the probability of a short-term sale, while the interaction with Post_t is also negative and significant, indicating a more negative sensitivity in the post-GFC period. Results are highly consistent across the two- and three-year resale horizons.[17]

[15] Because the dataset is large, we use a random 50% sample of transactions to reduce computation time.

[16] As robustness checks, we also compute rental yield using either the MSA rent-to-price ratio averaged over a two-year window or the ratio in the year beginning one year after purchase; the results are very similar.

[17] Appendix B reports year-by-year logit estimates. The same patterns hold within each boom: sensitivities to capital gains are stronger pre-GFC, while sensitivities to rental yields are more negative post-GFC.

6.5 Model versus Empirical Sensitivities of Short-Term Sales

Figure 4 panel (a) reports the model-implied effects of capital gains on the share of short-term sales. In the mortgage (pre-GFC) environment, the slope is 1.46 percentage points per one-percentage-point increase in annual price growth, while in the bond (post-GFC) environment it is only 0.23 percentage points. Under loose credit conditions, investors operate close to their resale threshold, so expected appreciation strongly raises short-term resale activity. By contrast, when outside asset returns are low, investors have stronger incentives to hold properties and rely more on rental income, making resale much less responsive to capital gains. Panel (b) reports the estimated probability of a short-term sale as a function of capital gains. In the pre-GFC period, the slope is 1.72 percentage points per one-percentage-point increase in annual price growth, while in the post-GFC period it is essentially zero.

Figure 5 panel (a) reports the model-implied effects of rental yields on the share of short-term sales. In the mortgage environment, the slope is -0.82 percentage points, whereas in the bond environment it becomes substantially more negative (-1.64 percentage points). Higher rental profitability makes holding properties more attractive, increasingly discouraging short-term resale when investors place greater weight on rental income relative to resale gains. Panel (b) reports the estimated probability of a short-term sale as a function of rental yields. In the pre-GFC period, the slope is -1.35 percentage points and becomes substantially more negative in the post-GFC period (-2.25 percentage points). As with capital gains, the three-year-horizon empirical results are very similar.[18]

Taken together, the model and empirical results point to a clear and economically intuitive shift in investor strategies around the GFC. Before the GFC, short-term resale activity was highly sensitive to expected capital gains and comparatively less sensitive to rental

[18] For the three-year resale horizon, the empirical slopes are 1.67 percentage points pre-GFC and approximately zero post-GFC for capital gains, and -1.46 percentage points pre-GFC and -2.46 percentage points post-GFC for rental yields. Reported slopes are evaluated at the mean price growth and rental yield in Table 4.

yields, consistent with resale-oriented strategies focused on timing appreciation.[19] After the GFC, short-term sales became far less responsive to capital gains and substantially more sensitive to rental yields, reflecting the rise of buy-and-hold behavior in which returns are driven primarily by rental cash flows. That the mortgage and bond environments reproduce this empirical shift provides support for the model's underlying mechanisms. It also reinforces the interpretation that changes in macro-financial conditions, such as credit availability and outside asset returns, fundamentally reshaped investor strategies.

7 Model Results: Boom Experiments and Investor Strategies

This section uses the model to study how housing investors adjust investment strategies, leverage, and portfolio composition during housing booms under alternative financing environments. We conduct counterfactual boom experiments that impose identical house price paths across model variants, while allowing rents and financing conditions to differ in ways that capture key institutional features of the pre- and post-GFC periods.

The section proceeds in four steps. We first describe the design of the housing boom experiments and the financing regimes considered. We then analyze investor behavior in the mortgage environment, where credit conditions vary and short-term resale strategies may emerge. Next, we study the bond environment, where changes in rental income and financial returns shape long-term, rent-oriented investment behavior. Finally, we examine how these strategies respond to a subsequent house price bust.

7.1 Housing Boom Experiments

We formalize the pre- and post-GFC boom experiments by specifying paths for house prices and rents common to all model environments, combined with alternative financing condi-

[19] See Haughwout et al. (2011) for evidence of speculative investor behavior prior to the GFC.

tions that capture the macro-financial environment of the 2000s and 2010s.

These paths are displayed in Figure 6. In both experiments, house prices grow steadily (panel (a)), implying a cumulative increase of 38% over six years, consistent with the empirical evolution of real house prices shown in Figure 1, panel (a). Rental prices evolve differently across experiments, shaping the path of the price-to-rent ratio. In the pre-GFC experiment, rental prices remain flat (panel (b)),[20] causing the price-to-rent ratio to rise sharply from 15 to 20.8 (panel (c)). In the post-GFC experiment, rental prices increase by 25% over the six-year period, limiting the rise in the price-to-rent ratio to a more moderate peak of 16.6. Both peaks closely mirror the empirical patterns in Figure 1, panel (b).

We study two alternative financing conditions. In the mortgage environment, credit conditions loosen in the pre-GFC experiment, with the LTV limit rising from 80% to 100% and the mortgage interest rate falling from 6% to 4% (panels (d)-(e) of Figure 6), while credit conditions remain normal in the post-GFC experiment. In the bond environment, real bond returns decline from 5.3% to -2.9% over six years in the post-GFC experiment (panel (f)), while remaining high in the pre-GFC experiment. These patterns are consistent with credit conditions and bond yield movements in Figure 1, panels (c) and (d).

By holding house price growth constant across experiments, we isolate the roles of rents (and therefore the price-to-rent ratio) and financing conditions. In the mortgage environment, sharp increases in the price-to-rent ratio coincide with relaxed credit, capturing key features of the pre-GFC housing boom. In the bond environment, rising rental income coincides with declining bond yields, reflecting the post-GFC financial environment. Appendix A provides details on the simulation setup and the computation of initial conditions.

[20] Aggregate real rents during the 2000s housing boom were essentially flat: the Rent of Primary Residence Index deflated by CPI All Items Less Shelter rose by only about 2% between 2002 and 2007.

7.2 The Emergence of Short-Term Investors

Using the housing boom experiments described above, we study how investors respond to changes in credit conditions and price-to-rent ratios in the mortgage environment. When credit loosens and the price-to-rent ratio rises sharply, investors endogenously shift toward capital-gains-oriented, flipping strategies: short-term sales surge, leverage-financed purchases expand, holding periods shorten, and rental activity collapses.

7.2.1 Model Results

We first examine investor behavior during the pre-GFC housing boom in the mortgage environment. Figure 7 shows the resulting dynamics. Panel (a) shows a sharp increase in sales activity: the share of the housing stock sold rises from 13% at the onset of the boom to 35% after six years, an increase of 22 percentage points. This reflects a shift toward faster resale and shorter holding periods as investors seek to capitalize on rising prices. Panel (b) shows that the share of the housing stock rented declines from 77% to 59% over the boom period (a drop of 18 percentage points), reflecting a reallocation away from rental income toward short-term resale strategies. Not renting preserves the option to sell in the next period, so this decline signals a growing inventory of units available for resale.

A key outcome is the surge in short-term sales, defined in the model as sales of properties held for three years or less (see equation (13)). Panel (c) shows a 20 percentage point increase in this share, from 5% to a peak of 25%. This pattern closely mirrors the empirical pattern in Figure 2, panel (a), supporting the model's underlying mechanisms.

The relaxation of credit fuels a surge in new purchases and mortgage originations during the early phase of the boom, enabling leverage-driven growth. Panel (d) shows that total housing holdings expand rapidly over the first half of the boom, peaking at 153% above the steady-state level, before unwinding as new purchases decline and investor sales intensify toward the price peak.[21] Panel (e) shows that total mortgage balances follow a similar pattern, reaching a peak of 280% above the steady-state level and then declining as new originations taper off and investors deleverage. As a result, leverage (measured by the ratio of mortgage debt to total housing value) rises from 0.66 in the steady state to 0.86 midway through the boom, before easing later as debt exposure is reduced.

Panel (f) shows that the mean holding period in the housing portfolio declines sharply, from 7.1 to 3.5 years, marking the emergence of short-term investment behavior. This shift reflects investors' response to rising house prices and easier credit, which incentivize resale over long-term rental and accelerate market turnover. Later in the boom, the average holding period rises gradually as new purchases taper off, but the recovery is modest (reaching only 4.6 years by the end of the boom), indicating that short-term behavior persists.

The decline in average holding periods does not arise from sales alone, but from the interaction between new purchases and subsequent resale behavior. Sold units are randomly drawn from the investor's portfolio, so higher sales activity by itself does not mechanically change the distribution of holding periods or increase the share of short-term sales.[22] Instead, the shift toward short-term behavior is driven by the surge in new housing purchases, which replace older holdings and are quickly resold. As a result, the share of short-term sales rises together with the decline in the mean holding period: flipping emerges as newly acquired properties cycle through investors' portfolios much faster than in the steady state.

We now examine investor behavior during the post-GFC housing boom, characterized by identical house price growth but rising rents and normal credit conditions. This combination yields a more muted increase in the price-to-rent ratio and higher rental yields. Despite house prices growing at the same rate as in the pre-GFC boom, short-term resale activity does not emerge. Panel (a) shows that sales rise only modestly (by 5 percentage points over six years) compared with a 22 percentage point increase during the pre-GFC boom. In the

[21] See Leamer (2007) for empirical evidence that housing market quantities often peak before prices.

[22] See Appendix A for details on how we track the distribution of holding periods.

post-GFC experiment, rental activity remains stable: the rental share declines only slightly, from 77% to 73%, unlike the sharp 18-point drop in the pre-GFC episode (panel (b)).

Consistent with this, the share of short-term sales remains essentially flat throughout the boom (panel (c)), unlike the fivefold increase observed under pre-GFC conditions. Subdued new purchases and mortgage originations lead to a gradual decline in total housing holdings (panel (d)) and mortgage balances (panel (e)). As a result, the mean holding period rises from 7.1 to 8.4 years (panel (f)), since sales are no longer offset by new purchases.

These responses reflect a long-term, rent-and-hold strategy that remains robust throughout the boom. Muted increases in the price-to-rent ratio (implying higher rental yields) sustain incentives to rent, largely offsetting gains from short-term resale. Absent credit loosening in this post-GFC experiment, investors lack an alternative source of financing to expand new purchases as prices rise, and housing accumulation remains subdued. In the bond environment, discussed next, we allow for such expansion by permitting investors to finance housing purchases through portfolio reallocation away from bond holdings.

Overall, these findings establish a central result: large increases in the price-to-rent ratio combined with relaxed credit conditions, as observed in the pre-GFC boom, are necessary for the emergence of short-term flipper investors. When credit does not loosen and rents rise alongside prices, as in the post-GFC boom, investors instead pursue long-term, rent-based strategies. Table 7 summarizes the boom experiments in the mortgage environment.

7.2.2 Returns and Leverage

Table 7 reports realized internal rates of return (IRRs) for housing investments initiated at $t = 1$, accounting for adjustment costs, depreciation, and endogenous rental and sale decisions, assuming that any remaining housing stock is liquidated at $t = 6$. In the pre-GFC boom, the realized levered IRR reaches 28%, compared with an unlevered counterfactual IRR of 8% (see Appendix A for details). By contrast, in the post-GFC boom (where no new borrowing occurs) the IRR is only 11%. These results indicate that, under loose credit conditions, short-term resale strategies generate higher returns through leverage. Such debt-amplified short-run gains were prominent in housing markets during the 2000s.

7.2.3 The Role of Credit

To isolate the role of credit, we conduct a boom experiment that holds fixed the pre-GFC paths of rising house prices and flat rents while maintaining a low LTV limit and a high mortgage rate. Relative to the baseline with credit loosening, new housing purchases decline. The share of houses sold increases by only 10 percentage points, and short-term sales rise by just 3 percentage points. The mean holding period rises to 8.3 years rather than declining, as sales are no longer offset by new purchases. As a result, total housing holdings contract by 43%. These results show that relaxed credit conditions are crucial for the emergence of flipping and leverage-driven expansion in this price-to-rent environment.

Table 7 also reports results from boom experiments initialized from the ergodic distribution; the main insights are unchanged (see Appendix A for details).

7.3 The Emergence of Long-Term Investors

Using the boom experiments described above, we examine how investors respond to changes in bond returns and price-to-rent ratios in the bond environment. When bond returns fall and rents rise sharply, investors endogenously shift toward long-term, rent-based strategies and reallocate portfolios away from financial assets and toward real estate. As a result, holding periods lengthen and housing positions expand.

7.3.1 Model Results

Figure 8 illustrates how investors respond to the post-GFC housing boom. A central margin of adjustment is portfolio allocation. The share of investors' total portfolio value allocated to housing increases gradually over the boom, rising from 7% in the steady state to 71% by the end of the six-year period. This reallocation reflects both a decline in bond holdings and a sustained expansion of housing positions, consistent with the post-GFC return environment of falling bond returns and rising rental income, which favors housing investment.

This portfolio shift is accompanied by a change in the use of housing assets. Panel (a) shows that the share of the housing stock rented increases gradually from 33% at the start of the boom to a peak of 84% (a rise of 51 percentage points) and remains elevated thereafter. This pattern reflects a reorientation toward buy-and-hold strategies centered on rental income. In the model, rented houses are ineligible for sale in the following period, so higher rental activity reduces the inventory of units available for potential resale, limiting the scope for realizing capital gains even as house prices continue to rise.

Consistent with this shift in investment strategy, short-term resale declines markedly. Panel (b) shows that the share of the housing stock sold falls by 54 percentage points (from 60% in the steady state to 6% after six years), while panel (c) shows a similar decline in short-term sales. In contrast to the resale-driven dynamics in the pre-GFC experiment of the mortgage environment, investors hold properties for longer and rent them out more frequently. Despite substantial price appreciation, investors respond to changes in the return environment by favoring stable rental income over short-term resale gains.

Asset holdings also evolve markedly over the boom. Panel (d) shows that total housing holdings rise steadily (by a factor of 7.6 relative to their initial level) over the six-year period. At the same time, panel (e) shows that bond holdings decline gradually, falling to 28% of their initial level by the end of the boom. As wealth is reallocated out of bond holdings, new housing purchases increase throughout the boom, peaking at 5.8 times their initial level.

Despite this strong purchase margin, the mean holding periods lengthen, from 1.7 to over 2.3 years (panel (f)). Given the substantial resale activity in the initial conditions, continued purchases would typically reduce holding periods; instead, the rise reflects that newly acquired units are retained as long-term rental assets rather than resold in the short term.

We contrast this behavior with the pre-GFC housing boom, characterized by the same house price growth but flat rents and high bond returns. This setting generates an increase in the price-to-rent ratio and a decline in rental yields. In this case, long-term rental strategies do not emerge, and the share of houses rented falls from 33% to nearly zero, compared with a 51 percentage point increase in the post-GFC boom (panel (a)). The share of houses sold rises by 35 percentage points (from 60% to 95%), and short-term sales increase by 41 percentage points (panels (b)-(c)). The share of portfolio value allocated to housing declines slightly, in contrast to the tenfold increase observed in the post-GFC boom (from 7% to 71%).

In the pre-GFC boom, panel (d) shows that total housing holdings decline by 39% relative to their initial level, while new housing purchases remain close to their initial value throughout the boom. Bond holdings increase modestly by about 5% by the end of the boom (panel (e)). Panel (f) shows that the mean holding period falls from 1.7 years to just above one year, driven by steady new purchases combined with a rise in sales activity during the boom.

In the absence of incentives to reallocate out of bonds and with low rental yields, investors respond to rising house prices through short-term resale, keeping houses vacant for rapid sale rather than adopting buy-and-hold strategies. High bond returns discourage investors from unwinding bond positions, leaving no alternative source of financing to support an expansion in new purchases as prices rise. By contrast, in the mortgage environment discussed above, leverage relaxes this constraint and allows housing positions to expand.

Taken together, these findings highlight a key result: large rent increases that limit the rise in the price-to-rent ratio, combined with falling bond returns, as in the post-GFC boom, are necessary for the emergence of long-term, rent-oriented investors. When bond returns remain high and rents stay flat while prices rise, as in the pre-GFC boom, investors instead pursue short-term resale strategies. Table 8 summarizes the bond environment experiments.

7.3.2 Returns and Portfolio Reallocation

Table 8 reports realized IRRs for housing investments initiated at $t = 1$, computed as in the mortgage environment. It also reports the IRR on a bond investment initiated at $t = 1$, with proceeds reinvested annually until $t = 6$ (see Appendix A for details). In the post-GFC boom, the realized housing IRR is 6.0%, substantially exceeding the 1.2% realized IRR on bonds. In contrast, during the pre-GFC boom, the housing IRR is 6.1% and much closer to the 5.3% IRR on bonds, although housing returns are riskier ex ante. These results indicate that, in a housing boom with rising rents and falling bond returns, reallocating from bonds to housing under a long-term rental strategy yields higher returns.

7.3.3 The Role of Bond Yields

To assess the role of bond yields, we conduct a boom experiment that holds fixed the post-GFC paths of rising house prices and rents while keeping bond returns at their initial high level. Relative to the baseline with declining bond yields, new housing purchases fail to expand and the shift toward long-term rental strategies disappears. The share of houses rented falls by 26 percentage points, while the share sold rises by 29. Mean holding periods remain short at 1.1 years rather than increasing, and total housing holdings contract by 39% instead of expanding. These findings indicate that declining bond returns are essential for the emergence of long-term, rent-and-hold investors within our calibration.

Table 8 also reports results from boom experiments initialized from the ergodic distribution; the main insights are similar (see Appendix A for details).

7.4 Housing Price Busts and Investor Strategies

Finally, we study how the investment strategies that emerge during the housing booms respond to a subsequent house price bust. The objective is to assess which strategies are fragile and which remain resilient when price dynamics and financing conditions reverse.

We proceed in two steps. First, in the mortgage environment, we introduce a house price bust following the pre-GFC boom. After six years of price growth ($t = 1$ to $t = 6$), prices begin to decline at $t = 7$ and revert toward their pre-boom level, reaching it four years after the peak.[23] Credit conditions tighten, with the LTV limit and mortgage rates returning to their pre-boom levels, while rental prices remain fixed. Second, in the bond environment, we introduce an analogous bust after the post-GFC boom. Bond returns gradually increase, partially reversing the decline observed during the boom, while rental prices remain at their peak level. Figures 9 and 10, panels (a)-(b), summarize the bust experiments.[24]

7.4.1 Mortgage Environment: Bust Dynamics

Figure 9 shows how investor strategies adjust to a house price bust in the mortgage environment. The reversal in prices, combined with tighter credit conditions, triggers a contraction in short-term resale activity. The share of short-term sales declines from 25% at the peak of the boom to 2% four years into the bust, a drop of 23 percentage points (panel (c)). Rent-oriented strategies strengthen: the share of the housing stock rented rises from 59% to 81%, an increase of 22 percentage points (panel (d)). The mean holding period rises from 4.7 to 6.3 years over the same period, reflecting reduced turnover and subdued new purchases.

The contraction in short-term selling reflects several reinforcing incentives. Selling at depressed prices reduces the net equity recovered upon sale, as mortgage balances do not

[23] Using the Case-Shiller Index for real house prices (deflated by CPI All Items Less Shelter), prices peak in 2006 and revert close to their pre-boom (2001) level after four years, about 4% above it.

[24] In both environments, bust experiments start from the endogenously generated state (portfolio composition, leverage, and investment strategies) reached at the peak of the preceding boom.

adjust with house values, weakening incentives to liquidate vacant units. Tighter LTV limits restrict access to new borrowing, making it costlier to replace sold properties. Lower expected resale values and activity, together with higher expected rental yields, further tilt incentives toward operating housing assets rather than selling them.

7.4.2 Bond Environment: Bust Dynamics

Figure 10 shows how investor strategies adjust to a house price bust in the bond environment. The decline in prices, together with the gradual reversion of bond yields, does not lead to a significant unwinding of long-term investment strategies: rental activity remains resilient. The share of the housing stock rented declines modestly, from 84% to 76% (panel (c)), and remains elevated relative to its pre-boom level (33%). Resale activity increases slightly: the share of short-term sales rises from 5% to 14% (panel (d)). Despite this increase, the mean holding period continues to lengthen, from 2.3 to 3.1 years, reflecting a contraction in new purchases as sold units are replaced more slowly, leading to portfolio consolidation.

The persistence of rental strategies is supported by several incentives. Rental cash flows remain high and stable during the bust, and rental yields increase, sustaining the continuation value of operating housing assets. Although bond yields recover following the bust, the adjustment is gradual, reducing incentives to unwind housing positions sharply. As a result, investors engage in limited portfolio rebalancing rather than a full reversal of strategy.

7.4.3 Bust Dynamics: Summary

These results reveal a clear asymmetry in the resilience of housing investment strategies. In the mortgage environment, short-term, leverage-financed flipping strategies unravel when house prices fall and credit conditions tighten. In contrast, in the bond environment, long-term, rent-oriented strategies remain largely intact despite falling house prices and improving outside returns. Initializing from the ergodic distribution yields similar dynamics.

8 Conclusions

This paper develops a dynamic portfolio model in which housing investors' holding periods emerge endogenously from macro-financial conditions. A selling margin and a vacancy-rental margin link expected capital gains, rental yields, credit availability, and outside asset returns to turnover and investment horizons. Thus, flippers and buy-and-hold landlords do not simply reflect fixed investor types, but also arise in response to fundamentals.

Using U.S. transaction-level deed records from 2000 to 2017, we document a sharp shift in investor behavior across the pre- and post-GFC housing booms. Despite comparable house price growth, pre-GFC investors relied heavily on mortgage credit and resold quickly, whereas post-GFC investors relied far more on cash purchases and held properties longer. Calibrated versions of the model reproduce this shift: a rising price-to-rent ratio and loose mortgage credit generate flipping, whereas higher rental yields and low safe returns generate buy-and-hold behavior. The model also captures the weaker sensitivity of short-term resale to capital gains and the stronger role of rental yields in the post-GFC boom.[25]

A key asymmetry emerges in bust dynamics. Leverage-driven, short-term strategies collapse when prices fall and credit tightens, whereas long-term, rent-oriented strategies prove more resilient. These findings suggest that the composition of investor strategies has first-order consequences for how housing markets respond to downturns and the transmission of monetary policy. When investor purchases are largely cash-financed, changes in interest rates operate less via borrowing costs and more via portfolio reallocation and holding-period adjustments. Studying these channels in general equilibrium is left to future research.

[25] The post-pandemic housing boom of the early 2020s, marked by rapid price growth, low interest rates, and subsequent monetary tightening, provides a natural setting to assess the model's predictions in future work.

References

Albanesi, S., G. DeGiorgi, and J. Nosal (2022). Credit growth and the financial crisis: A new narrative. *Journal of Monetary Economics 132*, 118–139.

Allen, M. T., J. Rutherford, R. Rutherford, and A. Yavas (2018). Impact of investors in distressed housing markets. *The Journal of Real Estate Finance and Economics 56*(4), 622–652.

Amherst Capital Management LLC (2018, April). An update on institutional single-family rental activity: 2017/2018 U.S. market trends support long-term growth & opportunity. Market commentary, Amherst Capital Management LLC.

Andersen, S., C. Badarinza, L. Liu, J. Marx, and T. Ramadorai (2022). Reference dependence in the housing market. *American Economic Review 112*(10), 3398–3440.

Bayer, P., C. Geissler, K. Mangum, and J. W. Roberts (2020). Speculators and middlemen: The strategy and performance of investors in the housing market. *The Review of Financial Studies 33*(11), 5212–5247.

Bhutta, N. (2015). The ins and outs of mortgage debt during the housing boom and bust. *Journal of Monetary Economics 76*, 284–298.

Boddin, D., D. te Kaat, C. Ma, and A. Rebucci (2023). A housing portfolio channel of QE transmission. Working paper.

Campbell, J. Y. and J. F. Cocco (2015). A model of mortgage default. *The Journal of Finance 70*(4), 1495–1554.

Chinco, A. and C. Mayer (2016). Misinformed speculators and mispricing in the housing market. *The Review of Financial Studies 29*(2), 486–522.

Chow, G. C. and A.-l. Lin (1971). Best linear unbiased interpolation, distribution, and extrapolation of time series by related series. *The Review of Economics and Statistics 53*(4), 372–375.

Daniel, K., L. Garlappi, and K. Xiao (2021). Monetary policy and reaching for income. *The Journal of Finance 76*(3), 1145–1193.

DeFusco, A. A., C. G. Nathanson, and E. Zwick (2022). Speculative dynamics of prices and volume. *Journal of Financial Economics 146*(1), 205–229.

Demers, A. and A. L. Eisfeldt (2022). Total returns to single-family rentals. *Real Estate Economics 50*(1), 7–32.

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

Elenev, V., T. Landvoigt, and S. Van Nieuwerburgh (2021). A macroeconomic model with financially constrained producers and intermediaries. *Econometrica 89*(3), 1361–1418.

Gao, Z., M. Sockin, and W. Xiong (2020). Economic consequences of housing speculation. *The Review of Financial Studies 33*(11), 5248–5287.

Garcia, D. (2022). Second-home buying and the housing boom and bust. *Real Estate Economics 50*(1), 33–58.

Gargano, A. and M. Giacoletti (2024, May). Individual investors' housing income and interest rate fluctuations. SSRN Working Paper No. 4251621.

Garriga, C., P. Gete, and A. Tsouderou (2023). The economic effects of real estate investors. *Real Estate Economics 51*(3), 655–685.

Ghent, A. (2012). Infrequent housing adjustment, limited participation, and monetary policy. *Journal of Money, Credit and Banking 44*(5), 931–955.

Gilbukh, S. and P. Goldsmith-Pinkham (2024). Heterogeneous real estate agents and the housing cycle. *The Review of Financial Studies 37*(11), 3431–3489.

Glaeser, E. L. and C. G. Nathanson (2017). An extrapolative model of house price dynamics. *Journal of Financial Economics 126*(1), 147–170.

Greenwald, D. L. and A. Guren (2025). Do credit conditions move house prices? *American Economic Review 115*(10), 3559–3596.

Gupta, A., C. Hansman, and P. Mabille (2025). Financial constraints and the racial housing gap. *Journal of Financial Economics 173*, 104142.

Gurun, U. G., J. Wu, S. C. Xiao, and S. W. Xiao (2023). Do Wall Street landlords undermine renters' welfare? *The Review of Financial Studies 36*(1), 70–121.

Han, L. and S.-H. Hong (2024). Cash is king? Understanding financing risk in housing markets. *Review of Finance 28*(6), 2083–2118.

Haughwout, A., D. Lee, J. S. Tracy, and W. Van der Klaauw (2011). Real estate investors, the leverage cycle, and the housing market crisis. Staff Report 514, Federal Reserve Bank of New York.

Jeske, K., D. Krueger, and K. Mitman (2013). Housing, mortgage bailout guarantees and the macro economy. *Journal of Monetary Economics 60*(8), 917–935.

Judd, K. L., F. Kubler, and K. Schmedders (2002). A solution method for incomplete asset markets with heterogeneous agents. Working Paper, Stanford University.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2019). Credit supply and the housing boom. *Journal of Political Economy 127*(3), 1317–1350.

Korevaar, M. (2023). Reaching for yield and the housing market: Evidence from 18th-century Amsterdam. *Journal of Financial Economics 148*(3), 273–296.

Lambie-Hanson, L., W. Li, and M. Slonkosky (2022). Real estate investors and the U.S. housing recovery. *Real Estate Economics 50*(6), 1425–1461.

Leamer, E. E. (2007). Housing is the business cycle. Working Paper 13428, National Bureau of Economic Research.

Loewenstein, L. and P. S. Willen (2023). House prices and rents in the 21st century. Working Paper 31013, National Bureau of Economic Research.

Mabille, P. (2023). The missing homebuyers: Regional heterogeneity and credit contractions. *The Review of Financial Studies 36*(7), 2756–2796.

Mills, J., R. Molloy, and R. Zarutskie (2019). Large-scale buy-to-rent investors in the single-family housing market: The emergence of a new asset class. *Real Estate Economics 47*(2), 399–430.

Piazzesi, M. and M. Schneider (2009). Momentum traders in the housing market: Survey evidence and a search model. *American Economic Review 99*(2), 406–11.

Rouwenhorst, K. G. (1995). Asset pricing implications of equilibrium business cycle models. In T. F. Cooley (Ed.), *Frontiers of Business Cycle Research*, pp. 294–330. Princeton, NJ: Princeton University Press.

Sufi, A. and A. M. Taylor (2022). Financial crises: A survey. *Handbook of International Economics 6*, 291–340.

Figures

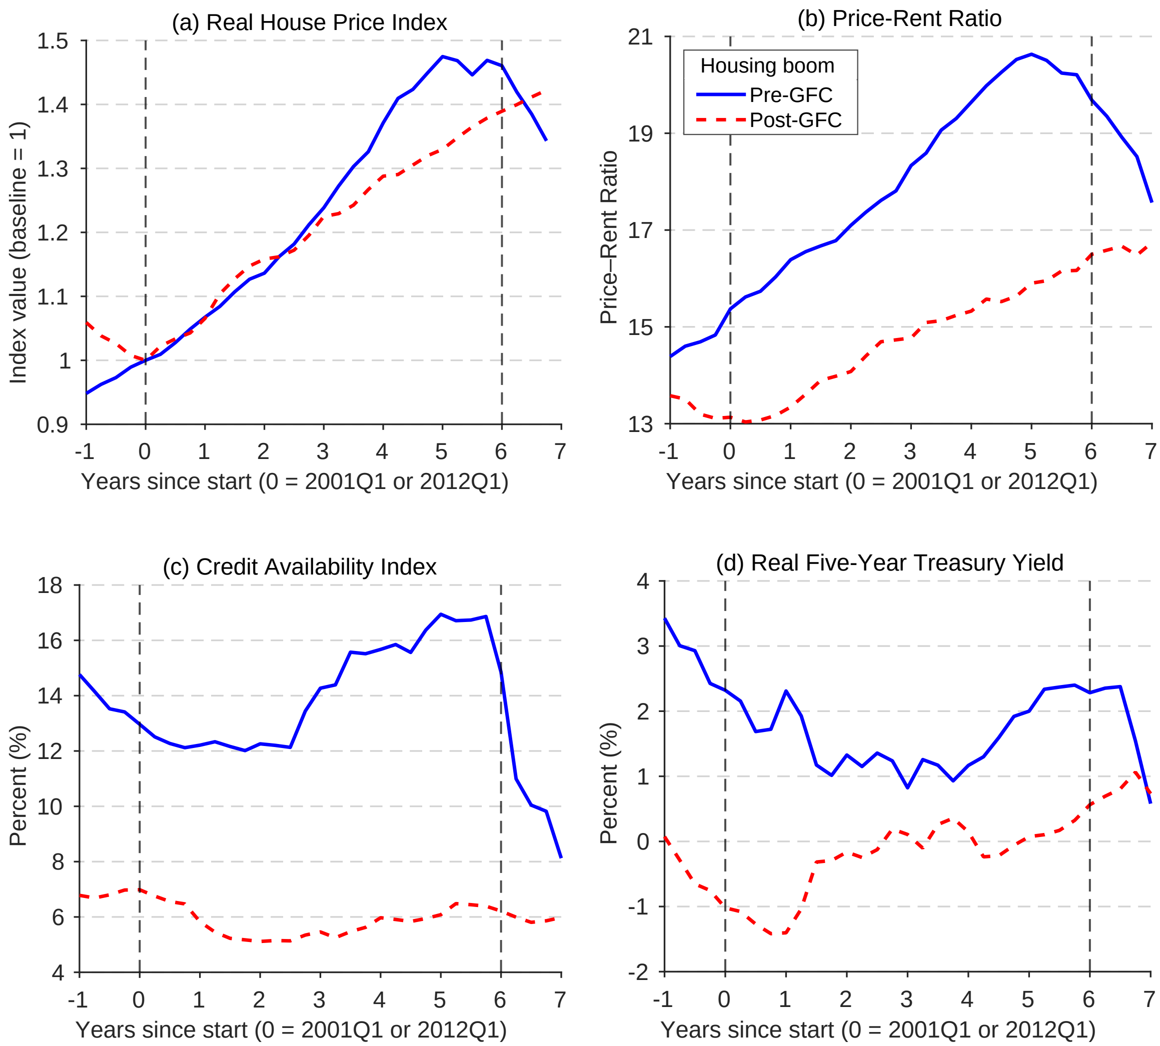

Figure 1: Macro-Financial Conditions Across Two U.S. Housing Booms

Note: Panel (a) shows the real house price index, constructed from the S&P CoreLogic Case-Shiller U.S. National Home Price Index deflated by CPI All Items Less Shelter. Panel (b) reports the price-rent ratio, measured as the ratio of the value of residential real estate owned by households to owner-occupied housing services. Panel (c) shows the Urban Institute's Housing Credit Availability Index (HCAI). Panel (d) displays the real five-year Treasury yield, using inflation-indexed yields or nominal yields adjusted for expected inflation. All series are quarterly. See Appendix B for details. Vertical dashed lines mark the start and end of each housing boom (2001Q1–2007Q1 and 2012Q1–2018Q1).

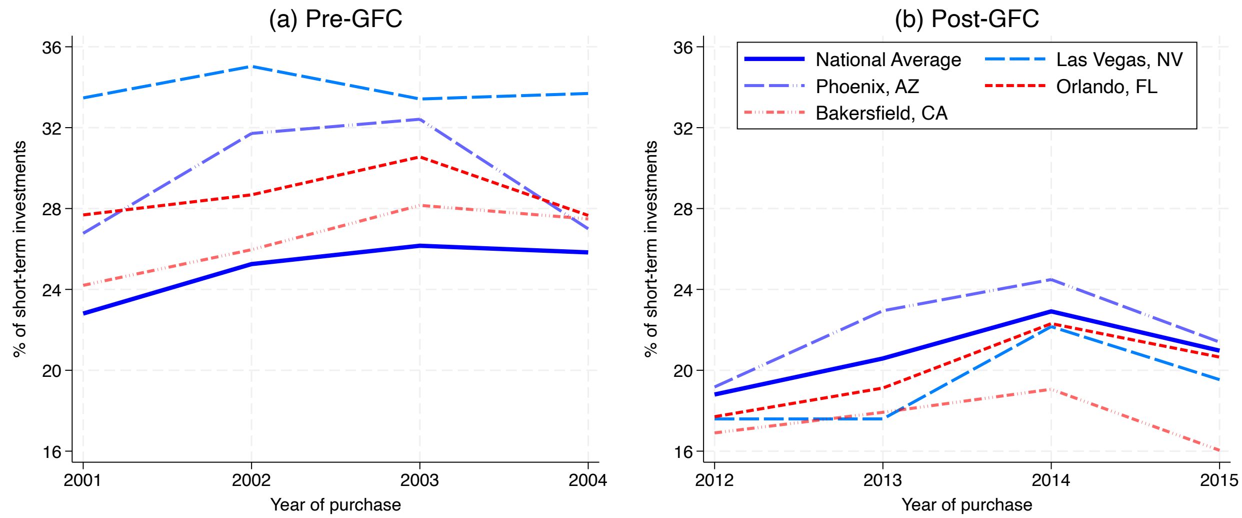

Figure 2: Short-Term Investor Sales in the Pre-GFC and Post-GFC Housing Booms

Note: This figure shows the percentage of houses purchased by retail investors in a given year that were resold within six months to three years. Panel (a) covers houses bought between 2001 and 2004, and Panel (b) covers houses bought between 2012 and 2015. Percentages are shown for the U.S. average and for selected metropolitan areas that experienced pronounced boom-bust cycles during the 2000s (e.g., DeFusco, Nathanson, and Zwick, 2022). The holding period is measured using deed-recorded purchase and sale dates from the CoreLogic dataset. Retail investors are defined as described in Section 5, with additional details in Appendix B.

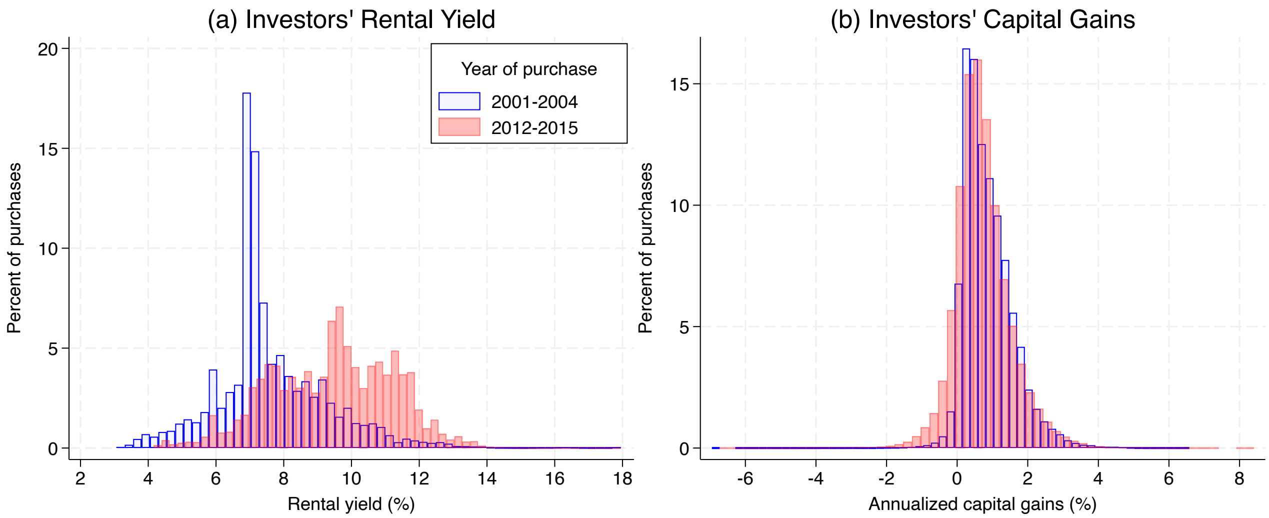

Figure 3: Distribution of Rental Yields and Capital Gains: Pre-GFC and Post-GFC Booms

Note: Panel (a) shows annual rental yields, and Panel (b) shows annualized capital gains across investor purchase locations. The histograms compare properties purchased during 2001-2004 with those purchased during 2012-2015. Investor purchases are identified in the CoreLogic deeds dataset, which provides the address needed to match each property to local price and rent indices. Rental yields and capital gains are measured using Zillow zip-code house price indices and MSA-level rent indices matched to each purchase. Details on variable construction are provided in Appendix B.

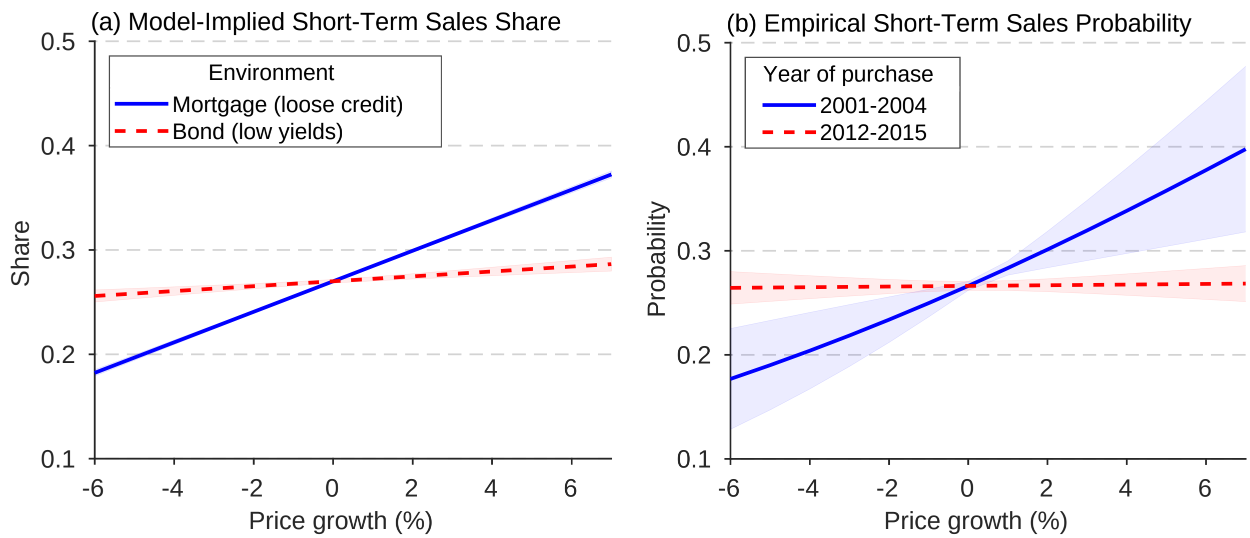

Figure 4: Short-Term Sales and House Price Growth: Model versus Empirical Evidence

Note: Panel (a) reports the model-implied short-term sale share (relative to start-of-period housing holdings) as a function of house price growth. Shares are obtained from OLS regressions on simulated data (see Section 6) and are normalized to match the empirical baseline level. Panel (b) reports the estimated probability of a two-year sale implied by the logit model (20), with price growth measured as zip-code house price appreciation during the year after purchase. Results are shown for pre-GFC (2001-2004) and post-GFC (2012-2015) purchases. Shaded areas denote 95% confidence intervals.

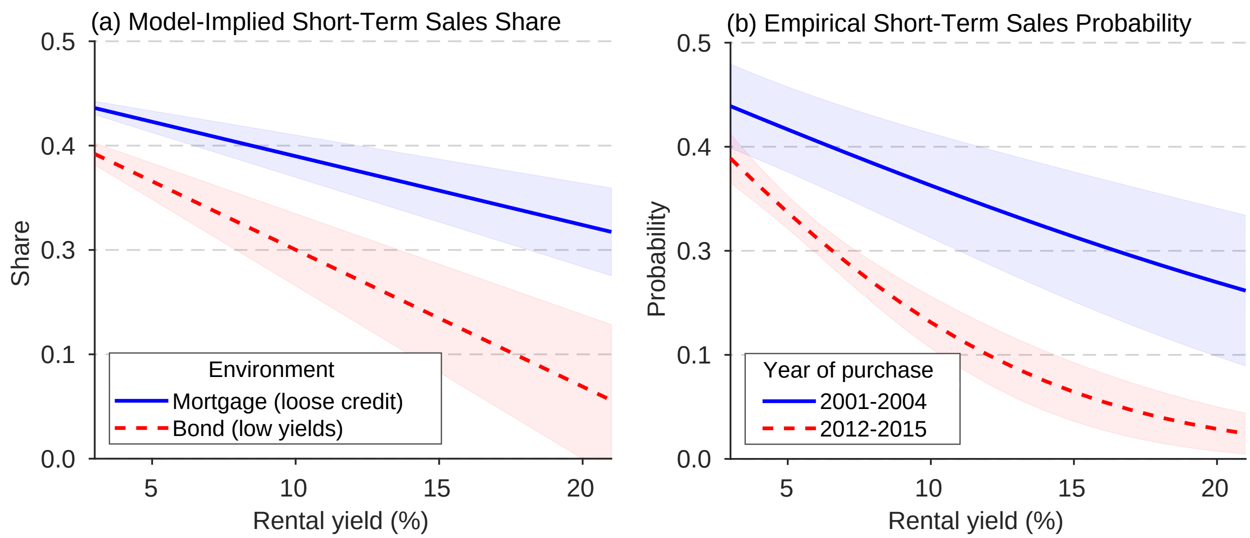

Figure 5: Short-Term Sales and Rental Yields: Model versus Empirical Evidence

Note: Panel (a) reports the model-implied short-term sale share (relative to start-of-period housing holdings) as a function of rental yields. Shares are obtained from OLS regressions on simulated data (see Section 6) and are normalized to match the empirical baseline level. Panel (b) reports the estimated probability of a two-year sale from the logit model (21), where rental yield is measured as the MSA-level rent-to-price ratio during the year after purchase. Results are shown for pre-GFC (2001–2004) and post-GFC (2012–2015) purchases. Shaded areas denote 95% confidence intervals.

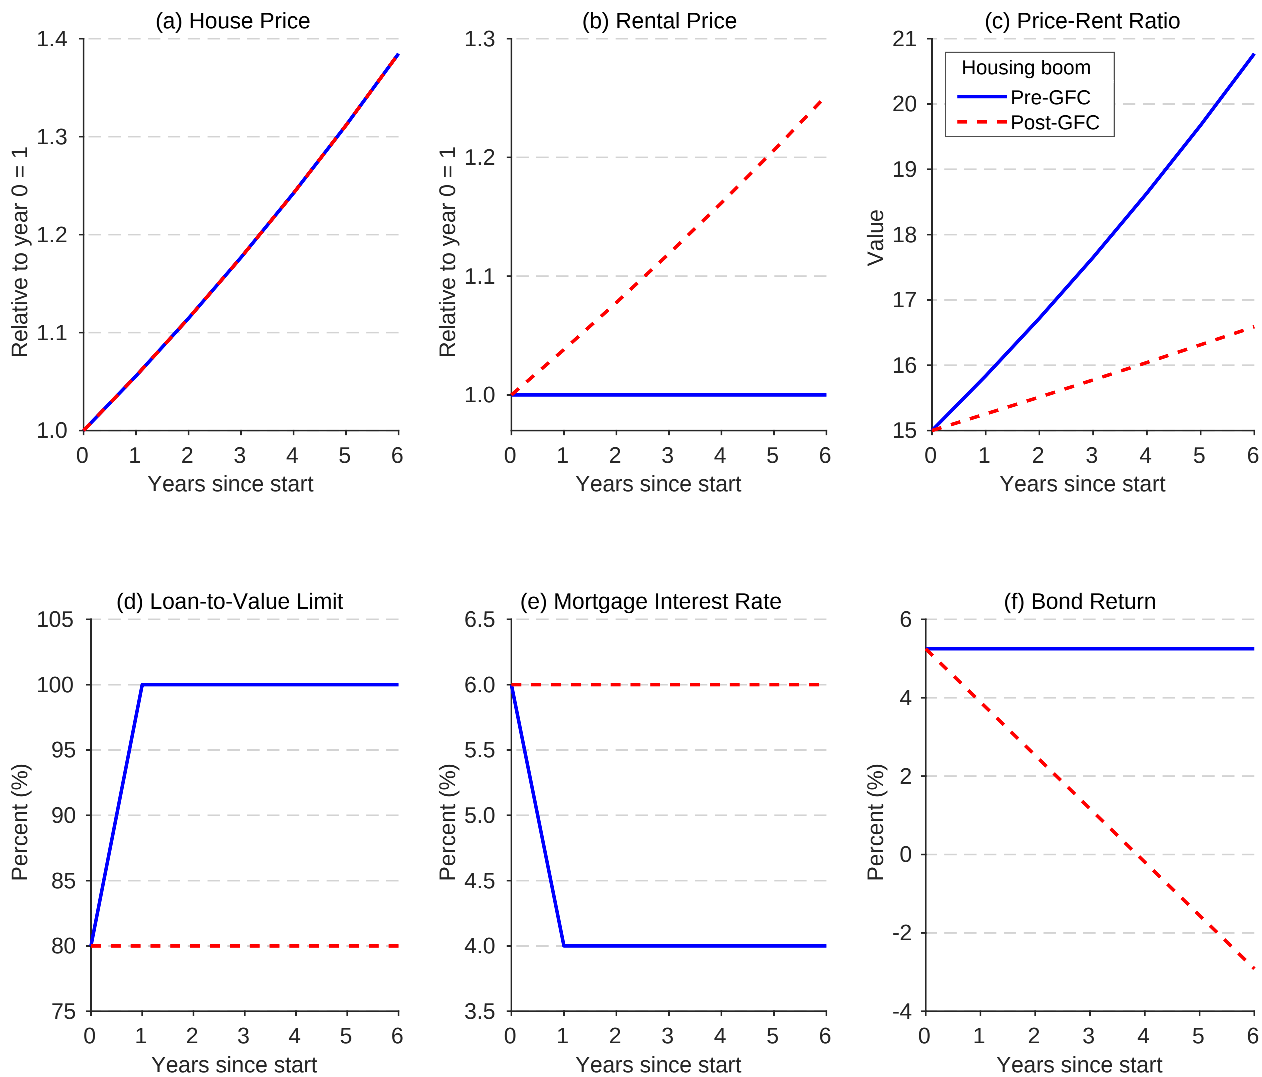

Figure 6: Simulation of Pre- and Post-GFC Housing Booms

Note: This figure shows the simulated exogenous paths used in the pre- and post-GFC boom experiments. Panels (a)-(c) display house prices, rental prices, and the price-to-rent ratio, while panels (d)-(f) show the LTV limit, the mortgage interest rate, and the real bond return. House price growth is identical across experiments. The pre-GFC boom features flat rents and looser credit conditions, whereas the post-GFC boom features rising rents and declining bond yields. House prices are normalized to 1 at $t = 0$, and time is measured in years. The discretization of the AR(1) processes for house prices, rents, and bond yields is described in Section 6 and Appendix A.

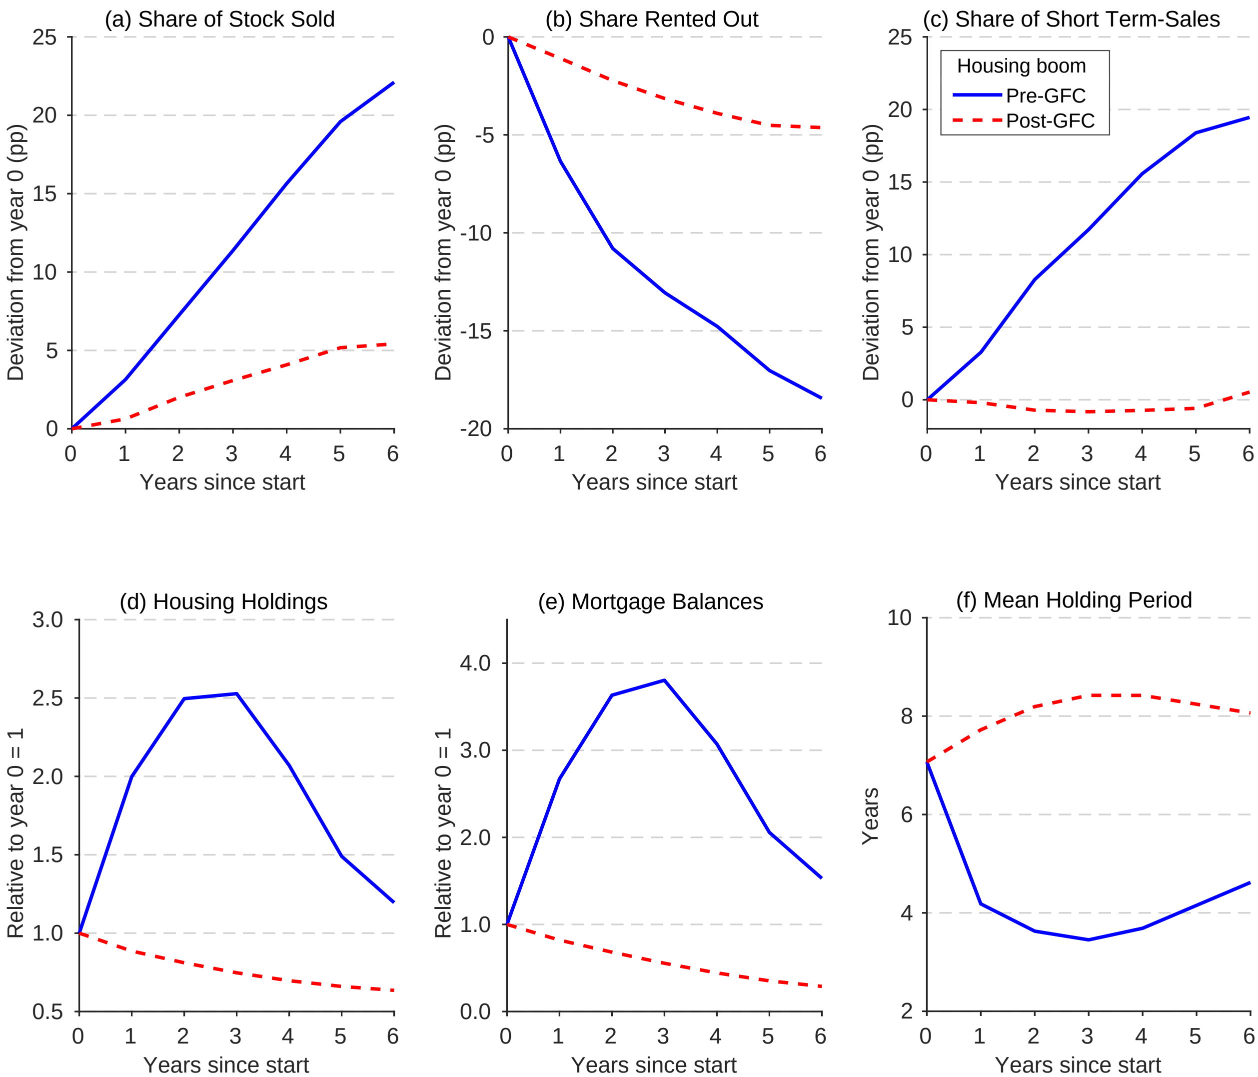

Figure 7: Investor Portfolio Dynamics in the Mortgage Environment

Note: This figure reports simulated investor responses under the pre- and post-GFC boom experiments. Panels (a)-(c) display the shares sold, rented, and resold in the short term. Panels (d)-(f) show housing holdings, mortgage balances, and the mean holding period. The pre-GFC boom features sharp increases in sales and short-term resales, declining rentals, shorter holding periods, and balance sheet expansion, whereas the post-GFC boom exhibits muted sales, stable rentals, longer holding periods, and declining housing and mortgage balances. Deviations are relative to the stochastic steady state at $t = 0$. *pp* denotes percentage points. See Section 7 and Appendix A for details.

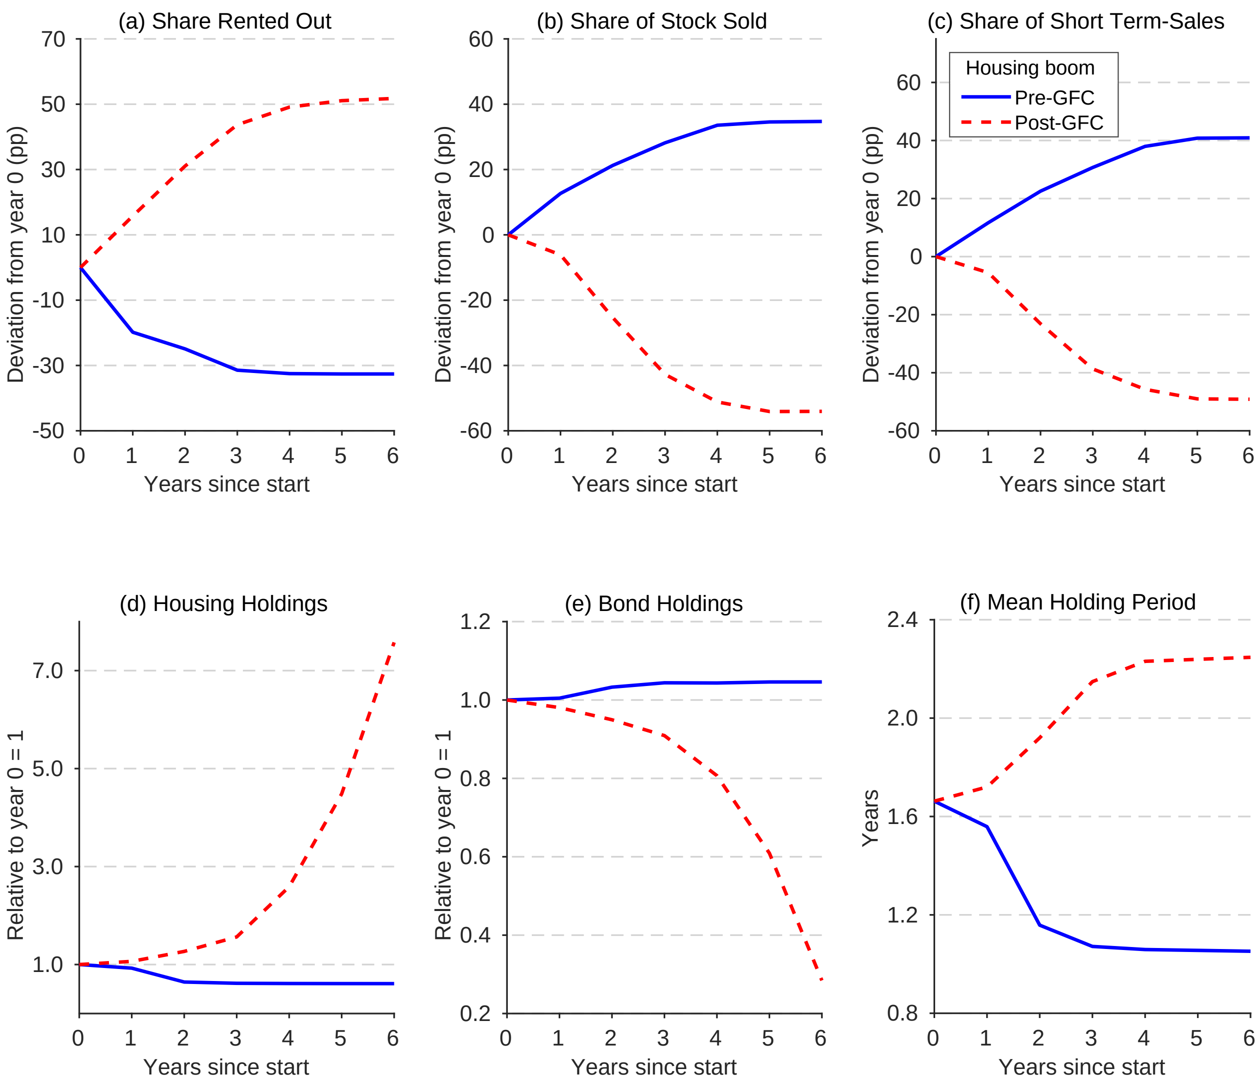

Figure 8: Investor Portfolio Dynamics in the Bond Environment

Note: This figure reports simulated investor responses under the pre- and post-GFC boom experiments. Panels (a)-(c) display the shares rented, sold, and resold in the short term. Panels (d)-(f) show housing holdings, bond holdings, and the mean holding period. The post-GFC boom features portfolio reallocation toward housing, higher rental shares, longer holding periods, and a sharp contraction in resale activity. In contrast, the pre-GFC boom shows limited portfolio reallocation, lower rentals, rising resale activity, and shorter holding periods. Deviations are relative to the stochastic steady state at $t = 0$. *pp* denotes percentage points. See Section 7 and Appendix A for details.

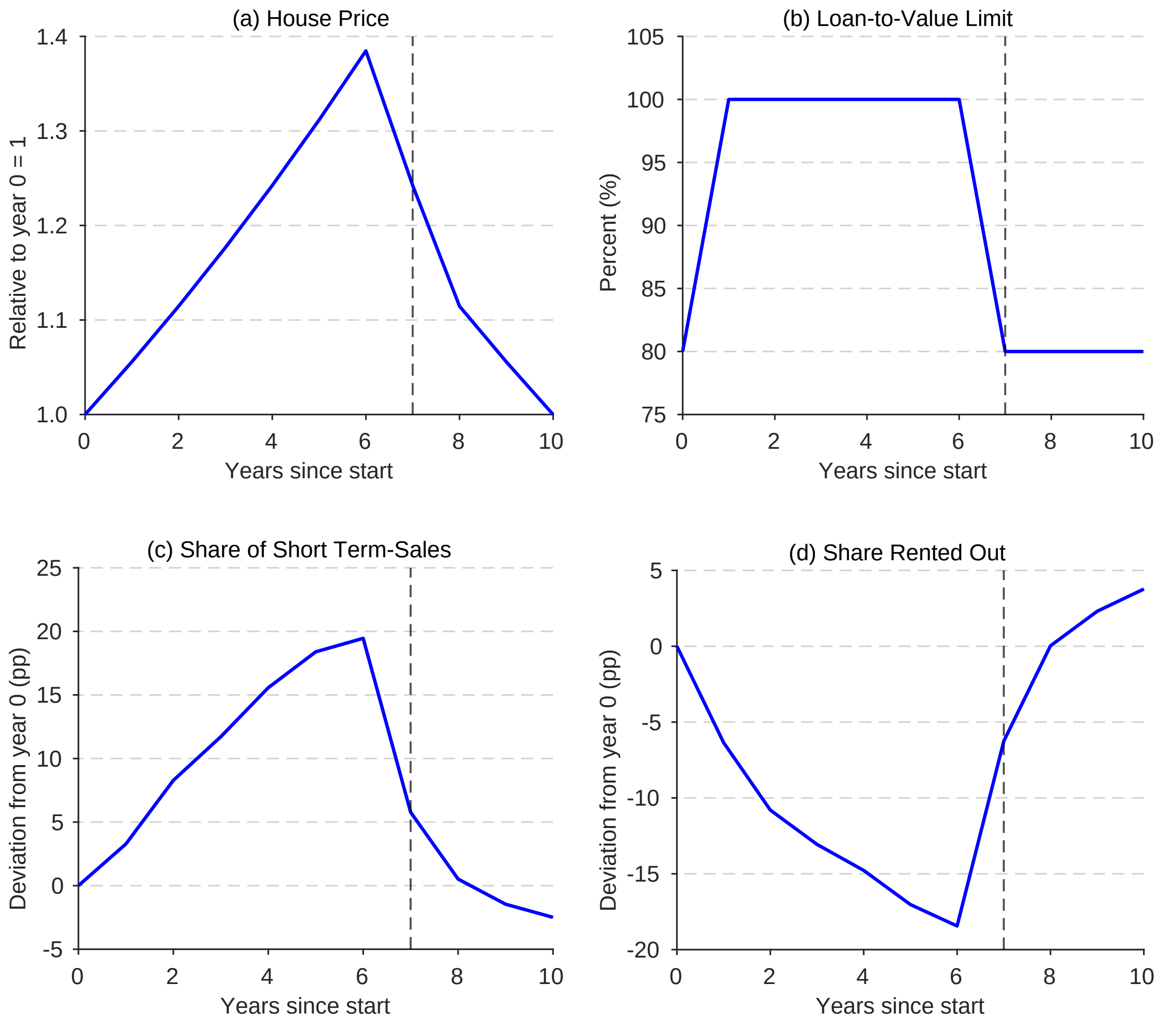

Figure 9: Boom-Bust Episode and Investor Portfolio Dynamics: Mortgage Environment

Note: Panels (a)-(b) show house prices and the LTV limit. Panels (c)-(d) show the share of short-term sales and housing units rented out. The boom phase from $t = 1$ to $t = 6$ is identical to that in the pre-GFC boom in Figures 6 and 7. At $t = 7$, house prices begin to decline and credit conditions tighten, while rental prices remain at their pre-GFC boom value. Short-term sales contract sharply and fall below their pre-boom level, while rental shares increase. Deviations are relative to the stochastic steady state at $t = 0$. *pp* denotes percentage points. See Section 7 and Appendix A for details.

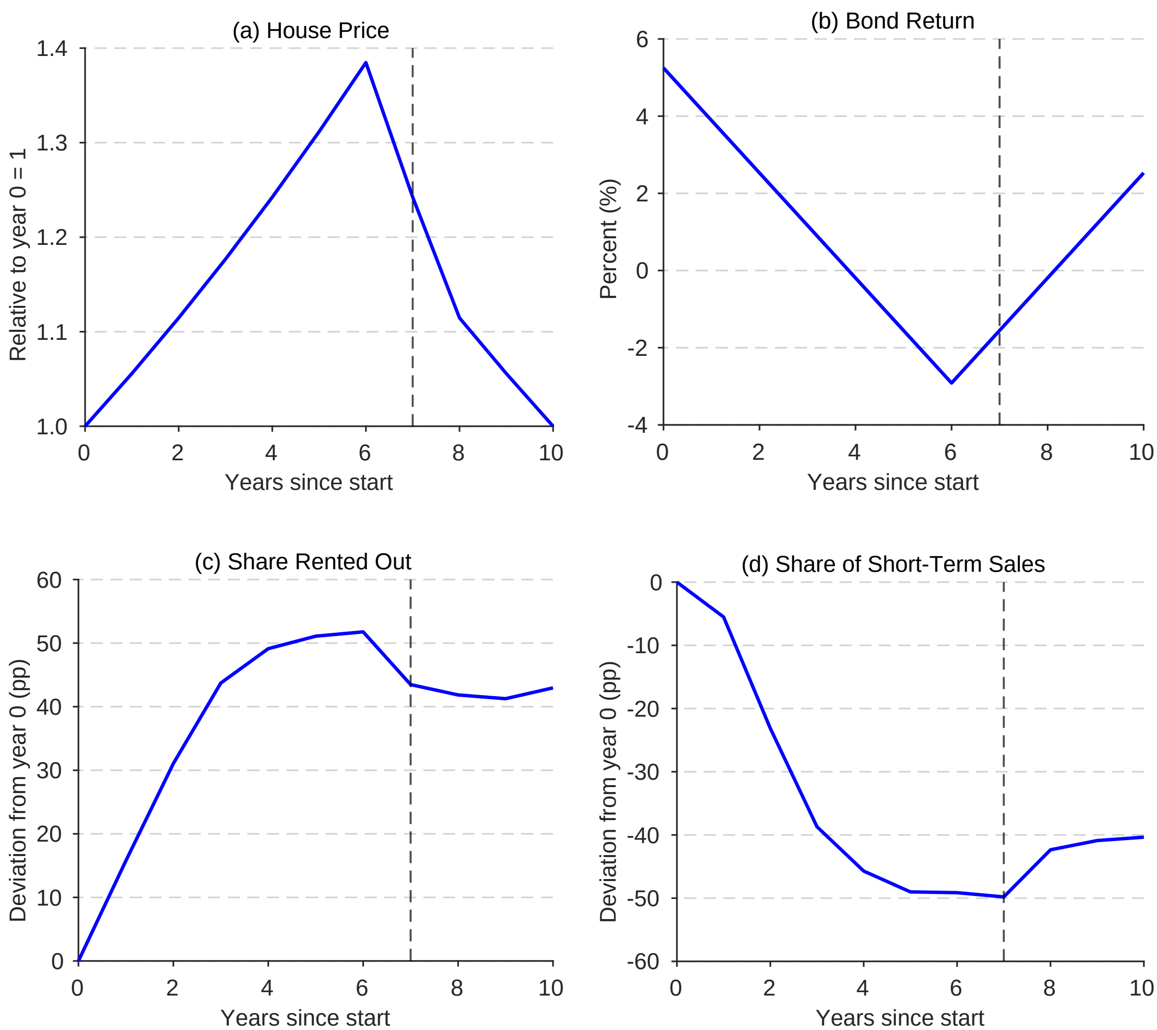

Figure 10: Boom-Bust Episode and Investor Portfolio Dynamics: Bond Environment

Note: Panels (a)-(b) show house prices and real bond returns. Panels (c)-(d) show the share of housing units rented out and short-term sales. The boom phase from $t = 1$ to $t = 6$ is identical to that in the post-GFC boom in Figures 6 and 8. At $t = 7$, house prices begin to decline and bond returns increase, while rental prices remain at their peak post-GFC boom value. Rental shares decline modestly and remain elevated, while short-term sales rise only slightly. Deviations are relative to the stochastic steady state at $t = 0$. *pp* denotes percentage points. See Section 7 and Appendix A for details.

Tables

Table 1: Mortgage Financing and Leverage Across Housing Booms

Pre-GFC				Post-GFC			
Year	Has mortgage (%)	LTV ratio (%)	Credit availability index	Year	Has mortgage (%)	LTV ratio (%)	Credit availability index
2001	64.5	84.9	12.4	2012	19.7	81.0	6.7
2002	64.1	84.2	12.1	2013	26.5	80.7	5.4
2003	64.4	83.3	12.5	2014	23.0	81.1	5.2
2004	65.4	81.7	14.9	2015	22.7	82.1	5.8
Average	64.7	83.3	13.0	Average	23.1	81.2	5.7

Note: This table reports, by year, the share of investor purchases financed with a mortgage and the average LTV conditional on having a mortgage. It also reports the Urban Institute's Housing Credit Availability Index (HCAI). Differences across booms are -41.6 percentage points for mortgage usage and -2.1 percentage points for average LTV ($p < 0.01$ for both), based on t-tests of equality of means.

Table 2: Calibration: Mortgage Environment

Parameter	Value	Interpretation	Endogenous
		Preferences, housing, and income	
β	0.952	Discount factor	Yes
σ	2	CRRA coefficient	No
δ	1.5%	Housing depreciation rate	No
ζ	0.050	Housing adjustment cost parameter	Yes
$h_{\min}$	0.5	Minimum housing purchase size	No
ν	0.033	Mortgage amortization rate	No
y	1	Exogenous income	No
		Housing unit heterogeneity	
μ_ω	-0.062	Rental income heterogeneity (location)	Yes
s_ω	0.392	Rental income heterogeneity (scale)	Yes
μ_κ	0.005	Sales income heterogeneity (location)	Yes
s_κ	0.036	Sales income heterogeneity (scale)	Yes
		Macro-financial shocks	
ρ_p	0.85	Log house price persistence	No
σ_p	0.035	Log house price standard deviation	No
$\bar{p}$	0	Log house price mean	No
$\Pi_{pp'}$	See Section 6	House price transition probabilities	No
ρ_q	0.90	Log rent persistence	No
σ_q	0.02	Log rent standard deviation	No
$\bar{q}$	-2.76	Log rent mean	No
$\Pi_{qq'}$	See Section 6	Rent transition probabilities	No
θ_n	80%	LTV limit: normal credit	No
θ_l	100%	LTV limit: loose credit	No
r_n^*	6%	Mortgage interest rate: normal credit	No
r_l^*	4%	Mortgage interest rate: loose credit	No
$\Pi_{nn} = \Pi_{ll}$	0.9	Credit regime persistence probability	No

Note: Endogenous parameters are jointly chosen to match moments of the model's ergodic distribution: the ratio of annual housing purchases to income, housing value-to-mortgage ratio, rental share, vacant-sale share, rental yield, and capital gains. Non-endogenous parameters follow literature values or AR(1) estimates for house prices and rents. See Section 6 and Appendix A for details.

Table 3: Calibration: Bond Environment

Parameter	Value	Interpretation	Endogenous
		Preferences, housing, and income	
β	0.990	Discount factor	Yes
σ	2	CRRA coefficient	No
δ	1.5%	Housing depreciation rate	No
ζ	0.052	Housing adjustment cost parameter	Yes
$h_{\min}$	0.5	Minimum housing purchase size	No
y	1	Exogenous income	No
		Housing unit heterogeneity	
μ_ω	-0.677	Rental income heterogeneity (location)	Yes
s_ω	0.064	Rental income heterogeneity (scale)	Yes
μ_κ	0.033	Sales income heterogeneity (location)	Yes
s_κ	0.012	Sales income heterogeneity (scale)	Yes
		Macro-financial shocks	
ρ_p	0.85	Log house price persistence	No
σ_p	0.035	Log house price standard deviation	No
$\bar{p}$	0	Log house price mean	No
$\Pi_{pp'}$	See Section 6	House price transition probabilities	No
ρ_q	0.90	Log rent persistence	No
σ_q	0.02	Log rent standard deviation	No
$\bar{q}$	-2.76	Log rent mean	No
$\Pi_{qq'}$	See Section 6	Rent transition probabilities	No
ρ_r	0.80	Bond return persistence	No
σ_r	0.01	Bond return standard deviation	No
$\bar{r}$	1.17%	Bond return mean	No
$\Pi_{rr'}$	See Section 6	Bond return transition probabilities	No

Note: Endogenous parameters are jointly chosen to match moments of the model's ergodic distribution: the ratio of bond holdings to income, housing portfolio share, rental share, vacant-sale share, rental yield, and capital gains. Non-endogenous parameters follow literature values or AR(1) estimates for house prices, rents, and bond returns. See Section 6 and Appendix A for details.

Table 4: Summary Statistics: Empirical Short-Term Sales Analysis

	Pre-GFC (2001-2004)				**Post-GFC** (2012-2015)			
	Mean	SD	P10	P90	Mean	SD	P10	P90
Sale within two years	0.17	0.38	0.00	1.00	0.16	0.37	0.00	1.00
Sale within three years	0.25	0.43	0.00	1.00	0.20	0.40	0.00	1.00
Price growth (%)	0.87	0.70	0.16	1.78	0.73	0.79	-0.11	1.73
Rental yield (%)	7.59	1.70	5.87	9.77	9.43	1.82	7.08	11.7
Sales price ($ thousands)	204	161	49.0	410	212	243	33.0	465
House age (years)	30.5	29.2	0.00	76.0	42.5	30.1	7.00	86.0
Log house size	9.09	0.94	8.07	10.3	9.09	0.96	8.16	10.3
Number of rooms	3.01	3.37	0.00	8.00	2.72	3.23	0.00	7.00
Local buyer	0.63	0.48	0.00	1.00	0.70	0.46	0.00	1.00
Foreign buyer	0.001	0.03	0.00	0.00	0.01	0.09	0.00	0.00
Legal entity	0.34	0.47	0.00	1.00	0.55	0.50	0.00	1.00
Population growth (%)	1.74	1.83	-0.33	4.33	1.04	0.88	-0.06	2.11
Unemp. rate change (pp)	0.24	0.74	-0.70	1.30	-1.05	0.53	-1.70	-0.40
Income growth (%)	5.20	3.74	0.83	9.77	4.31	3.16	0.06	7.76
Observations	1,557,752				1,387,427			

Note: This table reports summary statistics for the main variables used in the empirical analysis at the property level. The sample consists of single-family home purchases by retail investors matched to local price indices, rent measures, property characteristics, and economic controls. Statistics are reported separately for purchases made during 2001-2004 and 2012-2015, tracking each property until resale or the end of 2017. Price growth is measured at the property's zip-code level and rental yield at the MSA level. Population growth, unemployment rate change, and per capita income growth are annual county-level percentages. See Appendix B for detailed variable definitions.

Table 5: Logit Estimates: Short-Term Sales and Capital Gains

	Two-year horizon		**Three-year horizon**	
Price growth × Post-GFC	-0.052***	-0.095***	-0.035***	-0.072**
	(0.020)	(0.030)	(0.009)	(0.030)
Price growth	0.056***	0.097***	0.040**	0.081***
	(0.019)	(0.029)	(0.017)	(0.028)
Property characteristics	Yes	Yes	Yes	Yes
Investor type	Yes	Yes	Yes	Yes
Demand factors	Yes	Yes	Yes	Yes
MSA fixed effects	Yes	No	Yes	No
Year fixed effects	Yes	No	Yes	No
MSA × Year fixed effects	No	Yes	No	Yes
Observations	1,283,070	1,282,649	1,283,085	1,282,757

Note: This table reports logit coefficient estimates for equation (20). The main regressors are local price growth and its interaction with the post-GFC indicator. Robust standard errors clustered at the zip-code level are in parentheses. Controls include purchase price on the deed, house age, log house size, and number of rooms; investor-type indicators (local, foreign, legal entity); and county-level population growth, per capita income growth, and the change in the unemployment rate in the year prior to potential sale, as well as month-of-purchase fixed effects. Specifications include either MSA and year fixed effects or MSA×year fixed effects. *** $p < 0.01$, ** $p < 0.05$.

Table 6: Logit Estimates: Short-Term Sales and Rental Yields

	Two-year horizon		**Three-year horizon**	
Rental yield × Post-GFC	-0.084***	-0.108***	-0.070***	-0.084***
	(0.028)	(0.035)	(0.024)	(0.031)
Rental yield	-0.055**	-0.072***	-0.043*	-0.065***
	(0.025)	(0.016)	(0.022)	(0.015)
Property characteristics	Yes	Yes	Yes	Yes
Investor type	Yes	Yes	Yes	Yes
Demand factors	Yes	Yes	Yes	Yes
State fixed effects	Yes	No	Yes	No
Year fixed effects	Yes	No	Yes	No
State × Year fixed effects	No	Yes	No	Yes
Observations	885,977	885,893	885,977	885,977

Note: This table reports logit coefficient estimates for equation (21). The main regressors are rental yield and its interaction with the post-GFC indicator. Robust standard errors clustered at the MSA level are in parentheses. Controls include purchase price on the deed, house age, log house size, and number of rooms; investor-type indicators (local, foreign, legal entity); and county-level population growth, per capita income growth, and the change in the unemployment rate in the year prior to potential sale, as well as month-of-purchase fixed effects. Specifications include either state and year fixed effects or state×year fixed effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 7: Housing Boom Experiments: Mortgage Environment Results

	Stochastic Steady State		**Ergodic Distribution**	
Variable	Pre-GFC	Post-GFC	Pre-GFC	Post-GFC
Macro-financial fundamentals (exogenous)				
Change in house price (%)	38.5	38.5	38.5	38.5
Change in rental price (%)	0	25.2	0	25.2
Price-to-rent ratio	20.8	16.6	20.8	16.6
LTV limit at origination (%)	100	80	100	80
Mortgage interest rate (%)	4	6	4	6
Portfolio composition and turnover (endogenous)				
Change in leverage (pp)	20.2	-44.5	23.3	-42.0
Change in housing sales share (pp)	22.1	5.42	21.8	4.35
Change in rental share (pp)	-18.4	-4.63	-18.8	-4.52
Change in short-term sales share (pp)	19.5	-0.82	14.1	-5.84
Mean holding period (years)	3.45	8.42	3.93	9.41
Investment activity and performance (endogenous)				
Housing holdings (index)	2.53	0.64	2.89	0.67
New purchases (index)	7.75	0.25	8.38	0.32
Mortgage balances (index)	3.80	0.29	4.55	0.32
New originations (index)	10.2	0	11.2	0
Realized housing IRR, levered (%)	28.3	10.8	27.3	12.1
Realized housing IRR, unlevered (%)	8.08	10.8	8.04	12.1

Note: Results are grouped based on whether initial conditions ($t = 0$) are computed at the stochastic steady state or sampled from the ergodic distribution. Reported changes reflect differences between peak or trough values during the boom and initial values at $t = 0$, except for house and rental prices, which are reported as percentage changes. Quantities reported as "index" are normalized so that their value at $t = 0$ equals 1. The housing IRR corresponds to an investment initiated at $t = 1$, accounting for endogenous rental and sale proceeds, mortgage financing, scheduled amortization, interest payments, and prepayments upon sale. Any remaining stock is sold at $t = 6$. *pp* denotes percentage points. See Section 7 for discussion and Appendix A for simulation details.

Table 8: Housing Boom Experiments: Bond Environment Results

	Stochastic Steady State		**Ergodic Distribution**	
Variable	Pre-GFC	Post-GFC	Pre-GFC	Post-GFC
Macro-financial fundamentals (exogenous)				
Change in house price (%)	38.5	38.5	38.5	38.5
Change in rental price (%)	0	25.2	0	25.2
Price-to-rent ratio	20.8	16.6	20.8	16.6
Bond return (%)	5.25	-2.92	5.25	-2.92
Portfolio composition and turnover (endogenous)				
Change in housing portfolio share (pp)	-1.97	64.6	-23.1	53.3
Change in housing sales share (pp)	34.7	-54.2	45.2	-40.7
Change in rental share (pp)	-32.6	51.1	-22.9	56.9
Change in short-term sales share (pp)	40.9	-49.1	54.0	-33.5
Mean holding period (years)	1.05	2.25	1.05	2.68
Investment activity and performance (endogenous)				
Housing holdings (index)	0.61	7.58	0.17	2.25
New purchases (index)	1	5.82	1	6.23
Bond holdings (index)	1.05	0.28	1.46	0.25
Realized housing IRR (%)	6.11	5.98	7.14	6.42
Realized bond IRR (%)	5.25	1.15	5.25	1.15

Note: Results are grouped based on whether initial conditions ($t = 0$) are computed at the stochastic steady state or sampled from the ergodic distribution. Reported changes reflect differences between peak or trough values during the boom and their initial values at $t = 0$, except for house and rental prices, which are reported as percentage changes. Quantities reported as "index" are normalized so that their value at $t = 0$ equals 1. The housing IRR corresponds to an investment initiated at $t = 1$, accounting for endogenous rental and sale proceeds, with any remaining stock sold at $t = 6$. The bond IRR corresponds to an investment initiated at $t = 1$, with one-period bond returns reinvested annually until $t = 6$. *pp* denotes percentage points. See Section 7 and Appendix A for details.

Internet Appendix

Sections A and B provide additional details on the model and the empirical analysis, respectively. The appendix also includes supplementary figures and tables.

A Model Appendix

In this section, we derive the optimality conditions, describe the numerical solution method, and present the calibration details used in the quantitative analysis.

A.1 Derivation of Optimality Conditions

We derive the optimality conditions associated with the recursive problem defined in Section 3, which characterize the policy rules and form the basis for the numerical solution.

The first-order condition for consumption implies $u'(c_t) = \lambda_t$, where λ_t is the multiplier on the period-t budget constraint. The associated stochastic discount factor is $\Lambda_{t+1} = \beta \frac{u'(c_{t+1})}{u'(c_t)}$, which we use to express all intertemporal optimality conditions below.

A.1.1 Housing Purchases

The first-order condition for housing purchases, H_t^*, is

$$\frac{\beta}{\lambda_t}\mathbb{E}_t\left[\frac{\partial V(\mathcal{Z}_t, \mathcal{S}_{t+1})}{\partial H_t}\right] - p_t - \psi_{H^*}(H_t^*, H_{t-1}) + \frac{\tilde{\mu}_t^{\mathrm{LTV}}}{\lambda_t}\theta_t p_t + \frac{\tilde{\mu}_t^{H^*}}{\lambda_t} = 0, \tag{A1}$$

where $\tilde{\mu}_t^{\mathrm{LTV}}$ and $\tilde{\mu}_t^{H^*}$ are the Lagrange multipliers on the LTV limit and the minimum purchase constraint, respectively. Define the adjusted multipliers $\mu_t^{\mathrm{LTV}} = \frac{\tilde{\mu}_t^{\mathrm{LTV}}}{\lambda_t}$ and $\mu_t^{H^*} = \frac{\tilde{\mu}_t^{H^*}}{\lambda_t}$.

The envelope condition with respect to previous housing holdings, H_{t-1}, is

$$\frac{\partial V(\mathcal{Z}_{t-1}, \mathcal{S}_t)}{\partial H_{t-1}} = \lambda_t \left\{ -\delta p_t - \psi_H(H_t^*, H_{t-1}) + \phi_{t-1} \left(1 + \frac{S_\omega(\phi_{t-1})}{\phi_{t-1}}\right) q_t \right. \\ \left. + \rho_t \left(1 + \frac{S_\kappa(\tilde{\rho}_t)}{\tilde{\rho}_t}\right) p_t + \frac{\beta}{\lambda_t} \mathbb{E}_t \left[\frac{\partial V(\mathcal{Z}_t, \mathcal{S}_{t+1})}{\partial H_t}\right] (1 - \rho_t) \right\}. \tag{A2}$$

Using the acquisition cost $\mathcal{C}_t$ in (16), combining (A1) and (A2), iterating the envelope condition one period forward, and substituting back into (A1) yields

$$\mathcal{C}_t - \mu_t^{H^*} = \mathbb{E}_t \left[\Lambda_{t+1} \left\{ -\delta p_{t+1} - \psi_H(H_{t+1}^*, H_t) + \phi_t \left(1 + \frac{S_\omega(\phi_t)}{\phi_t}\right) q_{t+1} \right. \right. \\ \left. \left. + \rho_{t+1} \left(1 + \frac{S_\kappa(\tilde{\rho}_{t+1})}{\tilde{\rho}_{t+1}}\right) p_{t+1} + (1 - \rho_{t+1}) \left(\mathcal{C}_{t+1} - \mu_{t+1}^{H^*}\right) \right\} \right]. \tag{A3}$$

This condition equates the marginal acquisition cost to the expected discounted payoff from housing. The right-hand side collects the rental return, the resale value of vacant units, and the continuation value of retained housing, net of depreciation and adjustment costs.

Define P_{t+1}^H as the gross marginal payoff from holding one unit of housing from t to $t+1$, i.e., the term inside the expectation in (A3). With this notation, $\mathcal{C}_t - \mu_t^{H^*} = \mathbb{E}_t \left[\Lambda_{t+1} \left\{P_{t+1}^H\right\}\right]$.

A.1.2 New Mortgages

The first-order condition for new mortgages, M_t^*, is

$$\frac{\beta}{\lambda_t} \mathbb{E}_t \left[\frac{\partial V(\mathcal{Z}_t, \mathcal{S}_{t+1})}{\partial M_t}\right] + \frac{\beta}{\lambda_t} \mathbb{E}_t \left[\frac{\partial V(\mathcal{Z}_t, \mathcal{S}_{t+1})}{\partial X_t}\right] r_t^* + 1 - \frac{\tilde{\mu}_t^{\text{LTV}}}{\lambda_t} + \frac{\tilde{\mu}_t^{M^*}}{\lambda_t} = 0, \tag{A4}$$

where $\tilde{\mu}_t^{M^*}$ is the Lagrange multiplier on the non-negativity constraint for new mortgages.

Define the adjusted multiplier $\mu_t^{M^*} = \frac{\tilde{\mu}_t^{M^*}}{\lambda_t}$. Also define $\Omega_t^M = -\frac{\beta}{\lambda_t} \mathbb{E}_t \left[\frac{\partial V(\mathcal{Z}_t, \mathcal{S}_{t+1})}{\partial M_t}\right]$ and $\Omega_t^X = -\frac{\beta}{\lambda_t} \mathbb{E}_t \left[\frac{\partial V(\mathcal{Z}_t, \mathcal{S}_{t+1})}{\partial X_t}\right]$ as the adjusted marginal continuation costs of principal and interest payments, respectively. Then (A4) can be written as

$$\Omega_t^M + \Omega_t^X r_t^* + \ \mu_t^{\text{LTV}} - \mu_t^{M^*} = 1. \tag{A5}$$

This condition equates the expected discounted cost of servicing principal and interest to the marginal benefit of borrowing (one unit of resources), adjusted for binding constraints.

The envelope condition with respect to previous mortgage balances, M_{t-1}, is

$$\frac{\partial V(\mathcal{Z}_{t-1}, \mathcal{S}_t)}{\partial M_{t-1}} = \lambda_t \left\{ -\rho_t(1-\nu) - \nu + \frac{\beta}{\lambda_t}\mathbb{E}_t \left[\frac{\partial V(\mathcal{Z}_t, \mathcal{S}_{t+1})}{\partial M_t} \right] (1-\rho_t)(1-\nu) \right\}. \tag{A6}$$

Iterating (A6) one period forward, multiplying both sides by $-\frac{\beta}{\lambda_t}$, substituting the definition of Ω_{t+1}^M, and taking expectations conditional on time t, yields

$$\Omega_t^M = \mathbb{E}_t \left[\Lambda_{t+1} \left\{ \nu + \rho_{t+1}(1-\nu) + (1-\rho_{t+1})(1-\nu)\Omega_{t+1}^M \right\} \right]. \tag{A7}$$

This recursion captures the marginal continuation cost of mortgage principal: regular amortization, prepayment upon sale, and continuation if the mortgage remains outstanding.

The envelope condition with respect to previous interest obligations, X_{t-1}, is

$$\frac{\partial V(\mathcal{Z}_{t-1}, \mathcal{S}_t)}{\partial X_{t-1}} = \lambda_t \left\{ -1 + \frac{\beta}{\lambda_t}\mathbb{E}_t \left[\frac{\partial V(\mathcal{Z}_t, \mathcal{S}_{t+1})}{\partial X_t} \right] (1-\rho_t)(1-\nu) \right\}. \tag{A8}$$

The analogous recursion for the marginal continuation cost of interest obligations is

$$\Omega_t^X = \mathbb{E}_t \left[\Lambda_{t+1} \left\{ 1 + (1-\rho_{t+1})(1-\nu)\Omega_{t+1}^X \right\} \right]. \tag{A9}$$

This term reflects the marginal continuation cost of servicing interest payments, including the next period payment and continuation for properties that are not sold.

A.1.3 Selling Choice for Vacant Units

The first-order condition for the share of vacant houses to sell, $\tilde{\rho}_t$, is

$$-\frac{\beta}{\lambda_t}\mathbb{E}_t\left[\frac{\partial V(\mathcal{Z}_t,\mathcal{S}_{t+1})}{\partial H_t}\right]H_{t-1}-\frac{\beta}{\lambda_t}\mathbb{E}_t\left[\frac{\partial V(\mathcal{Z}_t,\mathcal{S}_{t+1})}{\partial M_t}\right](1-\nu)M_{t-1}$$
$$-\frac{\beta}{\lambda_t}\mathbb{E}_t\left[\frac{\partial V(\mathcal{Z}_t,\mathcal{S}_{t+1})}{\partial X_t}\right](1-\nu)X_{t-1}-(1-\nu)M_{t-1}+\left(1+\frac{\partial S_\kappa(\tilde{\rho}_t)}{\partial\tilde{\rho}_t}\right)p_tH_{t-1}=0. \quad \text{(A10)}$$

There is no need to include Lagrange multipliers on corner constraints, since the cumulative distribution Γ_κ has full support on the real line.

Using $\frac{\partial S_\kappa(\tilde{\rho}_t)}{\partial\tilde{\rho}_t}=\bar{\kappa}_t$ and defining the average mortgage rate on outstanding debt as $\bar{r}_{t-1}=\frac{X_{t-1}}{M_{t-1}}$, we substitute the definitions of Ω_t^M and Ω_t^X and apply (A1). This yields the cutoff

$$\bar{\kappa}_t=\frac{1}{p_t}\psi_{H^*}(H_t^*,H_{t-1})-\mu_t^{\mathrm{LTV}}\theta_t-\frac{\mu_t^{H^*}}{p_t}+\frac{(1-\nu)M_{t-1}}{p_tH_{t-1}}\left(1-\Omega_t^M-\Omega_t^X\bar{r}_{t-1}\right). \quad \text{(A11)}$$

Since $\tilde{\rho}_t=1-\Gamma_\kappa(\bar{\kappa}_t)$, a lower cutoff increases resale activity. The cutoff declines with higher house prices, looser credit conditions, lower mortgage prepayment burdens relative to portfolio value, and lower continuation values of the existing mortgage.

A.1.4 Vacancy-Rental Choice

The first-order condition for the share of houses to rent out, ϕ_t, is

$$\frac{\beta}{\lambda_t}\mathbb{E}_t\left[\frac{\partial V(\mathcal{Z}_t,\mathcal{S}_{t+1})}{\partial\phi_t}\right]=0. \quad \text{(A12)}$$

As before, there is no need to include Lagrange multipliers for corner constraints, since the cumulative distribution Γ_ω has full support on the real line.

The envelope condition with respect to the previous period's rental decision, ϕ_{t-1}, is

$$\begin{aligned}\frac{\partial V(\mathcal{Z}_{t-1}, \mathcal{S}_t)}{\partial \phi_{t-1}} = \lambda_t \Bigg\{ & \tilde{\rho}_t(1-\nu)M_{t-1} + \left(1 + \frac{\partial S_\omega(\phi_{t-1})}{\partial \phi_{t-1}}\right) q_t H_{t-1} - \tilde{\rho}_t \left(1 + \frac{S_\kappa(\tilde{\rho}_t)}{\tilde{\rho}_t}\right) p_t H_{t-1} \\ & + \frac{\beta}{\lambda_t}\mathbb{E}_t\left[\frac{\partial V(\mathcal{Z}_t, \mathcal{S}_{t+1})}{\partial H_t}\right]\tilde{\rho}_t H_{t-1} + \frac{\beta}{\lambda_t}\mathbb{E}_t\left[\frac{\partial V(\mathcal{Z}_t, \mathcal{S}_{t+1})}{\partial M_t}\right]\tilde{\rho}_t(1-\nu)M_{t-1} \\ & + \frac{\beta}{\lambda_t}\mathbb{E}_t\left[\frac{\partial V(\mathcal{Z}_t, \mathcal{S}_{t+1})}{\partial X_t}\right]\tilde{\rho}_t(1-\nu)X_{t-1}\Bigg\}. \end{aligned} \tag{A13}$$

Using $\frac{\partial S_\omega(\phi_t)}{\partial \phi_t} = \bar{\omega}_t$ and (A10) to eliminate continuation-value terms, we obtain

$$\frac{\partial V(\mathcal{Z}_{t-1}, \mathcal{S}_t)}{\partial \phi_{t-1}} = \lambda_t \left\{(1 + \bar{\omega}_{t-1})\, q_t + (\tilde{\rho}_t \bar{\kappa}_t - S_\kappa(\tilde{\rho}_t))\, p_t\right\} H_{t-1}. \tag{A14}$$

Iterating the envelope condition one period forward and substituting into (A12) yields

$$\bar{\omega}_t = \frac{\mathbb{E}_t\left[\Lambda_{t+1}\left\{(S_\kappa(\tilde{\rho}_{t+1}) - \tilde{\rho}_{t+1}\bar{\kappa}_{t+1})\,\frac{p_{t+1}}{p_t}\right\}\right]}{\mathbb{E}_t\left[\Lambda_{t+1}\left\{\frac{q_{t+1}}{p_t}\right\}\right]} - 1. \tag{A15}$$

Since $\phi_t = 1 - \Gamma_\omega(\bar{\omega}_t)$, a lower cutoff increases the rental share. The numerator captures the expected capital gains foregone by renting rather than preserving the option to sell next period, expressed as a return, while the denominator reflects the expected rental return.

A.1.5 Bond Purchases

The first-order condition for bond purchases, B_t, is

$$\frac{\beta}{\lambda_t}\mathbb{E}_t\left[\frac{\partial V(\mathcal{Z}_t, \mathcal{S}_{t+1})}{\partial B_t}\right] - \frac{1}{1+r_t} + \frac{\tilde{\mu}_t^B}{\lambda_t} = 0, \tag{A16}$$

where $\tilde{\mu}_t^B$ is the Lagrange multiplier on the non-negativity constraint for bond purchases. Define the adjusted multiplier $\mu_t^B = \frac{\tilde{\mu}_t^B}{\lambda_t}$. Using the standard envelope condition $\frac{\partial V(\mathcal{Z}_{t-1}, \mathcal{S}_t)}{\partial B_{t-1}} =$

λ_t, iterating one period forward, and substituting into (A16), we obtain

$$\frac{1}{1+r_t} - \mu_t^B = \mathbb{E}_t\left[\Lambda_{t+1}\right]. \tag{A17}$$

This condition states that the investor chooses bond purchases such that the expected discounted return equals the purchase cost, unless the non-shorting constraint is binding.

A.2 Distribution of Holding Periods

Let $\Phi_t(s = k)$ denote the share of the housing portfolio at the end of period t consisting of units with holding period exactly equal to k. Holding periods are measured at the end of period t. New purchases H_t^* therefore enter with holding period one, while units that remain in the portfolio $(1-\rho_t)H_{t-1}$ age by one period. The law of motion for the distribution is:

$$\Phi_t(s=1) = \frac{H_t^*}{H_t}, \tag{A18}$$

$$\Phi_t(s=k) = \frac{(1-\rho_t)H_{t-1}}{H_t}\,\Phi_{t-1}(s=k-1), \quad k \geq 2, \tag{A19}$$

where total housing holdings H_t evolve according to (2). The cumulative short-term share defined in (12) satisfies $\Phi_t(s \leq T) = \sum_{k=1}^{T} \Phi_t(s=k)$.

Conditional on being eligible for sale at the beginning of period t, the probability of sale is independent of holding duration. As a result, changes in the distribution arise from endogenous portfolio dynamics through purchases and survival.

A.3 Numerical Solution

We solve the model using a global projection method, adapting the computational approach of Elenev, Landvoigt, and Van Nieuwerburgh (2021). The solution is obtained by approximating the policy, transition, and forecasting functions over the relevant region of the state space. Global nonlinear methods are well suited to our environment, which features cutoff rules governing resale and rental decisions, as well as occasionally binding borrowing, minimum purchase, and non-negativity constraints, and other sources of strong nonlinearity.

A.3.1 State Space and Function Approximation

Let $s = (z, x)$ denote the state vector, where z collects endogenous states and x exogenous shocks, as defined in Section 3. The endogenous state space is discretized using bounded tensor-product grids with strictly increasing nodes in each dimension. The exogenous state follows a finite-state Markov chain constructed from discretized AR(1) processes described in Section 6. The full discretized state space consists of all combinations $\{s_j\}_{j=1}^{N_s} = \{z_k\}_{k=1}^{N_z} \times \{x_i\}_{i=1}^{N_x}$, where $N_s = N_z N_x$. Thus, each s_j corresponds to a unique pair (z_k, x_i).

The solution method approximates three sets of functions defined over the discretized state space. The first set consists of policy functions, which determine the investor's optimal choices at each state s. These include consumption $c(s)$, housing purchases $H^*(s)$, new mortgage borrowing $M^*(s)$, the rental share $\phi(s)$, the vacant-sale share $\tilde{\rho}(s)$, bond holdings $B(s)$, the Lagrange multiplier on the LTV constraint $\mu^{\text{LTV}}(s)$, and the multipliers associated with the minimum purchase and non-negativity constraints $\mu^{H^*}(s)$, $\mu^{M^*}(s)$, and $\mu^B(s)$.

The second set consists of transition functions, which describe the law of motion of endogenous state variables as a function of the current state s and the next-period realization of exogenous shocks x'. These functions determine the next-period housing holdings, rental shares, mortgage balances, interest payments, and bond positions.[26]

The third set consists of forecasting functions, which map the current state into the variables required to evaluate conditional expectations in the optimality conditions. These func-

[26] In our formulation, next-period endogenous states do not depend directly on next-period shocks, unlike in cash-on-hand or net worth formulations. We nevertheless retain the dependence on next-period shocks in the notation for generality and notational convenience in describing the system of nonlinear equations below.

tions partially overlap with the policy functions but also include additional variables. In particular, they include consumption $c(s)$, the selling cutoff $\bar{\kappa}(s)$, the LTV multiplier $\mu^{\text{LTV}}(s)$, the multiplier on the minimum purchase constraint $\mu^{H^*}(s)$, marginal continuation values of mortgage debt and interest payments $\Omega^M(s)$ and $\Omega^X(s)$, and the value function $V(s)$.

A.3.2 Solution Algorithm

The solution procedure consists of three steps.

Step 1. Initialization. Policy, transition, and forecasting functions are initialized at their non-stochastic steady-state values, which serve as the starting point for the iteration.

Step 2. Iteration. The following steps are performed at each iteration:

1. For each grid point s_j, evaluate the transition functions for every possible realization of next-period exogenous shocks x_i. This yields the implied next-period endogenous states $(H_{ij}, \phi_{ij}, M_{ij}, X_{ij}, B_{ij})$ associated with the current state s_j and shock realization x_i.

2. Next, evaluate the forecasting functions at these future state realizations. This yields the next-period forecasting variables $(c_{ij}, \bar{\kappa}_{ij}, \mu_{ij}^{\text{LTV}}, \mu_{ij}^{H^*}, \Omega_{ij}^M, \Omega_{ij}^X, H_{ij}^*, V_{ij})$ associated with the current state s_j and next-period shock realization x_i. These variables are sufficient to evaluate the conditional expectations entering the optimality conditions.

3. At each grid point s_j, solve the nonlinear system of optimality conditions, which con-

sists of eight equations in eight policy variables, given by

$$\hat{\mathcal{C}}_j - \hat{\mu}_j^{H^*} = \sum_i \pi_{ij} \left[\hat{\Lambda}_{ij} \left\{ -\delta p_i - \psi_H(H_{ij}^*, H_{ij}) + \hat{\phi}_j \left(1 + \frac{S_\omega\left(\hat{\phi}_j\right)}{\hat{\phi}_j} \right) q_i \right. \right.$$

$$\left. \left. + \hat{\rho}_{ij} \left(1 + \frac{S_\kappa\left(\hat{\bar{\rho}}_{ij}\right)}{\hat{\bar{\rho}}_{ij}} \right) p_i + \left(1 - \hat{\rho}_{ij}\right) \left(\mathcal{C}_{ij} - \mu_{ij}^{H^*} \right) \right\} \right], \tag{A20}$$

$$\hat{\Omega}_j^M + \hat{\Omega}_j^X r_j^* + \hat{\mu}_j^{\text{LTV}} - \hat{\mu}_j^{M^*} = 1, \tag{A21}$$

$$\frac{1}{1 + r_j} - \hat{\mu}_j^B = \sum_i \pi_{ij} \left[\hat{\Lambda}_{ij} \right], \tag{A22}$$

$$\hat{\mu}_j^{\text{LTV}} \left(\theta_j p_j \hat{H}_j^* - \hat{M}_j^* \right) = 0, \tag{A23}$$

$$\hat{\mu}_j^{H^*} \left(\hat{H}_j^* - h_{\min} \right) = 0, \tag{A24}$$

$$\hat{\mu}_j^{M^*} \hat{M}_j^* = 0, \tag{A25}$$

$$\hat{\mu}_j^B \hat{B}_j = 0, \tag{A26}$$

$$\hat{c}_j + \delta p_j H_j + p_j \hat{H}_j^* + \psi(\hat{H}_j^*, H_j) + \hat{\rho}_j (1 - \nu) M_j + \nu M_j + X_j + \frac{B_j}{1 + r_j}$$

$$= y_j + \left[\phi_j \left(1 + \frac{S_\omega\left(\phi_j\right)}{\phi_j} \right) q_j + \left(1 - \phi_j\right) \hat{\bar{\rho}}_j \left(1 + \frac{S_\kappa\left(\hat{\bar{\rho}}_j\right)}{\hat{\bar{\rho}}_j} \right) p_j \right] H_j + \hat{M}_j^* + \hat{B}_j. \tag{A27}$$

Equations (A20) and (A21) correspond to the Euler equation for housing purchases in (A3) and the optimality condition for new mortgages in (A5), respectively, while (A22) corresponds to the bond Euler equation in (A17). Equations (A23)-(A26) are the complementary slackness conditions associated with the LTV constraint (5), the minimum purchase constraint, and the non-negativity constraints on mortgage borrowing and bond holdings. Finally, equation (A27) corresponds to the budget constraint in (11).

Expectations are computed as weighted sums, where the weights are the transition probabilities π_{ij} to the exogenous state x_i conditional on the current state s_j. The hat notation ($\hat{\cdot}$) denotes the unknown variables solved for by the nonlinear equation solver, which finds the solution to (A20)-(A27) at each point s_j. All variables inside the expectation terms with subscripts (ij) are functions of the forecasting variables.

The forecasting variables are treated as fixed numbers when solving the nonlinear system, since they are computed in the previous sub-step of the iteration. For instance, the stochastic discount factor satisfies $\hat{\Lambda}_{ij} = \beta \frac{u'(c_{ij})}{u'(\hat{c}_j)}$, which depends on both the forecasted next-period consumption c_{ij} and the unknown $\hat{c}_j$ determined at state s_j.

The system of equations (A20)-(A27) relies on the following auxiliary relationships:

$$\hat{\rho}_j = (1-\phi_j)\hat{\tilde{\rho}}_j, \tag{A28}$$

$$\hat{\rho}_{ij} = (1-\hat{\phi}_j)\tilde{\rho}_{ij}, \tag{A29}$$

$$\hat{\mathcal{C}}_j = \left(1-\hat{\mu}_j^{\text{LTV}}\theta_j\right) p_j + \psi_{H^*}(\hat{H}_j^*, H_j), \tag{A30}$$

$$\mathcal{C}_{ij} = \left(1-\mu_{ij}^{\text{LTV}}\theta_i\right) p_i + \psi_{H^*}(H_{ij}^*, H_{ij}), \tag{A31}$$

$$\hat{\Omega}_j^M = \sum_i \pi_{ij}\left[\hat{\Lambda}_{ij}\left\{\nu + \hat{\rho}_{ij}(1-\nu) + (1-\hat{\rho}_{ij})(1-\nu)\Omega_{ij}^M\right\}\right], \tag{A32}$$

$$\hat{\Omega}_j^X = \sum_i \pi_{ij}\left[\hat{\Lambda}_{ij}\left\{1 + (1-\hat{\rho}_{ij})(1-\nu)\Omega_{ij}^X\right\}\right], \tag{A33}$$

$$\hat{\bar{\kappa}}_j = \frac{1}{p_j}\psi_{H^*}(\hat{H}_j^*, H_j) - \hat{\mu}_j^{\text{LTV}}\theta_j - \frac{\hat{\mu}_j^{H^*}}{p_j} + \frac{(1-\nu)M_j}{p_j H_j}\left(1-\hat{\Omega}_j^M - \hat{\Omega}_j^X \bar{r}_j\right), \tag{A34}$$

$$\hat{\bar{\omega}}_j = \frac{\sum_i \pi_{ij}\left[\hat{\Lambda}_{ij}\left\{\left(S_\kappa(\tilde{\rho}_{ij}) - \tilde{\rho}_{ij}\bar{\kappa}_{ij}\right)\frac{p_i}{p_j}\right\}\right]}{\sum_i \pi_{ij}\left[\hat{\Lambda}_{ij}\left\{\frac{q_i}{p_j}\right\}\right]} - 1. \tag{A35}$$

Equations (A28) and (A29) define the current and forecasted shares of houses sold (relative to total holdings). The former depends on the current vacant-sale share $\hat{\tilde{\rho}}_j$, while the latter depends on the current rental share $\hat{\phi}_j$. Equations (A30) and (A31) are the current and forecasted marginal acquisition costs, corresponding to (16). Equations (A32) and (A33) are the marginal continuation values associated with mortgage balance and interest payments (A7) and (A9). Finally, equations (A34) and (A35) are the optimal cutoffs for the selling and vacancy-rental choices (A11) and (A15).

4. Given the solution for the policy variables at state s_j, the policy functions are updated directly. The forecasting functions are also updated directly. The value function is updated according to $\hat{V}_j = u(\hat{c}_j) + \sum_i \pi_{ij}\beta V_{ij}$. The transition functions for the endogenous

state variables are then updated using the model's laws of motion: for each current state s_j and next-period exogenous state x_i, we compute

$$H_{ij} = \hat{H}^*_j + (1 - \hat{\rho}_j)H_j, \tag{A36}$$

$$\phi_{ij} = \hat{\phi}_j, \tag{A37}$$

$$M_{ij} = \hat{M}^*_j + (1 - \hat{\rho}_j)(1 - \nu)M_j, \tag{A38}$$

$$X_{ij} = \hat{M}^*_j r^*_j + (1 - \hat{\rho}_j)(1 - \nu)X_j, \tag{A39}$$

$$B_{ij} = \hat{B}_j. \tag{A40}$$

Equations (A36), (A38), and (A39) correspond to the laws of motion for housing, mortgages, and interest payments in (2), (6), and (7). Equations (A37) and (A40) reflect that the rental share and bond holdings are fully determined by current policy choices.

5. Check convergence using the sup norm of changes in transition and forecasting functions between successive iterates. If the distance falls below a predetermined tolerance, the algorithm terminates and the current functions are taken as the approximate solution. Otherwise, policy functions are updated directly, while forecasting and transition functions are updated using a convex combination of previous and newly computed values to improve numerical stability. The algorithm then returns to Step 2.

We use the previous-iteration policy functions as initial guesses for the nonlinear solver at each grid point. If the solver fails to converge at specific points, we reinitialize the system using the solution from the nearest successfully converged point and resolve the equations.

The nonlinear system is solved using GSL's multidimensional root-finding algorithm. To ensure strict positivity of consumption, we solve for its logarithm and recover consumption via exponentiation within the solver. Because nonlinear solvers may struggle to handle complementary slackness conditions associated with occasionally binding constraints, we implement the Kuhn-Tucker conditions as additive equations, following Judd, Kubler, and

Schmedders (2002). The nonlinear systems are solved independently at each grid point, and the procedure is parallelized across grid points within each iteration.

Step 3. Simulation. After convergence, we simulate the model forward using the approximated policy, transition, and forecasting functions. In the stochastic simulations, exogenous shocks evolve according to the discretized Markov chain and endogenous states are updated recursively using the model's laws of motion. We use a burn-in of 500 periods followed by 20,000 simulation periods and fix the random seed for reproducibility. When shocks are held fixed, the simulation converges to a stochastic steady state associated with that exogenous configuration. When shocks evolve according to the Markov chain, the long-run simulated sample approximates the ergodic distribution of endogenous states.

Deterministic boom experiments are initialized from these simulated objects. Under stochastic-steady-state initialization, the deterministic path starts from the final endogenous state of a fixed-shock simulation. Under ergodic-distribution initialization, we first simulate the model stochastically, cluster the resulting endogenous states, and sample initial conditions using cluster weights implied by the simulated stationary distribution. Auxiliary states governing holding-period distributions are initialized consistently with the same procedure. We then impose the deterministic paths for house prices, rents, and financing conditions described in Section 7 and simulate the model forward. For the ergodic-distribution initialization, the reported boom results correspond to the mean across simulated paths.

A.3.3 Computation of Realized Internal Rates of Return

To compute the housing IRR, we construct a sequence of cash flows $\{CF_t\}_{t=1}^{6}$ for an investment initiated at $t = 1$ and define r_H^{IRR} as the value that solves $\sum_{t=1}^{6} \frac{CF_t}{(1+r_H^{\mathrm{IRR}})^{t-1}} = 0$.

The initial cash flow at $t = 1$ is $CF_1 = -p_1 H_1^* - \psi(H_1^*, H_0) + M_1^*$, where H_1^* and M_1^* are the housing purchase and mortgage choices implied by the model solution, while H_0 denotes the initial housing position from which the boom experiment starts.

For this particular investment, the associated housing stock, mortgage balance, and interest payments evolve according to $H_t = (1-\rho_t)H_{t-1}$, $M_t = (1-\rho_t)(1-\nu)M_{t-1}$, and $X_t = (1-\rho_t)(1-\nu)X_{t-1}$ for $t = 2, \ldots, 6$, with initial conditions $H_1 = H_1^*$, $M_1 = M_1^*$, and $X_1 = M_1^* r_1^*$. The selling share $\rho_t = (1-\phi_{t-1})\tilde{\rho}_t$ is implied by the model solution.

For $t = 2, \ldots, 6$, cash flows incorporate the rental and selling surpluses:

$$\begin{aligned} CF_t = & -\delta p_t H_{t-1} - \rho_t(1-\nu)M_{t-1} - \nu M_{t-1} - X_{t-1} \\ & + \left[\phi_{t-1}\left(1 + \frac{S_\omega(\phi_{t-1})}{\phi_{t-1}}\right) q_t + (1-\phi_{t-1})\tilde{\rho}_t \left(1 + \frac{S_\kappa(\tilde{\rho}_t)}{\tilde{\rho}_t}\right) p_t\right] H_{t-1}. \end{aligned} \tag{A41}$$

The final cash flow CF_6 additionally includes the liquidation value of the remaining housing stock net of the remaining mortgage balance, $p_6 H_6 - M_6$.

We compute this IRR separately for each simulated path. Under ergodic-distribution initialization, the reported IRR corresponds to the mean across simulated paths. The unlevered counterfactual IRR is computed analogously, setting initial mortgage borrowing to zero.

The bond IRR is computed for an investment initiated at $t = 1$ and rolled over period by period until $t = 6$. Since bond holdings purchased at t for $\frac{B_t}{1+r_t}$ pay face value B_t at $t+1$, the bond IRR is the geometric mean of the gross returns: $r_B^{\text{IRR}} = \left[\prod_{t=1}^{5}(1+r_t)\right]^{1/5} - 1$.

A.4 Calibration Details

A.4.1 AR(1) Estimates for Real House Prices and Rents

To calibrate the stochastic processes for house prices (8) and rents (9), we estimate AR(1) processes for detrended real U.S. house price and rent series. House prices are measured using the S&P CoreLogic Case-Shiller U.S. National Home Price Index (FRED series CSUSHPISA, annual). Rents are measured using the CPI Rent of Primary Residence index (FRED series CUSR0000SEHA, annual). Both series are converted to real terms by deflating with the CPI

All Items Less Shelter (FRED series CUSR0000SA0L2). The estimation sample begins in 1987, reflecting the earliest availability of the national Case-Shiller index.

We take logs of the real series and remove low-frequency trends using the Hodrick-Prescott filter with smoothing parameter $\lambda = 100$, the standard value for annual data. The cyclical components of the detrended series are then used to estimate AR(1) processes by ordinary least squares. The estimates imply persistence parameters $\rho_p = 0.849$ and $\rho_q = 0.626$, and innovation standard deviations $\sigma_p = 0.040$ and $\sigma_q = 0.016$ for house prices and rents, respectively. The estimated correlation between the AR(1) innovations is small (0.056).[27]

To generate a peak price-rent ratio of 16.6 during the post-GFC boom in the model (see Figure 6), we set the persistence of the rent process to $\rho_q = 0.9$. Using the lower persistence estimated in the data would produce counterfactually high price-rent ratios in the model.

A.4.2 AR(1) Estimates for Real Bond Returns

To calibrate the bond return process (10), we estimate an AR(1) model for real one-year U.S. Treasury returns. Nominal one-year Treasury yields are taken from the one-year constant-maturity Treasury series (FRED series GS1, monthly). Inflation is measured using the CPI for All Urban Consumers (FRED series CPIAUCNS, monthly). To construct annual real returns, we use a December-to-December timing convention. Specifically, we combine the one-year Treasury yield observed in December of each year with CPI inflation measured between December of the previous year and December of the current year.

We estimate an AR(1) process for the resulting annual series by ordinary least squares using data for 1980-2018. The estimates imply persistence $\rho_r = 0.804$, an innovation standard deviation $\sigma_r = 0.014$, and an unconditional mean real return of approximately 1.16%. Using calendar-year averages instead to construct annual real returns yields similar estimates.

[27]Using the CPI Owners' Equivalent Rent index (FRED series CUSR0000SEHC) instead yields very similar persistence and volatility estimates, as well as a similarly small innovation correlation with house prices.

A.4.3 Rental and Resale Heterogeneity

We assume logistic distributions for rental and resale heterogeneity in valuations:

$$\Gamma_j(x) = \frac{1}{1 + \exp\left(-\frac{x-\mu_j}{s_j}\right)}, \qquad j \in \{\omega, \kappa\}, \tag{A42}$$

where x takes values on the entire real line, and μ_j and s_j denote the location and scale parameters, respectively. The location parameters shift the distributions, while the scale parameters determine their dispersion. This specification is computationally convenient and generates smooth interior cutoff decisions for resale and rental choices.

B Empirical Appendix

This section describes the construction of the macro-financial data and transaction-level datasets, and presents additional empirical results.

B.1 Aggregate Data for Macro Drivers

B.1.1 Real House Prices

For national house prices, we use the S&P CoreLogic Case-Shiller U.S. National Home Price Index (FRED series CSUSHPISA, monthly, seasonally adjusted). To convert nominal house prices into real terms, we deflate the index by the CPI All Items Less Shelter (FRED series CUSR0000SA0L2, monthly, seasonally adjusted). For each housing boom episode (2001-2007 and 2012-2018), we construct a boom-specific real house price series by dividing the nominal Case-Shiller index by CPI ex-shelter and aggregating the resulting series to the quarterly frequency using calendar-quarter averages. We then normalize each quarterly series to equal 100 in the first quarter of the corresponding boom, i.e., 2001Q1 or 2012Q1.

B.1.2 Price-Rent Ratio

We construct the national price-rent ratio as the ratio of the market value of residential real estate owned by households to the value of owner-occupied housing services. The market value of household real estate is taken from the Financial Accounts (FRED series BOGZ1FL155035013Q, quarterly, not seasonally adjusted). The corresponding rent measure (owner-occupied housing services) is available only at the annual frequency (FRED series A2013C1A027NBEA). To obtain a quarterly series for owner-occupied housing services, we apply the Chow and Lin (1971) method, using the log of annual owner-occupied housing services as the low-frequency target and the log of the product of total housing services

(FRED series DHUTRC1Q027SBEA, quarterly) and the national homeownership rate (FRED series RHORUSQ156N, quarterly) as the quarterly indicator. The resulting Chow-Lin quarterly series is then exponentiated to obtain the quarterly value of owner-occupied housing services. The price-rent ratio is constructed each quarter as the ratio of market value of household real estate to the imputed quarterly value of owner-occupied housing services.

B.1.3 Mortgage Credit Availability Index

Mortgage credit conditions are obtained from the Urban Institute's Housing Credit Availability Index (HCAI). The HCAI reports the percentage of owner-occupied purchase loans that are likely to default and serves as a proxy for the overall underwriting stance of mortgage lenders. We use the "Total Risk" index from the "Whole Market" panel of the HCAI dataset, which is available at the quarterly frequency. We use the index in levels (percent).

B.1.4 Real Five-Year Treasury Yield

Our primary measure of the real five-year U.S. Treasury yield is the inflation-indexed five-year constant-maturity Treasury yield (FRED series DFII5, monthly). When the TIPS-based real yield is unavailable, we construct a proxy real yield as the nominal five-year constant-maturity yield (FRED series DGS5, monthly) minus one-year expected inflation (FRED series EXPINF1YR, monthly). The resulting monthly real-yield series is then converted to the quarterly frequency by taking calendar-quarter averages. The real ten-year Treasury yield is constructed analogously, using FRED series DFII10 (inflation-indexed ten-year constant-maturity yield) and DGS10 (nominal ten-year constant-maturity yield).

B.2 Detailed Description of the Database

In this section, we describe the data sources, cleaning procedures, and key variables used in our empirical analysis.

B.2.1 Investors' Purchases

The transaction data come from the CoreLogic Deeds Dataset, which compiles county-recorded deeds documenting ownership transfers of single-family homes in the U.S. Each record includes the transfer date, property address, sale price, buyer and seller names, the buyer's mailing address, and the loan amount when a mortgage lien is filed. We restrict the sample to transactions recorded between January 1, 2000 and December 31, 2017.

We focus on single-family homes, which are the most commonly transacted residential units. In total, the raw deeds universe contains nearly 60 million such transactions by homeowners and investors over this period, before applying the filters described below.

We drop transactions with missing purchase price or with price below $10,000, a common practice with deeds data. We also drop all transactions in non-disclosure states, which do not require that the sale price be reported to the county office. Specifically, we remove all records from the 12 non-disclosure states: Alaska, Idaho, Kansas, Louisiana, Mississippi, Missouri, Montana, New Mexico, North Dakota, Texas, Utah, and Wyoming.

To identify legal entities, we search for buyer names that contain the terms "LLC," "LP," "Inc," "Trust," "Corporation," "Partners," or entity names such as "Invitation Homes." From this set, we filter out names of relocation companies, non-profit organizations, construction companies, national and regional authorities, banks, Ginnie Mae, Fannie Mae, Freddie Mac, other mortgage loan companies and credit unions, homeowner associations, and county, city, or municipal entities. To identify relocation companies, non-profit organizations, and construction companies, we use publicly available lists of the top firms in each category in the U.S. We also manually check the names of the 200 largest non-individual buyers in each state using online search engines to classify them into the correct categories, iterating this procedure several times to ensure accurate classification of the largest buyers.

To accurately classify the largest single-family institutional investors, which we exclude from the analysis, we collect the names of the 26 largest institutional investors in the single-family rental market from industry and news reports. For example, Amherst Capital's 2018 market commentary report[28] provides a comprehensive list of these institutions (often referred to as "Wall Street Landlords"), ranked by the number of homes they own. We then search for the names of these large investors and their subsidiaries in the deeds database to ensure they are classified as institutional investors. We use public SEC filings and other business websites to identify the names of their subsidiaries. This procedure allows us to compute the exact holdings of the top single-family institutional investors.

To classify non-institutional buyers as individual retail investors, we calculate the number of purchases associated with each individual name within an MSA in the given year and the year before. We define individual retail investors as those who purchase more than one property within the same MSA over this two-year window. Thus, individuals who purchase their primary residence and a vacation home in different MSAs are not counted as retail investors. We also classify buyers as retail investors if the purchaser is a legal entity (such as an LLC, LP, trust, or corporation) and does not belong to the set of the 26 largest institutional single-family rental operators identified in the preceding paragraph.[29]

We calculate the share of retail investors as the number of purchases made by individual investors or by small and medium-sized legal entities divided by the total number of purchases, including those made by retail investors, large institutional investors, and home-

[28]Amherst Capital Management LLC (2018), An Update on Institutional Single-Family Rental Activity: 2017/2018 U.S. Market Trends Support Long-Term Growth & Opportunity, Amherst Capital Market Commentary, April 2018. Available at: https://www.amherst.com/insights/category/research/.

[29]Large institutional investors account for a very small share of single-family rental activity in aggregate (Garriga, Gete, and Tsouderou, 2023). Including these entities does not affect our findings.

owners. The holding duration is measured as the time between the date the deed was signed for the purchase of the property and the date the deed was signed for its subsequent sale.

We merge the deeds data with property characteristics obtained from tax assessors. These characteristics include the year the housing unit was built, total number of rooms, number of bedrooms, the square footage of the house, and the year of the last renovation or update.

B.2.2 Demographic Variables

Because the deeds dataset does not include demographic information on buyers, we use the mailing address reported in each deed to proxy for investor attributes. We obtain data on median income, education levels, and owner-occupied house values at the census-tract level from the Federal Financial Institutions Examination Council (FFIEC) census and demographic files. We merge the census tract code of each buyer's mailing address, as reported in the deeds, with these FFIEC data. This allows us to link each investor to the demographic and housing market characteristics of the area in which they reside.

B.2.3 House Prices and Rents

We match the property addresses in the deeds with data from several sources to obtain the evolution of housing prices and rents in the areas where investors purchase properties. House prices at the zip-code level and monthly frequency come from the Zillow Home Value Index for single-family homes. This index is designed to capture the value of a typical home in each zip code (rather than only transacted homes) by using information from the full distribution of properties in the zip code. Housing rental data also come from Zillow. For rents, the smallest geographical unit available with coverage back to 2000 is the MSA. We collect housing prices and rents at the MSA level with monthly frequency and calculate the rental yield as 12 times the monthly rent divided by the house price.

B.2.4 Other Variables

We also match the property addresses in the deeds with local economic conditions. We collect the following county-year variables from 2000 to 2017: population from the U.S. Census Bureau, per capita income from the Bureau of Economic Analysis, and the unemployment rate from the U.S. Bureau of Labor Statistics. This allows us to link each property to the economic conditions of the county in which it is located, both at the time of purchase and at the time of a potential sale (for example, two years after purchase).

Table 4 summarizes the main variables used in the empirical analysis.

B.3 Additional Results

B.3.1 New Entrants After the GFC

We examine whether the investors active in the post-GFC boom were the same individuals who participated during the 2000s. Using name and mailing-address information from property deeds, we match buyers across the 2001-2004 and 2012-2015 purchase cohorts. We consider two matching criteria. First, we match investors by name and the MSA of their mailing address, assuming that repeated investors would appear under the same name and reside in the same MSA; only 1.8% of investors satisfy this criterion. Second, we match investors by their exact mailing address across periods, assuming that repeat investors would use the same address in deed records; only 2.4% of addresses match.[30]

These results indicate that the vast majority of post-GFC retail investors were new entrants to the single-family market, rather than repeat participants from the pre-GFC boom.

[30] Matching geocoded mailing addresses (latitude and longitude) yields similarly low overlap.

B.3.2 Additional Evidence on Investor Characteristics

We complement the wealth proxy used in the main text with additional socioeconomic measures based on investors' mailing-address census tracts. We proxy investor income using the ratio of median household income in the investor's census tract to the corresponding MSA median, and measure educational attainment using the share of residents with a bachelor's degree or higher and a master's degree or higher. Figure A1 shows that the distributions of wealth, income, and education proxies all shift to the right in the post-GFC period. Panels (a) and (b) indicate that investors increasingly reside in higher-wealth and higher-income areas, while panels (c) and (d) show that they are located in more educated neighborhoods.

These patterns are consistent with the evidence in the main text: post-GFC investors are drawn from more affluent and higher-socioeconomic-status areas.

B.3.3 Year-by-Year Estimates of Short-Term Sales Sensitivities

Tables A1 and A2 report year-by-year logit estimates based on equations (20) and (21). The same qualitative patterns observed in the pooled specifications also appear within each boom: sensitivities to capital gains are stronger in the pre-GFC period, while sensitivities to rental yields are more negative in the post-GFC period. The year-by-year estimates also suggest that these differences are not driven by any single purchase year.

Internet Appendix Figures

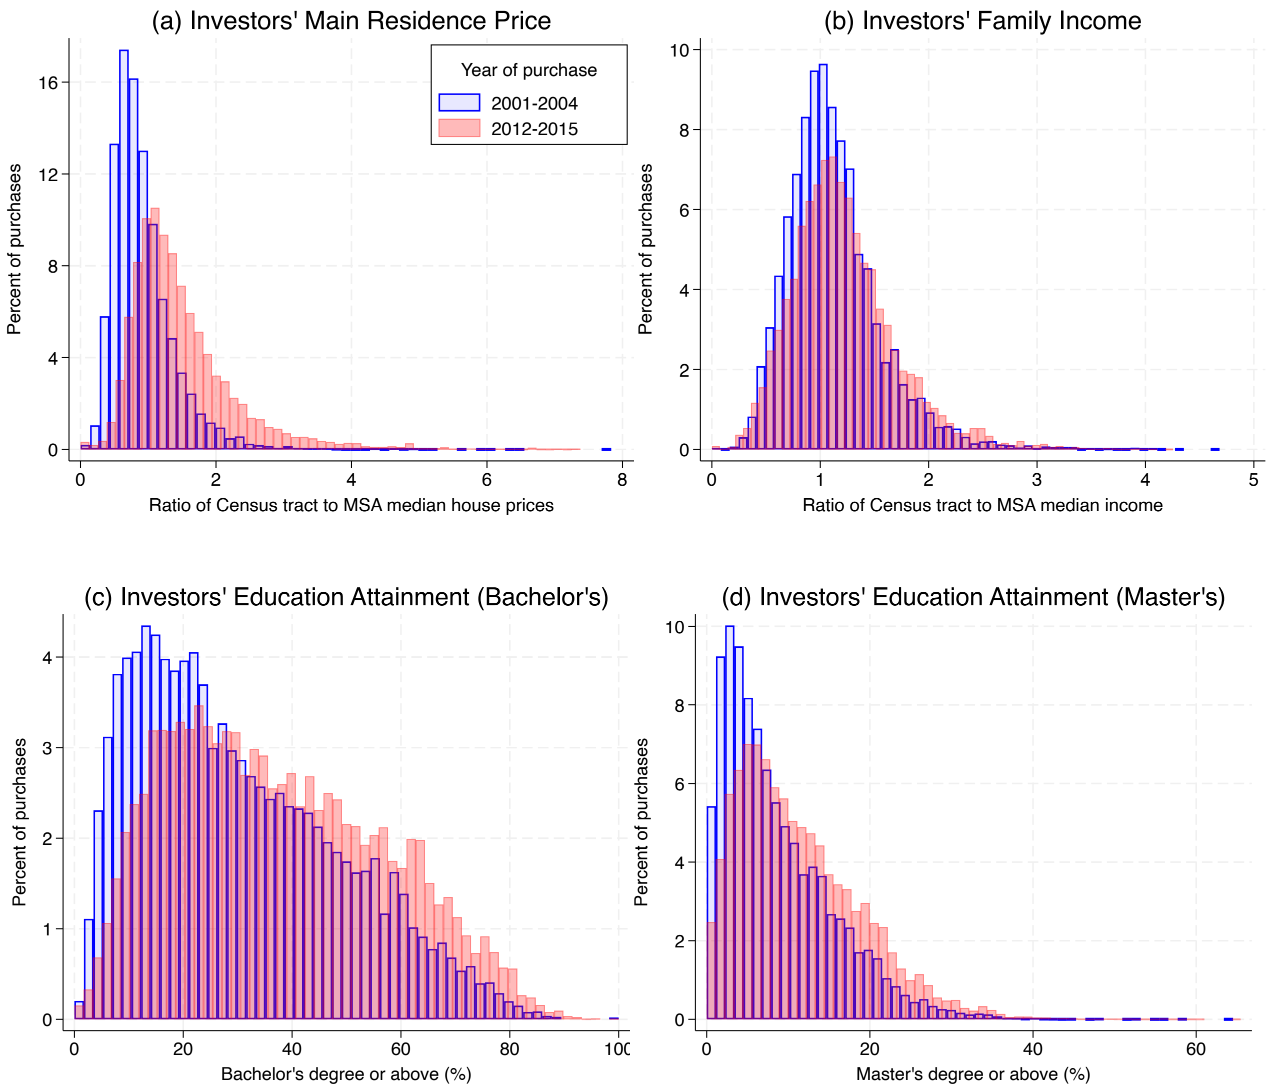

Figure A1: Distribution of Investor Characteristics

Note: The figure plots the distribution of investor characteristics across census tracts of investors' mailing addresses. Panel (a) shows the ratio of tract to MSA median house prices. Panel (b) shows the corresponding ratio for median household income. Panels (c) and (d) show the share of residents with a bachelor's and a master's degree or higher, respectively. Distributions are shown for investors purchasing in 2001-2004 and 2012-2015, corresponding to the pre- and post-GFC housing booms.

Internet Appendix Tables

Table A1: Logit Estimates: Short-Term Sales and Capital Gains by Year

	Pre-GFC				Post-GFC			
Year of purchase	2001	2002	2003	2004	2012	2013	2014	2015
	Two-year horizon				**Two-year horizon**			
Price growth	0.187***	0.179***	0.151***	0.174***	0.028**	-0.006	-0.024	-0.010
	(0.053)	(0.068)	(0.045)	(0.045)	(0.013)	(0.014)	(0.015)	(0.017)
	Three-year horizon				**Three-year horizon**			
Price growth	0.147***	0.154**	0.128***	0.156***	0.035***	0.005	-0.018	0.002
	(0.052)	(0.066)	(0.041)	(0.043)	(0.012)	(0.013)	(0.014)	(0.016)
Property characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Investor type	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demand factors	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	203,335	235,618	271,040	340,031	233,494	237,328	214,220	220,864

Note: This table reports year-by-year logit estimates based on equation (20). Robust standard errors clustered by zip code are in parentheses. The top panel reports results for sales within two years of purchase, and the bottom panel reports results for sales within three years. Property controls include purchase price from the deed, house age, log house size, and number of rooms. Investor-type controls include indicators for local investors, foreign investors, and legal entities. Demand controls include population growth, income growth, and the change in the unemployment rate. All specifications also include month-of-purchase dummies and MSA fixed effects. *** $p < 0.01$, ** $p < 0.05$.

Table A2: Logit Estimates: Short-Term Sales and Rental Yields by Year

	Pre-GFC				Post-GFC			
Year of purchase	2001	2002	2003	2004	2012	2013	2014	2015
	Two-year horizon				**Two-year horizon**			
Rental yield	-0.083***	-0.089***	-0.081***	-0.071***	-0.154***	-0.165***	-0.149***	-0.161***
	(0.016)	(0.018)	(0.021)	(0.023)	(0.036)	(0.034)	(0.031)	(0.037)
	Three-year horizon				**Three-year horizon**			
Rental yield	-0.066***	-0.068***	-0.063***	-0.070***	-0.128***	-0.140***	-0.123***	-0.152***
	(0.014)	(0.015)	(0.020)	(0.022)	(0.032)	(0.030)	(0.028)	(0.034)
Property characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Investor type	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demand factors	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	203,335	235,618	271,040	340,031	233,494	237,328	214,220	220,864

Note: This table reports year-by-year logit estimates based on equation (21). Robust standard errors clustered by MSA are in parentheses. The top panel reports results for sales within two years of purchase, and the bottom panel reports results for sales within three years. Property controls include purchase price from the deed, house age, log house size, and number of rooms. Investor-type controls include indicators for local investors, foreign investors, and legal entities. Demand controls include population growth, income growth, and the change in the unemployment rate. All specifications also include month-of-purchase dummies and state fixed effects. *** $p < 0.01$.