The Mortgage Credit Channel of Macroeconomic Transmission*

Daniel L. Greenwald†

July 13, 2026

Abstract

This paper investigates how payment-to-income (PTI) constraints influence house prices and household leverage. I document that PTI limits were widely binding in the 2010s post-bust period, but virtually unenforced during the 2000s boom period. A structural model reveals that many borrowers optimally choose to be jointly constrained by PTI and loan-to-value (LTV) limits, yielding sharply different demand dynamics from single-constraint models. In general equilibrium, this amplifies shocks affecting PTI constraints and dampens shocks affecting LTV constraints. A relaxation of PTI limits matching the data explains between 33% and 53% of the rise in price-rent ratios over the boom.

*This paper is a revised version of Chapter 1 of my Ph.D. dissertation at NYU. I thank the editor Toni Whited and an anonymous referee for their comments and feedback, which have greatly improved the paper. I also thank my thesis advisors Sydney Ludvigson, Stijn Van Nieuwerburgh, and Gianluca Violante for their invaluable guidance and support. The paper benefited greatly from conversations with Andreas Fuster, Mark Gertler, Andy Haughwout, Malin Hu, Erik Hurst, Tim Landvoigt, Virgiliu Midrigan, Jonathan Parker, and Johannes Stroebel, among many others, insightful conference discussions by Monika Piazzesi, Amir Sufi, Paul Willen, and Hongjun Yan, and many helpful comments from seminar audiences. I thank NYU and the Becker-Friedman Institute for financial support.

†Leonard N. Stern School of Business Kaufman Management Center, 44 West Fourth Street, New York, NY 10012; Email: dlg340@stern.nyu.edu.

1 Introduction

Mortgage debt is central to the modern economy. The sharp rise in residential mortgage debt observed in many countries at the start of the century has been credited with fueling a dramatic boom and subsequent collapse in house prices, ultimately triggering a financial crisis and severe recessions worldwide. This episode motivates two natural questions: what role have changes in credit limits played in housing and mortgage cycles, and how might macroprudential policies have dampened them?

These questions center on the *mortgage credit channel* of macroeconomic transmission: the path from primitive shocks, through mortgage credit issuance, to the broader economy. Characterizing this channel is challenging due to the complexity of mortgage finance. In the US, as in many countries, borrowers must satisfy multiple credit limits at origination that depend on both individual and aggregate economic variables. These limits can change due to shifts in underwriting standards, or through movements in economic fundamentals holding credit standards fixed. Different limits bind for different portions of the population, with these shares varying endogenously over time. Although general equilibrium models have typically abstracted from many of these institutional details, I argue they are pivotal in the dynamics of house prices and household leverage.

In this paper, I focus on the role of payment-to-income ("PTI") limits that cap the ratio of a borrower's mortgage payment to their income. While a vast literature has studied the impact of loan-to-value ("LTV") constraints that limit mortgage balances relative to house value, the influence of PTI limits has been largely unstudied, despite their widespread role in underwriting both in the US and abroad. I find that although an overwhelming majority of borrowers are constrained by LTV limits, a substantial minority are jointly constrained by *both* LTV and PTI limits. The housing demand of these jointly constrained borrowers reacts very differently to shocks compared to borrowers constrained by a single limit only. In particular, the presence of joint constraints amplifies shocks affecting PTI limits while dampening shocks affecting LTV limits.

I begin with a simple numerical example, building on Justiniano, Primiceri, and Tambalotti (2019), to demonstrate how the two constraints can interact to drive changes in housing demand. The key insight is that once borrowers hit their PTI limits, they transition from financing marginal housing mostly with credit to financing it entirely with cash. This discrete change leads to a natural "corner" solution for housing demand at which both constraints bind simultaneously. The housing demand of these jointly constrained households increases strongly as PTI constraints are loosened (e.g., after a decrease in

interest rates), but falls when LTV constraints are loosened.

Since the strength of this mechanism depends on the share of jointly constrained borrowers, I present empirical evidence using loan-level microdata from Fannie Mae. I first show three stylized facts: (i) most borrowers appear constrained by LTV limits in all periods, (ii) the PTI distribution exhibits bunching toward a standard limit in the 2014–2016 post-bust period, and (iii) PTI limits appear virtually unenforced during the 2000–2008 housing boom period. These facts suggest that PTI limits are influential in periods with normal underwriting standards, and that their loosening was the primary change to credit standards during the boom.

Next, I formally estimate the share of borrowers limited by each constraint. For LTV limits, I exploit that LTV-constrained borrowers bunch at known thresholds. For PTI limits, I use a "missing mass" approach based on the apparent non-enforcement of PTI limits during the boom period. This yields a counterfactual PTI distribution in the absence of binding PTI limits, adjusted for current economic fundamentals. Applying these results to the 2014–2016 post-bust period reveals that a majority of purchase loan borrowers (78.1%) face a binding constraint, and that a substantial minority (23.5%) of these are constrained by PTI limits. Importantly, around three-quarters of these PTI-constrained borrowers are *also* constrained by LTV limits, consistent with the intuition from my simple example, and key for the theory to follow.

I further study whether areas with a higher pre-boom share of borrowers facing binding PTI constraints experienced larger increases in house prices during the boom period, as the theory would predict following a large relaxation of PTI limits. Since the Fannie Mae loan-level data do not cover the pre-boom era, I proxy for the influence of PTI constraints with the share of recently moved mortgagors that have high PTI ratios in the 1990 Census Public Use Microdata Sample (PUMS). A cross-sectional regression reveals a strong and robust relationship between the pre-boom high-PTI share and subsequent house price growth, explaining most of the geographic variation in house price growth over the 1998–2006 period.

Motivated by these findings, I develop a tractable general equilibrium framework in which borrowers must satisfy both LTV and PTI limits when originating long-term mortgages. The model allows borrowers to be constrained by LTV limits only, PTI limits only, or both, and incorporates heterogeneity that creates endogenous and time-varying shares in each constraint group. I then aggregate the model over these heterogeneous choices, calibrate it to match US mortgage data at the aggregate and loan levels, and verify that

the model implies realistic shares for each constraint type. In addition, I develop sufficient statistics summarizing the impact of changes in credit standards and interest rates on house prices and credit that can be evaluated using only information on the shares limited by each constraint, which prove highly accurate in my quantitative experiments.

With the calibrated model in hand, I first establish the effects of marginal changes to LTV and PTI limits. To isolate the effects caused by the constraint structure, I contrast the benchmark model against otherwise identical economies that impose an LTV limit or PTI limit only. The economy with both constraints is unique in generating a large rise in the price-rent ratio following a loosening of the PTI limit, leading to a strong rise in household leverage despite 87.7% of households being constrained by LTV at steady state. In contrast, while a relaxation of LTV limits would cause a large rise in price-rent ratios in an economy with an LTV limit only, the economy with both constraints sees no increase, due to a fall in housing demand among jointly constrained households.

Next, I consider movements in credit limits caused by changes in fundamentals rather than credit standards. I show that an increase in long-term inflation endogenously tightens the PTI limit, producing much stronger house price responses than in economies with either constraint alone. Conversely, housing preference shocks that endogenously relax LTV limits have a *smaller* effect on house prices under both constraints than in the alternative economies, showing that the effect of the constraint structure depends crucially on the specific shock applied.

To demonstrate the robustness of these findings, I extend the model to incorporate a life cycle. Although this extension matches the age profile of housing, and features substantial variation in the response to shocks across age groups, the aggregate responses are nearly identical to those of the baseline model. This near equivalence arises because the optimal housing demand of jointly-constrained households — the key driver of the model's dynamics — depends only on the structure of the LTV and PTI constraints, and not on other model blocks such as preferences, income expectations, or the life cycle. I show analytically that a wide set of possible extensions should have approximately identical responses to shocks affecting PTI limits, provided that the model is correctly calibrated to match the empirical shares limited by each constraint.

To close, I study the role of credit standards in the 2000s boom and bust, arguing that a relaxation of PTI limits was essential to the events that unfolded. I find that a PTI liberalization calibrated to match the loan-level evidence explains 33% of the observed rise in the price-rent ratio and 26% of the observed rise in the aggregate loan-to-income

ratio over this period. In contrast, even a maximal relaxation of LTV limits would fail to explain rising price-rent ratios, which would actually fall when holding PTI limits fixed.

Beyond its direct effects, this PTI liberalization also amplified other factors boosting house prices at this time. To show this, I incorporate shocks to long-term inflation and mortgage spreads to match the observed fall in interest rates, and preference shocks (proxying for overoptimism) to match the peak increase in price-rent ratios. Compared to this baseline, a counterfactual macroprudential policy enforcing PTI limits at their historical levels would have reduced the size of the boom by 53%. This reduction is much larger than the direct effect of PTI liberalization alone, reflecting that shocks to housing preferences have stronger effects under loose PTI limits. I conclude that a cap on PTI ratios, rather than LTV ratios, would have been more effective in limiting this boom-bust cycle.

Literature Review. This paper builds on several strands of the literature. On the empirical side, it relates to a large body of work documenting links between mortgage credit and house prices, and patterns of credit growth in the boom.[1] My study complements these works by highlighting the specific role of PTI constraints, and analyzing the theoretical mechanisms behind these links in general equilibrium. To my knowledge, this paper was the first to empirically document the abnormally loose PTI limits imposed during the housing boom — a relaxation I show is consistent with the broad-based and largely prime expansion of credit evidenced by Adelino, Schoar, and Severino (2016), Albanesi, DeGiorgi, and Nosal (2022), and Foote, Loewenstein, and Willen (2021). I also connect to Foote, Loewenstein, and Willen (2018) and Johnson and Tzur-Ilan (2025), who provide evidence that a shift toward automated underwriting was the likely driver of the PTI liberalization I document.

My estimates of the shares constrained by each limit connect to prior work that also measures PTI distributions relative to a counterfactual. DeFusco, Johnson, and Mondragon (2020) do so to estimate the effect of additional costs for high-PTI jumbo loans imposed by the Dodd-Frank legislation. Bosshardt, Di Maggio, Kakhbod, and Kermani (2024) similarly measure the marginal effects of monetary tightening on high-PTI borrowers. While these works estimate the marginal effects of policy changes on subsets of affected borrowers, I measure the *total* shares constrained by PTI and LTV limits, which my theoretical results show are key to the dynamics of house prices and mortgage credit.

Additional empirical work has estimated the effects of changes in PTI or similar loan-

[1] See, e.g., Adelino, Schoar, and Severino (2025); Aladangady (2014); Di Maggio and Kermani (2017); Favara and Imbs (2015); and Mian and Sufi (2014).

to-income (LTI) limits on house prices, often finding large effects that support my theoretical results.[2] Johnson (2020) exploits variation in PTI limits between Fannie Mae and Freddie Mac during the boom and finds strong positive effects of looser PTI limits on house price growth. Higgins (2024) studies a change in LTI limits (a close analogue to PTI limits) in Ireland, and finds strong effects on the marginal borrowing of low-income borrowers. Jung and Lee (2017) use cross-sectional variation in regional exposure to macroprudential policy in Korea to show that PTI tightening dampens house prices more than LTV tightening. I complement these works by showing that the pre-period share of high-PTI loans provides a powerful cross-sectional predictor of boom house price growth, explaining the majority of variation across US regions.

Turning to structural models, a large heterogeneous-agent housing literature has developed, often incorporating idiosyncratic risk, costly financial transactions, and long-term mortgages.[3] I complement these works by introducing PTI limits as a novel mechanism driving house prices, and incorporating specialized heterogeneity that lets borrowers endogenously sort into constraint types while preserving the tractability of an aggregated model.[4] This tractable analysis also connects to the literature on DSGE models with housing and collateralized debt, again incorporating PTI limits as a novel feature.[5]

This work is also related to research using structural models to study the role of credit standards in the 2000s boom-bust cycle.[6] My findings support the importance of credit liberalization in the boom, but emphasize that a relaxation of PTI constraints appears key. Of particular relevance is Justiniano, Primiceri, and Tambalotti (2019), who find that the interaction of an LTV constraint with an exogenous lending limit can generate strong effects of movements in the non-LTV constraint on debt and house prices, and potentially reverse the effect of changes in LTV limits — results echoed in many of the findings below. By using an endogenous PTI constraint in place of an exogenous fixed limit on lending, I am able to connect these dynamics to interest rate transmission, which directly affects PTI limits, and to link the housing boom to observed relaxations of PTI standards

[2]See e.g., Akinci and Olmstead-Rumsey (2018); Kuttner and Shim (2016); Jung and Lee (2017); Peydró, Rodriguez-Tous, Tripathy, and Uluc (2024).

[3]See, e.g., Boar, Gorea, and Midrigan (2022), Guler (2015), Beraja, Fuster, Hurst, and Vavra (2019), Campbell and Cocco (2015), Chatterjee and Eyigungor (2015), Landvoigt (2015), Wong (2015), Elenev, Landvoigt, and Van Nieuwerburgh (2016), Kaplan, Mitman, and Violante (2020).

[4]An earlier version of this paper, Greenwald (2018), showed that this framework can be easily inserted into standard macroeconomic DSGE models.

[5]See, e.g., Iacoviello (2005), Monacelli (2008), Iacoviello and Neri (2010), Ghent (2012), Liu, Wang, and Zha (2013), Rognlie, Shleifer, and Simsek (2018).

[6]See, e.g., Campbell and Hercowitz (2005), Kermani (2012), Iacoviello and Pavan (2013), Favilukis, Ludvigson, and Van Nieuwerburgh (2017).

in the data. My framework also features endogenous shares of borrowers limited by each constraint, rather than assuming the entire economy switches between constraint types, which I show is key to the quantitative effects of shocks on debt and house prices.

This paper is also closely related to Corbae and Quintin (2015), which to my knowledge represents the only prior macroeconomic model to incorporate a PTI constraint and use its relaxation as a proxy for the housing boom. These authors use a model with exogenous house prices to study the links between endogenously priced default risk and credit growth. While their setup delivers important findings regarding default and foreclosure, both absent from my model, these authors do not study the implications of how PTI and LTV constraints interact to drive house prices, the central topic of this paper.

Last, this paper connects to companion work extending this framework. Greenwald and Guren (2025) endogenize the rental market, showing that the degree of frictions between households and landlords are pivotal to the effect of credit on house prices, ultimately estimating strong frictions and large effects of credit on house prices. Allen and Greenwald (2022) similarly extend this framework to model multiple mortgage submarkets with different credit standards, showing that the effects of tightening macroprudential limits in these submarkets are sharply different. This paper provides the foundational structure for these companion works, and shows analytically that the simpler constraint structures used in these papers well approximate the full dynamics modeled here.

Overview. The remainder of the paper is organized as follows. Section 2 uses a simple example to provide intuition. Section 3 presents facts from loan-level microdata, and estimates the shares limited by each constraint. Section 4 constructs, solves, and calibrates the theoretical model. Section 5 presents the model dynamics with respect to changes in credit limits, while Section 6 applies them to study the 2000s boom episode. Section 7 concludes. Additional details and results can be found in the appendix.

2 Simple Numerical Example

To provide intuition for the model's core mechanisms, I present a simplified example from an individual borrower's perspective, building on Justiniano, Primiceri, and Tambalotti (2019). I describe the results below, and formalize the setting in Appendix C.

Consider a prospective home-buyer who prefers financing their purchase with credit rather than cash, perhaps because they must save for the down payment and delaying

purchase is costly. Assume this borrower's annual income is \$63k, and they face a 28% PTI limit, which limits their combined mortgage payment, taxes, and insurance, to \$1,470 per month.[7] Assuming an interest rate of 6% on a 30-year fixed-rate mortgage, and tax and insurance payments equal to 1.75% of the loan balance, this maximum payment is associated with a maximum loan size of \$200k consistent with the PTI limit.[8] The borrower's maximum LTV ratio is 80% so that, including the minimum 20% down payment, they reach this maximum loan size at a house price of \$250k.

This \$250k house price represents the threshold at which the borrower switches from being LTV-constrained to PTI-constrained. This creates a kink in the borrower's required down payment as a function of house price, shown as the solid blue line in Figure 1. Below this threshold price, the borrower is constrained by the value of their collateral, and a marginal dollar of housing can be financed by 80 cents of credit, requiring only 20 cents in additional down payment. But above this kink, they are constrained by their PTI limit. In this region they cannot borrow further regardless of collateral value, and fund additional housing entirely with cash. This discrete change around the kink implies that a "corner solution" price of exactly \$250k is a likely optimum for this borrower.

To formalize this intuition, since the borrower's discount rate is higher than the market interest rate, they perceive the future payments on \$1 of mortgage balance as costing only $\phi < 1$ in present value. As a result, their perceived marginal cost of buying an additional \$1 of housing jumps discretely from $0.2 + 0.8 \times \phi < 1$ to the left of the kink to \$1 to the right of it, as shown in the blue solid line in Figure 1(c). Since the optimal housing choice equalizes this marginal cost to the marginal benefit of additional housing, it will be optimal for the borrower to purchase exactly \$250k in housing for a wide range of preference parameters, with one such example used for the plots in Figure 1.

From this starting point, let us now assume that credit standards change. First, consider an increase in allowed PTI ratios from 28% to 31%. This increases the borrower's maximum monthly payment by nearly 11%, to \$1,628. The maximum loan size and the kink house price are similarly increased by nearly 11% to \$221k and \$277k, respectively. As Figure 1(c) shows, because the discrete jump in marginal cost shifts to the right, the borrower will now optimally choose the new kink house price point of \$277k. As a result, the borrower increases their housing purchase in proportion to PTI limits, potentially

[7] For simplicity, I round most quantities to the nearest \$1k = \$1,000. The borrower's exact income used in the example is \$63,017, designed to yield a PTI limit of exactly \$200k.

[8] This assumption that taxes and insurance are proportional to loan balance is made for simplicity. In reality, taxes and insurance payments are associated with the value of the property, not the size of the loan.

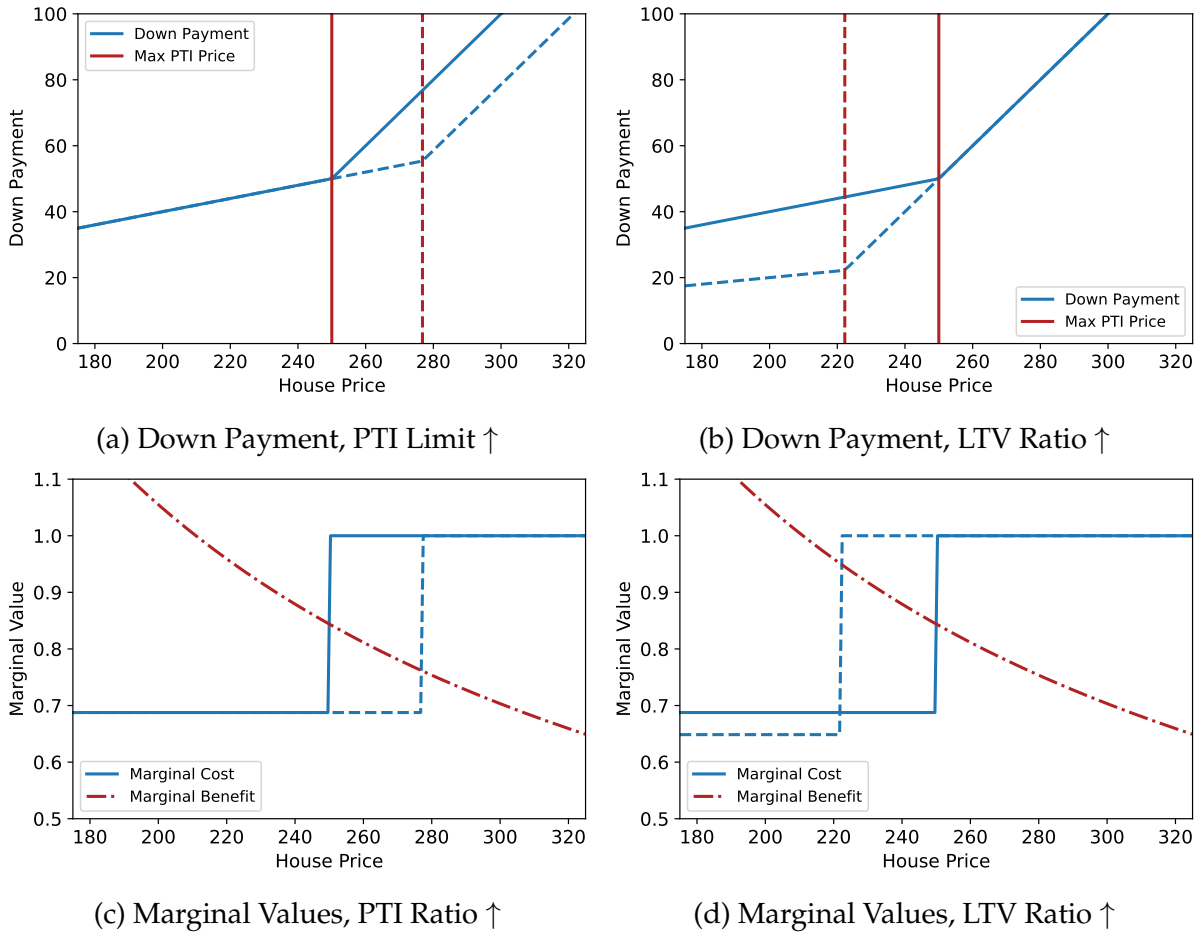

(a) Down Payment, PTI Limit ↑

(b) Down Payment, LTV Ratio ↑

(c) Marginal Values, PTI Ratio ↑

(d) Marginal Values, LTV Ratio ↑

Figure 1: Simple Numerical Example

Note: The figure displays a simple numerical example of a borrower's housing and mortgage choice. Panels (a) and (b) display the required down payment as a function of the house price under a baseline calibration and alternative relaxed credit standards. Panel (a) shows the effect of an increase in the maximum PTI ratio from 28% to 31%, while Panel (b) shows the effect of an increase in the maximum LTV ratio from 80% to 90%. Panels (c) and (d) display the corresponding perceived marginal cost and marginal benefit of additional housing under the baseline setting and relaxed credit standards. See Appendix C for details.

driving up house prices if many other borrowers do the same. This result, and the existence of the kink point more generally, depend crucially on the *interaction* of the LTV and PTI limits, and would not be present under either constraint in isolation.

In contrast, an increase in the maximum LTV ratio from 80% to 90%, shown in Figure 1(b), has a starkly different impact. In this case, the borrower's maximum loan size given their income is unchanged, at \$200k. But with only a 10% down payment, the house price associated with this loan falls to \$222k, an 11% *decrease*. Figure 1(d) shows that if the borrower once again follows their corner solution, the result is a fall in their housing demand, potentially contributing to a decline in house prices.

To understand this result, note that prior to the LTV loosening, moving from a $250k house to a $222k house would have let the borrower keep only $5.6k in cash, since they would have been forced to cut their loan size. But after the relaxation, the borrower can keep the entire $28k difference, increasing their cash savings from downsizing five-fold.[9] This result is not found in models in which borrowers face only an LTV constraint, where lower down payments typically increase housing demand, again highlighting the role of interactions between the two limits.[10]

These demand dynamics arise not only from changes in credit standards, but also from shifts in economic fundamentals. For instance, if the interest rate decreases from 6% to 5%, the borrower's original maximum payment of $1,470 will now be associated with a larger loan of $218k. In this case, the kink house price will again shift right to $272k, with nearly identical implications for demand as in the PTI relaxation example above. This demand response, with a semielasticity to interest rates exceeding 8, points to strong transmission from interest rates to house prices in this environment. More generally, shocks to economic fundamentals can interact with this channel in varying ways, depending on how they interact with the two constraints, as explored further in Section 4.

3 Empirical Analysis

This section documents the empirical properties of the LTV and PTI constraints. I present descriptive statistics from mortgage microdata, estimate the share of borrowers limited by each constraint, and study the relationship between pre-boom PTI ratios and boom house price growth.

3.1 Descriptive Statistics

I first describe basic patterns in loan-level data from the Fannie Mae Loan Performance dataset. Figure 2 displays distributions of combined LTV (total mortgage debt divided by house value — hereafter CLTV) and PTI ratios on newly originated conventional fixed-

[9] Alternatively, consider that a relaxation of the LTV limit increases the effective supply of collateral, since each unit of housing can collateralize more debt, but does not increase the demand for collateral, since the borrower's overall loan size is still constrained by their PTI limit. An increase in supply, with demand held fixed, pushes down the price of collateral, depressing the value of housing.

[10] With only an LTV constraint, the marginal cost curve would continue as a flat line without discretely jumping to $1. Since the value of this line would fall from $0.2 + 0.8\phi$ to the smaller value $0.1 + 0.9\phi$, this lower marginal cost would now intersect with a decreasing marginal benefit curve further to the right, boosting housing demand.

rate mortgages used for home purchases. The data are plotted for two samples: the boom period (2000–2008) and a post-bust period (2014–2016) following the stabilization of the housing market.[11] Additional histograms by LTV and PTI ratio, loan type, and credit score can be found in Appendix Figures B.12–B.15. The CLTV distributions in Panels (a) and (b) reveal two patterns of interest. First, most borrowers cluster in large spikes at known institutional limits and cost discontinuities.[12] Second, the CLTV distribution changes little between the two periods, and shows no sign of looser CLTV standards during the boom.

The PTI plots tell a markedly different story. Panel (d) shows that the post-bust distribution of PTI ratios does not feature discrete spikes, but instead builds up steadily toward the 45% institutional limit before being sharply truncated. This smooth shape likely reflects search frictions, as borrowers who prefer the kink-point house price described in Section 2 may be unable to find a home at this precise value. If borrowers are willing to buy a house at or below, but not above, the threshold price, a joint pattern of LTV spikes and a truncated PTI distribution will emerge naturally.[13] This pattern thus suggests that a meaningful share of borrowers were influenced by PTI limits in the post-bust period.

In complete contrast, the 2000–2008 boom data show no evidence of a PTI limit imposed at *any* level. Instead, the PTI histogram varies smoothly until 65%, at which point the data are truncated by the provider. By the boom's peak (2006:Q4), 55% of debt for home purchases went to loans violating the traditional PTI limit of 36%, while 19% of debt went to loans with PTI ratios exceeding 50%. These data suggest that PTI limits were unusually loose during the boom relative to historical norms and post-bust data, in contrast to the more stable distribution of CLTV ratios. Foote, Loewenstein, and Willen (2018) and Johnson and Tzur-Ilan (2025) attribute this relaxation to the introduction of automated underwriting systems by Fannie Mae and Freddie Mac in the mid-1990s, which effectively de-emphasized PTI limits in favor of alternative risk predictors.

Since these data cover the prime mortgage market (conforming loans to borrowers with typical credit scores), this evidence aligns with the broad-based expansion of credit documented by Adelino, Schoar, and Severino (2016), Albanesi, DeGiorgi, and Nosal (2022), and Foote, Loewenstein, and Willen (2021), who show that credit growth in the

[11] The boom period is chosen as the start of the data sample through the onset of the financial crisis in 2008. The post-bust period is chosen as a period with a stable housing and mortgage environment, and consistently enforced PTI limits. Appendix Figures B.16 and B.17 show that the bunching pattern displayed in Figure 2(d) begins to appear in 2009 and is fully in place from 2010:Q2 onward.

[12] The largest spikes occur at 80%, beyond which borrowers must pay for private mortgage insurance.

[13] Bank preapproval letters often cap the price at which a buyer can make an offer to exactly this threshold price by default, potentially explaining this asymmetry.

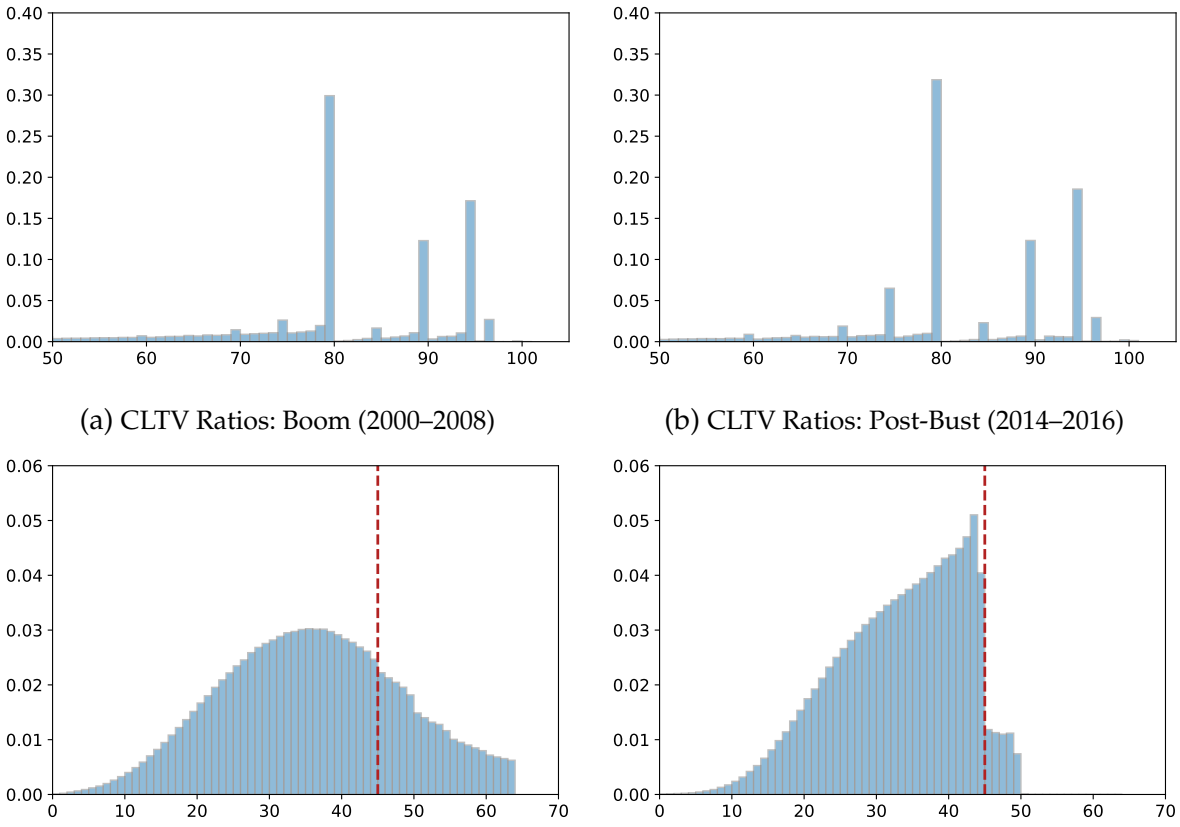

(a) CLTV Ratios: Boom (2000–2008)

(b) CLTV Ratios: Post-Bust (2014–2016)

(c) PTI Ratios: Boom (2000–2008)

(d) PTI Ratios: Post-Bust (2014–2016)

Figure 2: LTV and PTI Ratios, Fannie Mae Data

Note: The figure displays histograms of combined loan-to-value (CLTV) ratios and payment-to-income (PTI) ratios for newly originated purchase loans in the Fannie Mae Single-Family Loan Performance Data. The distributions are weighted by original unpaid principal balance. The sample is split into a "boom" period (2000–2008) and a "post-bust" period (2014–2016). The dashed vertical line in the PTI panels denotes the typical 45% PTI limit applied in the post-bust period.

boom was not concentrated among subprime or low-income borrowers. In Appendix B.4 I show that borrowers with high PTI ratios in the 2014–2016 period, who are more likely to be PTI-constrained, have somewhat lower average incomes than their low-PTI peers, but the two income distributions largely overlap. Moreover, high-PTI and low-PTI borrowers share similar house values, average local borrower income levels, and credit scores. As PTI ratios appear to bind across the income and credit score distributions, a relaxation of PTI limits is consistent with the broad-based expansion of credit found in the literature.

3.2 Estimating Shares at Each Constraint

Having documented the empirical patterns of LTV and PTI constraints, I now estimate the share of borrowers limited by each constraint in the post-bust period, as well as the share jointly limited by both constraints. This is non-trivial for the PTI limit, which clearly truncates the distribution of PTI ratios at 45%, but does not induce bunching at a single point. As a result, the mass to the left of the 45% cutoff contains a mixture of constrained borrowers who would have chosen higher PTI ratios if allowed, and unconstrained borrowers unaffected by the limit.

The ideal experiment would compare the actual distribution of PTI ratios to that of a counterfactual economy in which PTI constraints were not enforced. To proxy for this counterfactual, I exploit that PTI limits appear virtually unenforced during the boom, as shown in Figure 2(c). However, because PTI ratios are influenced by economic conditions such as house prices and interest rates, the boom distribution cannot be directly used as a counterfactual for the post-bust period. Instead, following Chernozhukov, Fernández-Val, and Melly (2013), I employ a distribution regression approach. By regressing the mass in different areas of the distribution on explanatory variables with high predictive power across locations and time, I can construct a fitted counterfactual distribution that adjusts for observed fundamentals in the post-bust sample.

To implement this, I estimate the three-digit ZIP code (ZIP3) by year regression

$$F_{i,t}^{PTI \leq j} = \alpha_i^j + \beta^j x_{i,t} + \varepsilon_{i,t}^j, \tag{1}$$

where $F_{i,t}^{PTI \leq j}$ is the share of borrowers in ZIP3 i at time t with PTI ratio at or below j, α_i^j is a ZIP3 fixed effect, and $x_{i,t}$ is a vector of two predictors. I estimate this regression over the 2000–2007 boom period, to capture the relationship between fundamentals and the PTI distribution when PTI limits are effectively unenforced. For $x_{i,t}$, I construct a first predictor $e_{i,t}^*$ to capture variation in geographic and macroeconomic fundamentals as

$$\log e_{i,t}^* = \log p_{i,t} + \log \hat{r}_t^* - \log y_{i,t}, \tag{2}$$

where $p_{i,t}$ is the FHFA house price index, $y_{i,t}$ is local average wages and salaries from the IRS Statistics of Income, and $\hat{r}_t^*$ is the effective payment rate on a new mortgage — the sum of principal and interest payments per dollar of debt on a 30-year fixed-rate loan at the national average interest rate, plus 175bp reflecting taxes and insurance.[14] Intuitively,

[14]This notation is inspired by equation (A.31), which derives a nearly identical expression as the cutoff

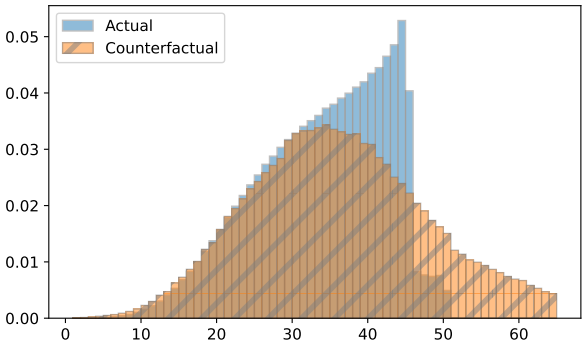

Figure 3: Actual vs. Counterfactual PTI Distributions, 2014–2016 Sample

Note: The figure compares the actual distribution of PTI ratios for purchase loans in the Fannie Mae data for the post-bust sample (2014–2016) against a counterfactual fitted distribution, computed using the distribution regression methodology described in Section 3.2.

as house prices and interest rates rise relative to incomes, more borrowers become constrained by PTI. For the second predictor, I use the share of borrowers with PTI ratios at or above 29%, denoted $F_{i,t}^{PTI \geq 29}$. This controls for shifts in the PTI distribution far from the limit, where the constraint should not be influential, to ensure that the counterfactual distribution is correctly anchored in this unconstrained region.[15]

Figure 3 compares the actual and counterfactual PTI distributions for the 2014–2016 sample using 1pp-width bins, with full implementation details in Appendix B.2. These sample dates follow the stabilization of PTI ratios after the housing recovery, but precede a loosening of PTI limits in 2017.[16] The actual distribution is close to the counterfactual for lower PTI ratios, but displays sizable excess mass just below, and missing mass just above, the 45% limit, consistent with a substantial fraction of PTI-constrained borrowers.

To quantify the constrained share, I re-estimate (1) using the share of loans with $PTI \geq 46$ as the dependent variable. Estimation results, displayed in Appendix Table B.2, show strong predictive power across a variety of specifications. Figure 4(a) displays the ag-

value of the idiosyncratic shock at which borrowers switch from LTV-constrained to PTI-constrained in the simplified model of Appendix A.5. The cutoffs in equations (19) and (20) in the full model below deliver similar functional forms.

[15]Since the effect of the PTI constraint should be concentrated in the proximity of the limit, the counterfactual and observed distributions should agree for PTI ratios well before the limit. In practice, while $\log e_{i,t}^*$ has substantial predictive power, Appendix Figure B.4 shows that a counterfactual distribution using $\log e_{i,t}^*$ as the sole predictor would overstate the share of borrowers with PTI ratios of 28 or less in the post-bust sample, and hence artificially understate the expected share in the high-PTI region, biasing the estimates. Adding $F_{i,t}^{PTI \geq 29}$ allows the counterfactual to effectively normalize the relative mass in the region far from the constraint, allowing a more accurate picture of distortions in the region closer to the constraint.

[16]Maximum PTI limits for many loans were expanded from 45 to 50 for loans underwritten by the automated Desktop Underwriter software in 2017.

gregated fitted values against the data, showing that the regression accurately fits the rise in this share during the boom. Importantly, these predicted values imply a decline in the $PTI \geq 46$ share in the post-bust period even absent binding constraints, making the raw change in shares inappropriate for measuring the fraction constrained. However, Figure 4(b) shows a large and significant deviation between the actual and counterfactual distributions, which stabilizes over the 2014–2016 period. Aggregating over this 2014–2016 period implies a total missing mass of 16.3% in the $PTI \geq 46$ region.

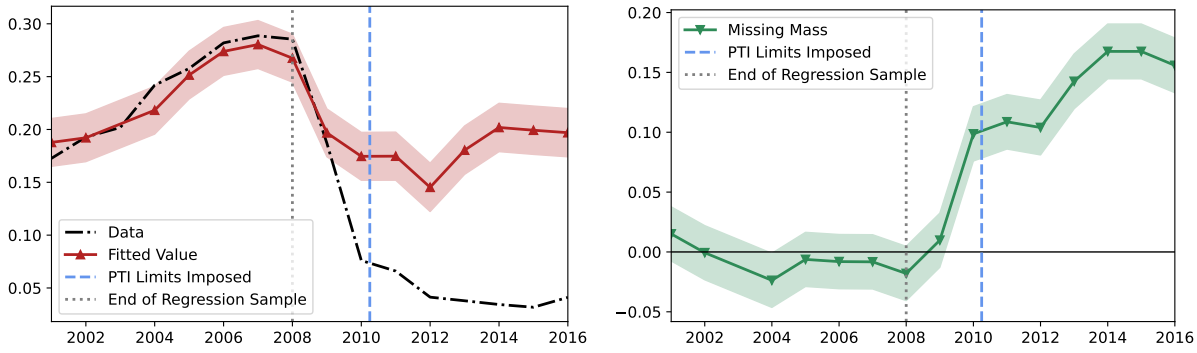

(a) Share PTI ≥ 46

(b) Missing Mass, PTI ≥ 46

Figure 4: High PTI Shares, Model vs. Data

Note: The figure displays the time series of the share of purchase loans with PTI ratios ≥ 46. Panel (a) compares the actual share in the data (dash-dotted line) to the counterfactual fitted values (solid line with markers) implied by the regression (1) with $j = 45$. Panel (b) displays the estimated missing mass, computed as the difference between the fitted and actual shares. Shaded regions denote 95% confidence intervals accounting for uncertainty in the regression estimates. The vertical dotted line indicates the end of the regression sample (2007), and the vertical dashed line indicates the imposition of PTI limits (2010). See Appendix B.2 for additional details.

To obtain the total fraction constrained, I adjust for borrowers with $PTI \in [46, 50]$, who may be constrained if they would have chosen $PTI > 50$ absent the PTI limit (see Appendix B.2). Assuming that eligibility for this higher limit is uncorrelated with unconstrained PTI choices yields a total PTI-constrained share of 18.3%, while assuming this group is either 0% or 100% constrained yields bounds of 16.3% and 19.9%, respectively.

To estimate the LTV-constrained share, I exploit that LTV-constrained borrowers appear in Figure 2(b) as large single-bin spikes, concentrated in the set $S = \{70, 75, 80, 85, 90, 95, 97\}$. To account for borrowers who would have chosen these values absent LTV limits, I subtract the average of the mass in the two neighboring bins from the mass in each bin in S, then sum over S to obtain the share constrained by LTV.[17] This procedure indicates

[17]For example, the excess mass at $LTV = 80$ is equal to the actual share of loans with $LTV = 80$ minus one-half the share of loans with $LTV \in \{79, 81\}$.

that a large majority (73.7%) of borrowers were LTV-constrained in the 2014–2016 period.

To measure the overlap between the LTV- and PTI-constrained groups, I estimate the same distribution regression separately for borrowers in the LTV limit bins (S), and for borrowers in the adjacent bins, again using the latter as a proxy for the non-LTV-constrained mass at those values. Full regression results are displayed in Appendix Tables B.4 and B.5. Combining these estimates provides a measure of the share jointly constrained by both limits, as well as the shares limited by each constraint individually.

Table 1: Share Constrained: Post-Bust Sample

	Purchase Loans				
Share	All Purch.	Repeat	First Time	Cash-Out	All
		Share of Total			
Constrained	78.1%	74.8%	81.8%	42.3%	67.2%
		Share of Constrained			
LTV	94.3%	92.7%	96.6%	67.6%	89.0%
PTI	23.5%	25.4%	17.8%	46.8%	29.5%
Both	17.8%	18.1%	14.4%	14.4%	18.5%
LTV Only	76.5%	74.6%	82.2%	53.2%	70.5%
PTI Only	5.7%	7.3%	3.4%	32.4%	11.0%

Note: The table reports the estimated share of borrowers constrained by loan-to-value (LTV) and/or payment-to-income (PTI) limits for the post-bust sample (2014–2016), computed using the distribution regression methodology described in Section 3.2. The top panel reports the fraction of loans in each column constrained by at least one limit. The subsequent panels decompose these constrained loans into mutually exclusive shares limited by both constraints, LTV only, and PTI only. These group shares sum to 100% of the constrained loans in each column. The column "All" combines purchase and cash-out refinance loans, but omits "rate" refinances that do not change the balance on the loan and hence are not affected by constraints. See Appendix B.2 for additional details.

The results are displayed in Table 1. The first column covers purchase loans, showing that 78.1% of purchase borrowers are constrained by at least one limit. Of these, virtually all (94.3%) are constrained by LTV, leaving only 5.7% exclusively limited by PTI. However, the groups overlap considerably, with 17.8% of constrained borrowers jointly limited by both constraints. As shown in the numerical example of Section 2, and in the structural model of Section 4, this jointly constrained group is central to the strong influence of PTI limits on house prices and credit. In total, a substantial minority (23.5%) of constrained purchase borrowers are PTI-constrained.

The next two columns break purchase loans into repeat and first-time buyers. First-

time borrowers, who lack equity from a prior home sale, are more likely to be LTV-constrained, while repeat buyers are more likely to be PTI-constrained. Still, PTI constraints remain important for both groups, with 14.4% of constrained first-time buyers and 18.1% of constrained repeat buyers jointly limited by both constraints.

The fourth column repeats the exercise for cash-out refinance loans, which exhibit a much higher PTI-constrained share. Unlike for purchase loans, the vast majority of PTI-constrained cash-out refinances are not jointly constrained by LTV. This aligns exactly with the predictions of my theory, as being jointly constrained requires the borrower to choose the specific kink-point house price. Because house prices, interest rates, and incomes are all stochastic, a borrower's current home is unlikely to be appraised at this exact value at the time of refinancing, reducing the share of jointly constrained borrowers among cash-out refinancers.

The last column pools purchase and cash-out refinance loans, showing that 29.5% of all constrained borrowers are PTI-constrained.[18] Since this paper's analysis focuses primarily on purchase loans, these results suggest that my findings on the influence of PTI limits are if anything conservative.

3.3 Links Between PTI Ratios and House Price Growth

I next study cross-sectional links between house price growth during the 2000s boom and the pre-boom PTI-constrained share. The theory of Section 4, motivated by the numerical example of Section 2, predicts that house prices in areas with more jointly constrained households should be more responsive to changes in PTI limits. Given the evidence for a large relaxation of PTI limits documented above, this suggests higher boom house price growth in areas with more PTI-constrained households prior to the boom.

To test this requires a proxy of the PTI-constrained share prior to the boom. Unfortunately, the Fannie Mae loan-level data used above are only available from 2000 onward, by which point PTI limits already appear fully relaxed (see Appendix Figure B.16(a)). As a substitute, I use data from the 1990 Census PUMS.[19] These household-level data contain detailed information on the components of the PTI ratio, including the household's first and second mortgage payments, property taxes, and property insurance payments (see Appendix B.1.5 for details). While these data are surveyed in the cross-section, rather

[18] I omit "rate" refinances where the borrower does not change their balance, as these loans have a fixed size that does not depend on credit constraints.

[19] I thank the referee for suggesting this data source.

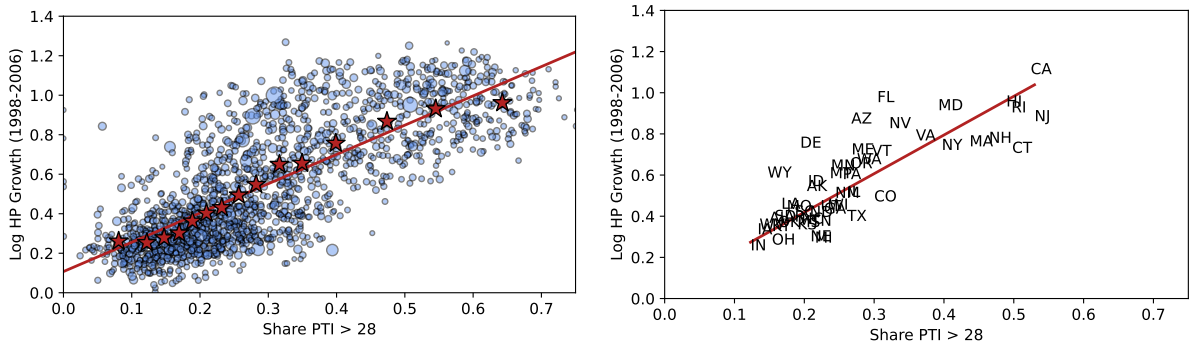

(a) Census: PUMA Data

(b) Census: State Data

Figure 5: High-PTI Share vs. House Price Growth, Census Data

Note: Each circle represents one PUMA in Panel (a) and each pair of letters represents one state in Panel (b). The y-axis variable is log growth in the FHFA house price index from 1998 to 2006. For Panel (a), house price growth is constructed by collapsing FHFA all-transactions 5-digit ZIP code house price indices to PUMAs using the 1990 ZIP-to-PUMA household crosswalk. For Panel (b), house price growth is measured directly using the state-level FHFA index. The x-axis variable is the 1990 Census PUMS share of households with housing payment-to-income ratios above 28% among households with positive mortgage payments and income that moved within the last year, where housing payments include first and second mortgage payments, property insurance, and property taxes. The line in each panel reflects fitted values from a weighted regression of the y-axis variable on the x-axis variable and a constant, using summed 1990 Census PUMS household weights over all households by PUMA or state. The stars in Panel (a) represent a binscatter of averages using 15 equally sized bins.

than computed at origination of the loan, I limit the sample to mortgagor households that moved into their residence within the last year, ensuring that current PTI ratios should be very close to those at origination. I then compute the share of these households with PTI ratio above 28% in each Public Use Microarea (PUMA).

Figure 5 plots this 1990 high-PTI share against boom (1998–2006) log house price growth at the PUMA level (Panel (a)) and state level (Panel (b)), with full results available in Appendix Table B.6. Both panels show strong, approximately linear relationships, with each explaining most of the variation (adjusted R^2) in boom house price growth. The estimated slopes are similar at both geographic levels, predicting 1.5% and 1.9% additional cumulative house price growth, respectively, for each 1pp increase in the high-PTI share.

A natural concern is that the initial share with PTI > 28 is correlated with demand factors unrelated to credit standards, such as local income. To address this, Table 2 presents results after incorporating a battery of controls for non-credit housing fundamentals, including pre-period income, boom income growth, pre-period house price growth, and boom rental growth. These variables provide a powerful set of controls that account for both pre-existing local conditions (via pre-period income and house price growth) and

Table 2: House Price Growth vs. High-PTI Share, Census Data

	(1)	(2)	(3)	(4)	(5)
Constant	0.107***	-0.588***	-1.227***	-1.207***	-1.763***
	(0.010)	(0.222)	(0.226)	(0.225)	(0.300)
Share PTI $>$ 28	1.482***	1.406***	1.323***	1.255***	1.244***
	(0.029)	(0.040)	(0.046)	(0.048)	(0.064)
1990 Census Household Income		0.070***	0.112***	0.116***	0.150***
		(0.022)	(0.022)	(0.022)	(0.029)
Boom Wage Income Growth			1.057***	1.032***	1.047***
			(0.083)	(0.083)	(0.108)
House Price Growth (1993-1998)				-0.246***	0.019
				(0.057)	(0.066)
Boom Rental Growth					0.460***
					(0.107)
N	1,708	1,708	1,546	1,537	830
Adj. R^2	0.573	0.576	0.643	0.649	0.657

Note: This table reports estimates of weighted least squares regressions at the PUMA level. The dependent variable is log growth in the FHFA house price index from 1998 to 2006, constructed by collapsing FHFA all-transactions 5-digit ZIP code house price indices to the PUMA level using the 1990 crosswalk. The main independent variable is the 1990 Census PUMS share of households with housing payment-to-income ratios above 28% among households with positive mortgage payments and income that moved within the last year, where housing payments include first and second mortgage payments, property insurance, and property taxes. 1990 Census Household Income is average household income in the 1990 Census PUMS, aggregated using household weights. Boom Wage Income Growth is the change in log wage and salary income per return from 1998 to 2006 in the IRS data, aggregated from ZIP-level to PUMA-level using the 1990 crosswalk. House Price Growth (1993–1998) is the log change in the FHFA house price index from 1993 to 1998, aggregated from ZIP-level to PUMA-level using the 1990 crosswalk. Boom Rental Growth is the log difference in median gross rent between the 2005–2009 5-year American Community Survey and the 2000 Decennial Census, aggregated from census tract to PUMA using the 2000 crosswalk. Regressions are weighted using summed 1990 Census PUMS household weights over all households by PUMA. Standard errors in parentheses are heteroskedasticity-robust.

subsequent changes in demand factors (via boom income growth and rental growth). In particular, controlling for rental growth should capture changes in housing preferences, demand fundamentals, or supply elasticity, that would affect the price of housing services (rent) but are not driven by a relaxation of PTI limits.

This control specification is aggressive, in that the inclusion of boom-era variables can potentially allow reverse causality, for instance from boom house price growth to boom income growth through local demand, or to boom rental growth through the user cost of housing. Nonetheless, we observe that the core relationship is highly robust to the inclusion of these controls, with the coefficient in the fully controlled specification in column

(5) retaining more than 80% of its uncontrolled value in column (1). The relationship is precise and highly significant in all specifications, with t-statistics on the main coefficient in excess of 19 in all columns.

Next, to check robustness to my construction of the PTI ratio, I perform two alternative estimations. First, Appendix Figure B.8 and Appendix Table B.7 reproduce these results using the share of loans with implied PTI ratios above 28% in the 1990–1995 Home Mortgage Disclosure Act (HMDA) data. While these data do not contain information on the interest rate, property taxes, or property insurance costs, they have the major advantage that they are available at origination and effectively contain the universe of US mortgages issued over this period. I impute values for these missing variables and screen out likely second liens, with full details provided in Appendix B.1.3. The resulting county-level estimates are highly similar in magnitude and predictive power to my benchmark Census results, and are again robust to a variety of controls.

Second, Appendix Figure B.9 and Appendix Table B.8 repeat this exercise using the 2014–2016 share of Fannie Mae loans with $PTI \geq 40$, the region containing the largest fraction of excess mass in Figure 3. While this proxy uses post-bust rather than pre-boom data, it generates highly similar results using directly measured rather than imputed PTI ratios at origination, again supporting the original estimates.

While not causal, these results support a strong link between the pre-boom PTI-constrained share and boom house price growth, consistent with the theory below.

4 The Structural Model

In this section, I develop a structural model designed to capture the empirical patterns evidenced above. The model features both PTI and LTV constraints, along with a tractable form of heterogeneity that yields endogenous and time-varying shares of households constrained by PTI only, LTV only, or both limits. I will refer to these groups as "PTI-only," "LTV-only," and "jointly constrained" households, respectively.

Trend Growth. I assume the true aggregate income process is $Y_t^g = G^t \bar{Y}$ for constants G and $\bar{Y}$. To maintain stationarity, I map all variables to detrended equivalents for an economy with output $\bar{Y}$, as described in Appendix A.2. Unless noted, all variables and equations in the main text correspond to the detrended economy.

Demographics and Preferences. Households exist in two infinitely-lived types: relatively impatient "borrowers," denoted with subscript B, and relatively patient "savers," denoted with subscript S. Households trade a complete set of contracts among their own type, providing perfect insurance against idiosyncratic risk, but cannot trade these securities with households of the other type.

Each borrower i maximizes expected lifetime utility

$$\mathbb{E}_t \sum_{k=0}^{\infty} \beta_B^k \Big[\log(c_{i,t+k}) + e_i^h (1+\omega) \xi \log(h_{i,t+k-1}) \Big], \tag{3}$$

where c is nondurable consumption, h is housing services, e_i^h is an idiosyncratic shifter to the preference for housing, ξ is the base preference for housing, and ω is a utility boost from owned housing. While at equilibrium all households own their homes and therefore receive this utility boost, separating ξ and ω is useful, as ξ pins down the rental price, while ω is used in experiments to capture secular shifts in housing demand.

Each saver i similarly maximizes

$$\mathbb{E}_t \sum_{k=0}^{\infty} \beta_S^k c_{i,t+k}, \tag{4}$$

which differs from the borrower's problem in three ways. First, the saver has discount factor $\beta_S > \beta_B$, which motivates lending and borrowing at equilibrium. Second, savers derive no utility from housing, which they hold in fixed quantity, ensuring that borrowers are the marginal buyers of housing. I discuss the motivation for this assumption further below. Third, savers have linear utility. This makes mortgage interest rates effectively exogenous, allowing me to match observed changes in interest rates without specifying the slope of the aggregate credit supply curve.

Endowments. Each borrower household i earns exogenous labor income equal to

$$y_{i,t} = e_i^y \bar{Y}_B$$

where e_i^y is an idiosyncratic shock with $\mathbb{E}[e_i^y] = 1$, $\bar{Y}_B = \chi_B \bar{Y}$ is detrended borrower income, and χ_B is the borrower income share.

Heterogeneity. Borrower households face two types of idiosyncratic heterogeneity: the preference for housing (e_i^h) and labor income (e_i^y). For parsimony, I assume that these shocks are separate components of a single idiosyncratic draw e_i, so that

$$\log e_i^h = -\frac{1}{2}\gamma^2\sigma_e^2 + \gamma \log e_i \tag{5}$$

$$\log e_i^y = -\frac{1}{2}(1-\gamma)^2\sigma_e^2 - (1-\gamma)\log e_i \tag{6}$$

for $\log e_i \sim N(0, \sigma_e^2)$. The constant terms in (5) and (6) ensure that $\mathbb{E}[e_i^h] = \mathbb{E}[e_i^y] = 1$. Together, e_i^h and e_i^y provide variation in the value-to-income (VTI) ratio, and ultimately determine which debt limit binds, as households become PTI-constrained if their VTI ratio is sufficiently high. Higher VTI ratios can reflect higher housing demand (e_i^h) or lower income (e_i^y), both of which occur as e_i rises. The parameter γ controls the difference in house size across constraint types, with higher γ increasing housing demand for PTI-constrained households relative to LTV-constrained households.

Mortgage Structure. Households trade a mortgage contract with balance m, which is an asset for the saver and a liability for the borrower. Following Chatterjee and Eyigungor (2015), I model the mortgage as a nominal perpetuity with geometrically declining payments. I consider a fixed-rate mortgage, which is the predominant contract in the US. Under this contract, the borrower receives \$1 at time t, and repays $\$(1-\nu)^{k-1}(r_t^* + \nu - \eta)$ at time $t+k$, for all $k > 0$ until prepayment, where r_t^* is the interest rate at origination, ν is the amortization rate, and η is a tax or cost that will generate a constant mortgage spread at equilibrium. For notation, stars (e.g., r_t^*) distinguish values for newly originated loans from those for existing loans, a distinction necessary under long-term fixed-rate debt.

Borrowers become "active" with probability ρ, selling their homes and purchasing new ones. Active borrowers prepay the remaining principal balance on their existing loans, and obtain new mortgages subject to their credit limits.[20]

[20] An earlier version of this paper included an endogenous probability of transacting that varied with house prices and other fundamentals. Allowing for endogenous choices to transact was important for the speed at which rising house prices translated into credit, and for the output response in a New Keynesian setting. However, endogenous prepayment had little impact on the dynamics of house prices or long-term credit growth, the main focus of this paper. To understand this, note that house prices are determined by the households that are currently transacting, regardless of the volume of households that are currently active. As such, the probability of transacting only affects the demand of currently active households through their forecasts of the expected time to wait before transacting again (the ρ term in equation (25)). In practice, the influence of this term is quantitatively small given reasonable variation in ρ.

Debt Limits. A new loan for borrower i must satisfy both an LTV and a PTI constraint:

$$m_{i,t}^* \leq \theta^{LTV} p_t h_{i,t}^* \tag{7}$$

$$\hat{r}_t^* m_{i,t}^* \leq \theta^{PTI} \alpha y_{i,t} \tag{8}$$

where $m_{i,t}^*$ is the balance on the new loan, θ^{LTV} and θ^{PTI} are the maximum LTV and PTI ratios, $\hat{r}_t^* \equiv r_t^* + \nu + \kappa$ is the effective payment rate, and κ reflects taxes and insurance payments. These constraints are treated as institutional and exogenous.[21]

The LTV constraint (7) limits the loan balance as a fraction of house value, equal to the product of the house price p_t and the quantity of housing purchased $h_{i,t}^*$. The key property of the LTV constraint is that it moves proportionally with p_t, increasing as house prices rise. The PTI constraint (8) limits the total initial mortgage payment, composed of principal, interest, taxes, and insurance, as a fraction of adjusted income ($\alpha y_{i,t}$). Scaling income by α accounts for two considerations: (i) recurring payments on non-mortgage debt are counted toward the numerator of the PTI ratio; and (ii) households obtaining new purchase loans at a given date may differ systematically from the average mortgagor (e.g., by being younger). As shown in Section 2, a key property of this limit is that it is highly sensitive to movements in interest rates.

These expressions imply the maximum debt balances

$$\bar{m}_{i,t}^{LTV} = \theta^{LTV} p_t h_{i,t}^* \tag{9}$$

$$\bar{m}_{i,t}^{PTI} = \frac{\theta^{PTI} \alpha y_{i,t}}{\hat{r}_t^*} \tag{10}$$

consistent with each of the two limits. These constraints are applied at loan origination only, so that borrowers are not forced to delever if they exceed these constraints later on.[22]

Housing. Housing produces a service flow each period equal to its stock, and can be owned by both types. A constant fraction δ of house value must be paid as a maintenance cost at the start of each period. As discussed above, only a random and exogenous fraction ρ of borrower households are able to purchase new homes in a given period.

[21] This choice is motivated by the observation that industry standards for these ratios can persist for decades, despite large changes in economic conditions.

[22] At equilibrium, this constraint will bind for all newly issued loans, consistent with Table 1, which shows few unconstrained borrowers at origination. However, households usually wait years between prepayments in the model, during which time they are typically away from their borrowing constraints and accumulating home equity.

To simplify the analysis, I fix the total housing stock to be $\bar{H}$, so that the price of housing p_t fully characterizes the state of the housing market.[23] I further assume that all housing is owner-occupied. The impact of this assumption is discussed extensively in Greenwald and Guren (2025), who show that segmentation frictions between owned and rented housing are very strong, so that house price dynamics are well approximated by a "fully segmented" model with no rental market.

To focus on the use of housing as a collateral asset, I assume that saver demand is fixed at $h_{S,t} = \bar{H}_S$, so that a borrower is always the marginal buyer of housing. By market clearing, this implies $h_{B,t} = \bar{H}_B \equiv \bar{H} - \bar{H}_S$ in each period. As discussed in Greenwald and Guren (2025), this assumption prevents disproportionate housing flows between the two types, which under divisible housing would occur unrealistically along the intensive margin of house size, leading to excessive dampening of house prices.[24]

Government. Both types are subject to a proportional income tax rate τ, whose revenue is used for spending outside the model. Borrower interest payments are tax deductible. The monetary authority achieves a constant rate of inflation π.

Borrower's Problem. Each borrower i maximizes (3) subject to the budget constraint

$$\begin{aligned} c_{i,t} \leq & \underbrace{(1-\tau)y_{i,t}}_{\text{labor income}} - \underbrace{(\pi G)^{-1}\big((1-\tau)x_{i,t-1} + \nu m_{i,t-1}\big)}_{\text{mortgage payment}} - \underbrace{\delta p_t h_{i,t-1}}_{\text{maintenance}} \\ & - \underbrace{p_t(h_{i,t} - h_{i,t-1})}_{\text{new housing}} + \underbrace{\big(m_{i,t} - (1-\nu)(\pi G)^{-1} m_{i,t-1}\big)}_{\text{new debt}} \\ & + \underbrace{a_{i,t-1} - \sum_{s_{t+1}} q_t(s_{t+1}) a_{i,t}(s_{t+1})}_{\text{Arrow securities}} . \end{aligned}$$

[23] Modeling a fixed housing stock precludes the dampening effect of supply on prices. However, from the perspective of credit growth, the key variable is total collateral value: the product of price and quantity. Under a flexible housing supply, smaller movements in price are compensated by larger movements in quantity, leading to similar overall effects. Moreover, my numerical results focus on price-rent ratios. These should not be strongly affected by supply responses, which typically move prices and rents in parallel. For results on spending and output, borrowing used for nondurable consumption in this model would be instead spent on residential investment in a flexible supply specification.

[24] This choice is further consistent with evidence from Landvoigt, Piazzesi, and Schneider (2015). These authors find that overall house price movements over the boom-bust period were primarily driven by the lower end of the price distribution, where borrowers tend to be more credit constrained.

and the debt limits (9) and (10), where $m_{i,t-1}$ and $x_{i,t-1}$ are start-of-period mortgage balance and interest payments. The term $(\pi G)^{-1}$ adjusts nominal mortgage balances and payments for inflation and trend growth between periods, to keep them in real detrended units. Each period, a borrower becomes active with probability ρ and freely selects their housing and mortgage allocations $h_{i,t} = h^*_{i,t}$ and $m_{i,t} = m^*_{i,t}$ subject to their constraints. With probability $(1-\rho)$, the borrower remains inactive with $h_{i,t} = h_{i,t-1}$ and $m_{i,t} = (1-\nu)(\pi G)^{-1} m_{i,t-1}$. The promised interest payment $x_{i,t}$ is then updated using

$$x_{i,t} = \begin{cases} r^*_t m^*_{i,t} & \text{if active} \\ (\pi G)^{-1}(1-\nu) x_{i,t-1} & \text{if inactive.} \end{cases} \tag{11}$$

Saver's Problem. The saver's problem aggregates to that of a representative saver who maximizes lifetime utility (4) subject to the budget constraint

$$c_{S,t} \leq \underbrace{(1-\tau)\chi_S \bar{Y}}_{\text{labor income}} + \underbrace{(\pi G)^{-1}\left(X_{t-1} + (\nu - \eta) M_{t-1}\right)}_{\text{mortgage payment}} - \underbrace{\delta p_t \bar{H}_S}_{\text{maintenance}} - \rho \underbrace{\left(M^*_t - (1-\nu)(\pi G)^{-1} M_{t-1}\right)}_{\text{new debt}}$$

where $\chi_S = 1 - \chi_B$ is the saver income share, and total new mortgage balances M^*_t, mortgage debt M_t, and interest payments X_t aggregate $m^*_{i,t}$, $m_{i,t}$, and $x_{i,t}$, respectively.

4.1 Model Solution

The core of the model solution is the optimality condition for active borrower housing. I present it here, with additional solution details in Appendix A.1. This condition is

$$h^*_{i,t} = \frac{e^h_i \Psi_{1,t}}{(1 - \mu^{LTV}_{i,t} \theta^{LTV}) p_t - \Psi_{2,t}} \tag{12}$$

where $\mu^{LTV}_{i,t}$ is the multiplier on the LTV constraint (9), and $\Psi_{1,t}$ and $\Psi_{2,t}$ are defined by

$$\Psi_{1,t} = \mathbb{E}_t \left\{ \Lambda_{B,t+1} \left[(1+\omega)\xi C_{B,t+1} + (1-\rho)\Psi_{1,t+1} \right] \right\} \tag{13}$$

$$\Psi_{2,t} = \mathbb{E}_t \left\{ \Lambda_{B,t+1} \left[(\rho - \delta) p_{t+1} + (1-\rho)\Psi_{2,t+1} \right] \right\}. \tag{14}$$

For intuition, $\Psi_{1,t}$ is proportional to the present value of the service flows generated by housing, while $\Psi_{2,t}$ is the present value of future financial transactions associated with the house, namely its maintenance and ultimate sale proceeds. In Appendix A.1, I show that the solution to (12) has three regimes:

$$h_{i,t}^{*} = \begin{cases} e_i^h h_t^{LTV} & \text{for } e_i < \bar{e}_t^{LTV} \\ e_i^y h_t^{both} & \text{for } \bar{e}_t^{LTV} \leq e_i \leq \bar{e}_t^{PTI} \\ e_i^h h_t^{PTI} & \text{for } e_i \geq \bar{e}_t^{PTI}, \end{cases} \tag{15}$$

corresponding to LTV-only, jointly constrained, and PTI-only borrowers, for

$$h_t^{LTV} = \frac{\Psi_{1,t}}{(1-\mu_t\theta^{LTV})p_t - \Psi_{2,t}} \tag{16}$$

$$h_t^{both} = \frac{\theta^{PTI}\alpha\bar{Y}_B}{\hat{r}_t^{*}\theta^{LTV}p_t} \tag{17}$$

$$h_t^{PTI} = \frac{\Psi_{1,t}}{p_t - \Psi_{2,t}} \tag{18}$$

where μ_t is the borrower's shadow value of mortgage debt. The cutoffs $(\bar{e}_t^{LTV}, \bar{e}_t^{PTI})$ are

$$\bar{e}_t^{LTV} = K\left(\frac{h_t^{both}}{h_t^{LTV}}\right) \tag{19}$$

$$\bar{e}_t^{PTI} = K\left(\frac{h_t^{both}}{h_t^{PTI}}\right) \tag{20}$$

where $K = \exp\{-(1-2\gamma)\sigma_e^2/2\}$ is a Jensen term close to unity. The key result, echoing Justiniano, Primiceri, and Tambalotti (2019), is that $\bar{e}_t^{LTV} < \bar{e}_t^{PTI}$, implying a nontrivial mass of borrowers jointly limited by both constraints.[25] Given this solution, the house price is determined by the market clearing condition

$$\bar{H}_B = h_t^{LTV}\int^{\bar{e}_t^{LTV}} e_i^h\, d\Gamma(e) + h_t^{both}\int_{\bar{e}_t^{LTV}}^{\bar{e}_t^{PTI}} e_i^y\, d\Gamma(e) + h_t^{PTI}\int_{\bar{e}_t^{PTI}} e_i^h\, d\Gamma(e), \tag{21}$$

which implicitly pins down the price of housing. The key dynamics then follow from differences in the housing demand functions (15) across the constraint groups, as well as

[25]The proof in Appendix A.1 shows that assuming either constraint binds in isolation for $\bar{e}_t^{LTV} < e_i < \bar{e}_t^{PTI}$ would lead to a housing demand choice that causes the other constraint to bind, creating a contradiction.

variation in their shares.

The optimal allocations for households limited by LTV only ($e_i^h h_t^{LTV}$) and PTI only ($e_i^h h_t^{PTI}$) share a similar structure, the main difference being that demand by LTV-constrained households is higher and increasing in θ^{LTV}, due to the marginal value of collateral. In contrast, the optimal demand $e_i^y h_t^{both}$ for the jointly constrained households does not depend on the housing fundamentals $\Psi_{1,t}$ and $\Psi_{2,t}$, and instead varies with components of the debt limits themselves. In particular, it is proportional to the PTI threshold θ^{PTI} and borrower income $\bar{Y}_B$, and inversely proportional to the effective payment rate $\hat{r}_t^*$, the LTV threshold θ^{LTV}, and the house price p_t. Notably, this implies that housing demand for jointly constrained borrowers is decreasing in θ^{LTV}, the reverse of the pattern for LTV-only borrowers. For additional intuition, see Appendix A.3.

Aggregate Mortgage Quantities. Aggregating, the average new mortgage balance is

$$M_t^* = \int^{\bar{e}_t^{LTV}} \bar{m}_{i,t}^{LTV} \, d\Gamma(e) + \int_{\bar{e}_t^{LTV}} \bar{m}_{i,t}^{PTI} \, d\Gamma(e), \tag{22}$$

which integrates over households constrained by LTV (first term) and PTI (second term, including both PTI-only and jointly constrained households), and can be solved in closed form when e is log-normal (see Appendix A.1). Aggregate mortgage balances and promised interest payments evolve according to

$$M_t = \rho M_t^* + (1-\rho)(1-\nu)(\pi G)^{-1} M_{t-1} \tag{23}$$

$$X_t = \rho r_t^* M_t^* + (1-\rho)(1-\nu)(\pi G)^{-1} X_{t-1}. \tag{24}$$

These expressions reflect that a fraction ρ of borrowers obtain new mortgages of average size M_t^* and interest payment $r_t^* M_t^*$, while the remaining fraction $(1-\rho)$ of borrowers retain their existing mortgages with average balance M_{t-1} and interest payments X_{t-1}, subject to adjustments for amortization (ν), inflation (π), and trend growth (G).

4.2 Approximations

I next present two approximations that use statistics on the share of borrowers in each constraint group to summarize house price levels and changes. While the model's quantitative results will be computed exactly in the experiments below, these simplifying approximations provide robust and largely model-free quantitative predictions, as well as

important economic intuition.

Approximating the House Price in Levels. To begin, I present an approximation in levels that derives a valuable summary statistic and connects this framework to earlier models. In Appendix A.4, I show that the equilibrium house price can be approximated by

$$p_t \simeq \frac{\mathbb{E}_t \left\{ \Lambda_{B,t+1} \left[(1+\omega) u^h_{B,t+1} / u^c_{B,t+1} + \Big(1 - \delta - (1-\rho)\mathcal{C}_{t+1}\Big) p_{t+1} \right] \right\}}{1 - \mathcal{C}_t} \tag{25}$$

where $u^h_{B,t}$ and $u^c_{B,t}$ are the marginal utilities of housing services and nondurable consumption for a representative borrower, for

$$\mathcal{C}_t \equiv \mu_t \hat{F}^{LTV}_t \theta^{LTV} \tag{26}$$

$$\hat{F}^{LTV}_t \equiv F^{LTV}_t + \frac{1}{2} F^{both}_t, \tag{27}$$

where F^{LTV}_t and F^{both}_t are the shares of new borrowers constrained by LTV only and by both constraints, respectively.

The term $\mathcal{C}_t$ represents the marginal collateral value of housing — the benefit to the average borrower from a marginal dollar of housing through its ability to relax their debt limit. Division by $1 - \mathcal{C}_t$ in (25) reflects a collateral premium, raising house prices when $\mathcal{C}_t > 0$.[26] If housing could not be used as collateral, we would have $\mathcal{C}_t = 0$, and (25) would imply that the price of housing is the discounted sum of its next-period housing services ($u^h_{B,t+1}/u^c_{B,t+1}$) and the continuation value of housing net of maintenance costs ($(1-\delta)p_{t+1}$). In an economy with an LTV constraint only, such as Iacoviello (2005), we would have $\mathcal{C}_t = \mu_t \theta^{LTV}$, the product of additional borrowing capacity per dollar of housing purchased (θ^{LTV}) and the rate at which the borrower values additional credit (μ_t).

Under both constraints, however, additional housing does not relax the debt limit for PTI-constrained borrowers. As a result, marginal collateral value $\mathcal{C}_{i,t} = \mu^{LTV}_{i,t} \theta^{LTV}$ is borrower-specific, with aggregate $\mathcal{C}_t$ equal to the integral of $\mathcal{C}_{i,t}$ over borrowers. The term $\mu^{LTV}_{i,t}$ is the shadow value of relaxing the LTV constraint. For LTV-only households, this is their sole binding constraint, and they receive the full collateral benefit ($\mu^{LTV}_{i,t} = \mu_t$). For PTI-only households, the LTV constraint is slack, and they receive no benefit ($\mu^{LTV}_{i,t} = 0$).

[26] The appearance of $\mathcal{C}_{t+1}$ in the numerator of (25) reflects that with probability $1 - \rho$ the borrower will not be active in the next period. In this state of the world, the borrower is unable to use their housing to collateralize a new loan, and does not receive the collateral benefit of housing.

For jointly constrained households, $\mu_{i,t}^{LTV}$ varies continuously between 0 and μ_t (see Appendix Figure A.2) reflecting an intermediate marginal value of collateral. Approximating this integral using the trapezoid rule immediately delivers (26) and (27).

This approximation shows that the effect of the constraint structure on house prices can be summarized by a single statistic: $\hat{F}_t^{LTV}$. Since collateral values are highest for LTV-only households, zero for PTI-only households, and intermediate for jointly constrained households, this statistic captures how rising shares of LTV-constrained households will increase house prices, a mechanism I term the *constraint-switching effect*.

Appendix A.5 further shows that (25) closely corresponds to the optimality condition for a simplified model in which households choose housing before observing their idiosyncratic draw e_i, and therefore choose symmetric allocations. The sole difference is that $\hat{F}_t^{LTV}$ in the baseline model becomes F_t^{LTV} in the simplified model, and I show that these quantities are nearly identical in practice.[27] Since this simplified framework nests the precursor to this paper (Greenwald, 2018), and its extensions in Greenwald and Guren (2025) and Allen and Greenwald (2022), these simpler constraint structures can accurately reproduce house price dynamics, despite lacking a nontrivial mass of jointly constrained households.

Elasticities Using Sufficient Statistics. I next derive a simple but powerful approximation of the effect of shocks on house prices and debt. I log-linearize equation (21) in steady state with respect to credit standards, interest rates, and ownership preference to obtain

$$\begin{aligned} d\log p = s_H^{both}\left[d\log\theta^{PTI} - d\log\theta^{LTV} - d\log\hat{r}^*\right] \\ + s_H^{LTV}\left(\frac{\mu\theta^{LTV}}{1-\mu\theta^{LTV}-\psi}\right) d\log\theta^{LTV} + (1-s_H^{both})\, d\omega, \end{aligned} \tag{28}$$

where $\psi \equiv \beta(\rho-\delta)/\left(1-\beta(1-\rho)\right)$, and where s_H^{LTV} and s_H^{both} are the shares of housing purchased by LTV-only and jointly constrained borrowers, respectively.[28] The first term reflects the demand response of jointly constrained borrowers, as h_t^{both} is proportional to θ^{PTI}, but inversely proportional to θ^{LTV} and $\hat{r}^*$. The second term reflects the demand response of LTV-only borrowers, as h_t^{LTV} is increasing in θ^{LTV}. The last term captures that demand for LTV-only and PTI-only borrowers is proportional to ω, but demand for

[27] This is true even though F_t^{both} is substantial, because the jointly constrained mass in the full model ends up evenly divided between the LTV-only and PTI-only groups in the simplified model. See Appendix A.5.

[28] I choose these variables and parameters as they are the main subject of my analysis below. It would be straightforward to approximate this expression around other variables or parameters.

jointly constrained borrowers is invariant to ω in steady state. Since housing demand for all borrowers declines in proportion to price, the price must increase by the same proportion as demand to restore market clearing.

We can also approximate the long-run change in mortgage credit by

$$\begin{aligned} d\log M = (s_M^{both} + s_M^{PTI}) \left(d\log\theta^{PTI} - d\log\hat{r}^* \right) \\ + s_M^{LTV} \left\{ \left(\frac{1-\psi}{1-\mu\theta^{LTV}-\psi} \right) d\log\theta^{LTV} + d\omega \right\}. \end{aligned} \tag{29}$$

For intuition, the first term represents the direct effect of changes in the components of the PTI limit on jointly constrained and PTI-only households. For the second term, corresponding to LTV-only households, note that the coefficient multiplying $s_M^{LTV}\, d\log\theta^{LTV}$ is exactly equal to the equivalent term in (28) plus one, as LTV standards affect the log LTV limit both directly via θ^{LTV} and indirectly via h_t^{LTV}. Last, note that changes to housing preferences (ω) only affect credit for LTV-constrained borrowers, as credit for PTI-constrained borrowers does not depend on house value.

These approximations prove accurate when compared to exact nonlinear perfect foresight paths in my quantitative experiments below, implying that the effects of various shocks can be well approximated using only measures of constraint shares (e.g., Table 1). For instance, a permanent 1% increase in the PTI threshold θ^{PTI} raises house prices by roughly s_H^{both} percent, while a permanent 1pp increase in the log effective payment rate $\log\hat{r}^*$ lowers house prices by a similar amount. A permanent change to the LTV threshold θ^{LTV} has ambiguous sign due to its counteracting effects, increasing housing demand through h_t^{LTV} and decreasing demand through h_t^{both}. Last, a secular increase in preferences ω will have a weaker effect when more borrowers are jointly constrained.

4.3 Calibration

I calibrate the model at quarterly frequency, with parameter values displayed in Table 3. The main calibration targets the mid-1990s pre-boom period using data from 1994–1996. Variable definitions and sources can be found in Appendix Table B.1, while additional details on calibration procedures are contained in Appendix A.7.

Preferences. Following Greenwald and Guren (2025), I set β_B so that borrowers are indifferent in steady state between obtaining an 80% LTV mortgage at the standard interest

Table 3: Parameter Values, Baseline Calibration

Parameter	Name	Value	Internal	Target/Source
		Preferences		
Borr. discount factor	β_B	0.9806	Y	PMI premium = 0.8% (ann.)
Saver discount factor	β_S	0.9920	Y	Real rate target
Housing preference	ξ	0.2500	N	Davis and Ortalo-Magné (2011)
Ownership utility	ω	0.0000	N	Not used in steady state
		Heterogeneity		
Shock dispersion	σ_e	0.4917	N	Fannie Mae Loan Performance Data*
Borrower heterogeneity	γ	0.1466	N	See text*
Income adjustment	α	0.8713	N	See text*
		Housing		
Housing depreciation	δ	0.0050	N	Standard
Log housing stock (borr.)	$\log \bar{H}_B$	1.0000	Y	Normalization
Log housing stock (saver)	$\log \bar{H}_S$	0.3755	Y	Borrower housing share
		Mortgages		
Probability active	ρ	0.0328	N	See text*
Mortgage amortization	ν	0.0038	N	See text*
Max PTI ratio	θ^{PTI}	0.3600	N	Historical limit
Max LTV ratio	θ^{LTV}	0.8000	N	Historical limit
PTI offset (taxes, ins.)	κ	0.0026	Y	See text
Term premium (mean)	η	0.0032	N	Avg. mortgage rate
		Technology and Government		
Growth rate	G	1.0050	N	2% annual growth
Detrended output	$\bar{Y}$	1.000	N	Normalization
Borrower income share	χ_B	0.6210	N	Survey of Consumer Finances
Income tax rate	τ	0.2040	N	Elenev et al. (2016)
Steady state inflation	π	1.0082	N	Avg. infl. expectations, 1998–2002

Note: Stars (*) in the Target/Source column indicate that additional details can be found in Appendix A.7.

rate r^*, and a hypothetical 95% LTV mortgage that also requires private mortgage insurance, priced using the actual FHA scheme.[29] This is supported by the fact that nearly 40% of Fannie Mae purchase borrowers paid for private mortgage insurance to raise their LTV limits during the 2000–2008 boom period, suggesting that the median borrower is close to indifferent between these options. This procedure yields $\beta_B = 0.9806$. I set $\beta_S = 0.9920$

[29] I use the FHA scheme, as its pricing is much simpler than for Fannie Mae or Freddie Mac loans. Goodman and Kaul (2017) show that the total costs of these insurance schemes are similar.

so that the model's implied steady state nominal interest rate π/β_S is equal to the average 10-year treasury yield (6.48% annualized). I set the base housing preference weight ξ to 0.2500 to generate a housing consumption share of 20%, consistent with Davis and Ortalo-Magné (2011). I normalize ownership utility ω to zero; this parameter is used later to simulate a secular increase in housing demand, but is not needed in steady state.

Heterogeneity. To calibrate σ_e, I exploit that PTI limits were effectively slack in the boom period. In the absence of PTI limits, all borrowers are constrained by LTV only, meaning that each loan size is $\theta^{LTV} p_t e_i^h h_t^{LTV}$, and each PTI ratio is

$$PTI_{i,t} = \left(\frac{\hat{r}_t^* \theta^{LTV} p_t h_t^{LTV}}{\alpha Y_t^B} \right) \left(\frac{e_i^h}{e_i^y} \right) \propto \frac{e_i^h}{e_i^y}$$

This proportionality implies $\log PTI_{i,t} = \log \overline{PTI}_t + \log e_i$ for some $\overline{PTI}_t$. Since the distribution of PTI limits is truncated at 65% in the Fannie Mae data, I jointly estimate $\log \overline{PTI}_t$ and σ_e using maximum likelihood, which yields $\log \overline{PTI}_t = -0.9252$ and $\sigma_e = 0.4917$ for the 2004–2006 peak boom period.

I set $\alpha = 0.8713$ so that the model's average PTI ratio at the peak of the "Full Boom" experiment in Section 6 matches the estimated value of $\log \overline{PTI}_t$ for the 2004–2006 Fannie Mae sample. For γ, I similarly exploit that during the boom period, $h_{i,t}^*$ should be proportional to e_i^h, but $PTI_{i,t}$ should be proportional to e_i. I therefore estimate γ by regressing the log house value on the log PTI ratio, which yields $\gamma = 0.1466$ (see Appendix Table A.1).

Housing. I set depreciation to a standard value of $\delta = 0.0050$. I normalize borrower housing to $\bar{H}_B = 1$, and set saver housing to $\bar{H}_S = 0.3755$ to match the share of owner-occupied housing held by mortgagors in the 1998 Survey of Consumer Finances.[30]

Mortgages. The parameter ρ governs the speed at which mortgage credit adjusts to changes in debt limits (M_t^*) in equation (23). I estimate ρ by fitting a nominal and non-detrended version of this equation using aggregate data from the Financial Accounts of the United States. This approach yields $\rho = 0.0328$, which Appendix Figure A.3 shows fits the data well. I calculate ν as the weighted average amortization rate for 30-year fixed-rate mortgages, assuming a constant prepayment rate of ρ, yielding $\nu = 0.0038$. This value is

[30]The parameter $\bar{H}_S$ is of little importance, as it only affects the maintenance payments made by savers. These in turn have no effect on interest rates or credit supply due to the linear utility assumption in (4).

well below the inverse of the loan's maturity ($1/120 = 0.0083$ quarterly), reflecting that amortization schedules are very backloaded, and that most borrowers prepay well before maturity. Last, I set the additional mortgage cost to $\eta = 0.0032$ to match the average 30-year mortgage rate for the 1994–1996 period.

Debt Limits. I set $\theta^{LTV} = 80\%$, which is the most common bunching point for LTV ratios in the data, equals the median combined LTV ratio in the Fannie Mae data in all boom years (see Appendix Figure B.18(a)) and is very close to the average combined LTV ratio of 80.3% among purchase borrowers in these data from 2000–2007. Since the PTI limit was effectively unenforced over this period, these LTV moments provide appropriate targets for θ^{LTV}. I set $\theta^{PTI} = 36\%$ to match the pre-boom underwriting standard.

The parameter κ approximates taxes and insurance, which count toward the PTI ratio's numerator. I calibrate $\kappa = 0.0026$ so that the effective payment rate $\hat{r}_t^*$ in steady state exactly reproduces the total principal, interest, taxes, and insurance on a true 30-year fixed-rate mortgage, given the steady state mortgage interest rate and a standard ratio of annual tax and insurance payments to mortgage balance of 1.75%.

Technology and Government. I set $G = 1.0050$, consistent with 2% real output growth per year, and normalize detrended output $\bar{Y}$ to unity. I set $\chi_B = 0.6210$ to match the income share of mortgagor households in the 1998 Survey of Consumer Finances. I set the income tax rate to $\tau = 0.2040$, following Elenev, Landvoigt, and Van Nieuwerburgh (2016). I set $\pi = 1.0082$ to match 10-year inflation expectations from the Federal Reserve Bank of Cleveland.

4.4 Steady State Properties

Before turning to the main results, I briefly summarize the model's steady-state properties and compare them to the empirical patterns in Section 3. The steady-state price-rent ratio is 15.20, very close to the average empirical value of 14.87 over the calibration period. The steady-state shares of LTV-only, jointly constrained, and PTI-only borrowers are 50.3%, 37.4%, and 12.3%, respectively.[31] These model PTI-constrained shares are higher than those in Table 1 due to differences in aggregate fundamentals: a lower PTI limit (36% pre-boom vs. 45% in the 2014–2016 period), and higher prevailing interest rates.

[31] Note that these population shares "F" differ from housing shares "s_H" as households demand different amounts of housing by type.

To interpret these shares, note that in 1995 the median household income was $34,080, the median house sales price was $133,475, and the average mortgage interest rate was 7.93%.[32] Assuming an 80% LTV ratio and that taxes and insurance are 1.75% of loan balance, the total payment (principal, interest, taxes, and insurance) on a mortgage for this median house would be $914 per month, or 32.2% of the median borrower's income, just below the historical 36% limit. Since this calculation excludes non-mortgage debt payments, which also count toward PTI, it is intuitive that roughly half of borrowers were PTI-constrained at this time, with the vast majority also constrained by LTV.

For more direct validation, I recalibrate the model to the 2014–2016 post-bust period, replicating the procedure in Section 4.3 to match moments from this alternative sample (see Appendix Table A.2). I also recalibrate ω so that the proportional change in the price-rent ratio between the two steady states matches the change between 1994–1996 and 2014–2016 in the data. In this alternative steady state, the jointly constrained share — the key to the model's dynamics — is 12.5%. This is similar to but slightly smaller than the data equivalent of 17.8%, demonstrating that the model's shares of jointly constrained borrowers are realistic and, if anything, conservative.

5 Model Responses to Changes in LTV and PTI Limits

In this section, I present the model's dynamics, showing how the constraint structure affects the response of house prices and debt to various shocks. I first evaluate the effects of changes in credit standards directly, then show how the constraint structure shapes responses to inflation and housing preference shocks through endogenous shifts in borrowing limits. To isolate the role of each credit limit, I compare three economies: (i) the Benchmark Economy described above, (ii) an LTV Economy imposing only an LTV constraint ($\theta^{PTI} \to \infty$), and (iii) a PTI Economy imposing only a PTI constraint ($\theta^{LTV} \to \infty$). To make these economies as parallel as possible, I adjust the θ^{LTV} parameter for the LTV Economy and the θ^{PTI} parameter for the PTI Economy, so that all three economies share the same aggregate mortgage balances in steady state.

[32]The median household income and home sales price are obtained from the US Census Bureau (FRED codes MEHOINUSA646N and MSPUS). The mortgage interest rate is obtained from Freddie Mac (see Table B.1) and is averaged over the year.

5.1 Response to Changes in Credit Standards

To begin, I study how the constraint structure influences the response of house prices and debt to changes in underwriting standards. Figure 6 plots responses to permanent 1% increases in θ^{PTI} (top panel) and θ^{LTV} (bottom panel) across the three economies. Each panel displays nonlinear perfect foresight paths for the price-rent ratio, the sufficient statistic $\hat{F}_t^{LTV}$, and the aggregate loan-to-income (LTI) ratio, computed using the "L-B-J" solution technique of Juillard, Laxton, McAdam, and Pioro (1998). The price-rent and LTI ratio plots also overlay the approximate long-run responses implied by (28) and (29).[33]

Panel (a) shows the response to a 1% increase in θ^{PTI}. The left plot shows that the price-rent ratio rises substantially in the Benchmark economy. This increase is driven by jointly constrained borrowers, whose housing demand rises proportionally with θ^{PTI}. The total rise closely matches the steady-state jointly-constrained share $s_H^{both} = 0.355$, as predicted by (28). The center plot provides further intuition, showing that $\hat{F}_t^{LTV}$ increases as looser PTI limits cause LTV constraints to bind for more households, raising the average marginal value of collateral $\mathcal{C}_t$, and pushing prices up via the constraint switching effect.

In contrast, increasing θ^{PTI} has no observable effect on the price-rent ratio in either the LTV Economy or PTI Economy, confirming that the *interaction* of the constraints, via jointly constrained borrowers, is key for these dynamics. This result is trivial in the LTV Economy, which features no role for this parameter, but is more subtle in the PTI Economy. Without an LTV constraint, mortgage balances in the PTI Economy are largely untethered from housing, and are closer to personal loans, with little effect on housing demand.[34]

Turning to the LTI ratio in the right plot, we observe no change in the LTV Economy (in which PTI limits are irrelevant), and a large increase in the PTI Economy, eventually matching the 1% rise in θ^{PTI}. The Benchmark Economy instead displays a long run change very close to the steady-state share of debt held by all PTI-constrained households $s_M^{PTI} + s_M^{both} = 0.436$, as predicted by (29). Underlying this strong response is the additional credit flowing to the jointly constrained households. Although these households' maximum LTV ratios have not changed, their higher housing demand provides additional collateral for the larger loans available under the higher PTI limit. Thus, even

[33] While (28) was derived for the log price, not the log price-rent ratio, the long-run response should be identical, as steady state rents are a function of aggregate housing and income, and are not affected by the θ^{LTV} and θ^{PTI} parameters. Similarly, (29) also approximates the change in the aggregate log LTI ratio $\log(M/\bar{Y})$, whose response is identical to that of $\log M$ as $\bar{Y}$ does not vary with these parameters.

[34] To see this formally, note that in the PTI Economy, each household purchases $h_{i,t}^* = e_i^h h_t^{PTI}$, and that h_t^{PTI} does not depend on θ^{PTI} or μ_t.

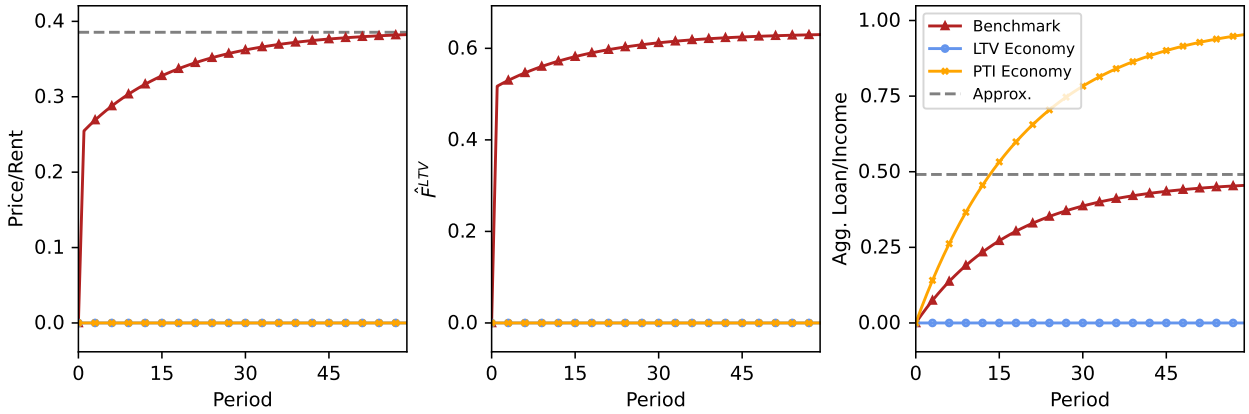

(a) Shock to θ^{PTI}

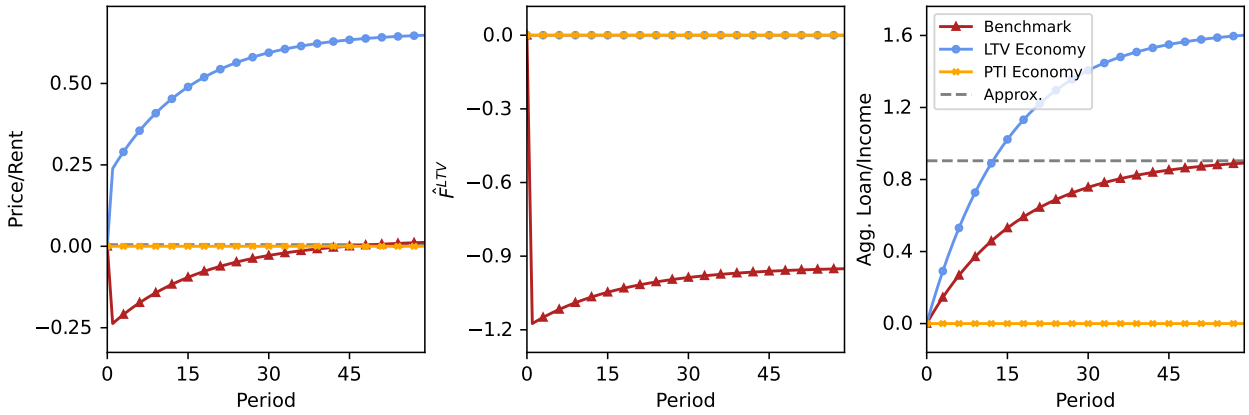

(b) Shock to θ^{LTV}

Figure 6: Responses, Shocks to Credit Standards

Note: The figure displays model responses to a permanent increase of 0.01 to $\log \theta^{PTI}$ (Panel (a)) or $\log \theta^{LTV}$ (Panel (b)). The responses are shown for the Benchmark Economy, the LTV Economy, and the PTI Economy. A value of 1 represents a 1% increase relative to steady state, except for $\hat{F}^{LTV}$, which is measured in percentage point deviations from steady state. The price-rent ratio is defined as $p_t/(u^h_{B,t}/u^c_{B,t})$, where the denominator is the implied price of rental services. Aggregate LTI is defined as $M_t/\bar{Y}$. The horizontal dashed gray lines display the analytical approximations for the long-run changes in the price-rent ratio and aggregate loan-to-income ratio, derived in equations (28) and (29).

though a large majority of the population (87.7%) is LTV-constrained in steady state, the Benchmark economy features a total response of credit to a relaxation of the PTI limit roughly one-half the size we would observe in an economy where all borrowers are PTI-constrained, due to the interaction of the two constraints.

Panel (b) displays the same variables following a permanent 1% increase in θ^{LTV}. Beginning with the price-rent ratio (left plot), we observe a large increase in the LTV Economy, as an increase in housing's effectiveness as collateral pushes up housing demand through h^{LTV}_t. In the Benchmark economy, however, the price-rent ratio is essentially un-

changed at long horizons — the net effect of an increase in demand by LTV-only households, and a decrease in demand by jointly constrained households. My log-linear approximation, again accurate for this case, implies the first effect alone would lead to an increase in the log price-rent ratio of 0.39, while the second effect alone would lead to a change of -0.38, which cancel nearly exactly to zero.

These results can also be understood through the sufficient statistic $\hat{F}_t^{LTV}$ and the constraint-switching effect. In the LTV Economy, marginal collateral value is $\mathcal{C}_t = \mu_t \theta^{LTV}$, which unambiguously increases with θ^{LTV}, raising house prices. In the Benchmark Economy, however, we have $\mathcal{C}_t \simeq \mu_t \hat{F}_t^{LTV} \theta^{LTV}$. The center plot shows that a higher θ^{LTV} decreases $\hat{F}_t^{LTV}$ in the Benchmark Economy, as looser LTV limits cause PTI limits to bind for more households.[35] The net effect of θ^{LTV} increasing and $\hat{F}_t^{LTV}$ decreasing leaves $\mathcal{C}_t$ roughly unchanged, explaining the neutral price-rent response.

For the LTI response, we observe a large increase in the LTV Economy, no change in the PTI Economy (by construction), and an intermediate increase in the Benchmark economy that is again roughly halfway between the two, even though the vast majority of borrowers are LTV-constrained. This is again due to the behavior of jointly constrained borrowers. Although their maximum LTV ratios are higher, these borrowers reduce their housing demand proportionally, leaving their maximum LTV loan size $\bar{m}^{LTV}$ unchanged. Comparing the two panels reveals that jointly constrained households do not simply amplify the response of debt to any change in credit standards. Instead, their presence pulls the responses to changes in *either* credit standard toward those of the PTI Economy, despite the extremely high LTV-constrained share.

5.2 Endogenous Changes via Macroeconomic Variables

Beyond direct changes to credit standards themselves, credit limits can also change endogenously via movements in economic fundamentals. While many variables can influence LTV and PTI limits in this way, I focus on two examples of particular relevance to recent economic events: inflation and secular movements in housing demand.

Figure 7(a) displays responses to a permanent 1pp increase in annualized inflation — of high policy salience following the post-pandemic rise in inflation — for the same set of economies and variables as in Figure 6. Although higher inflation does not directly

[35] More precisely, looser LTV limits cause some LTV-only households to become jointly constrained, and some jointly constrained households to become constrained by PTI only, both of which reduce marginal collateral demand and $\hat{F}_t^{LTV}$.

affect θ^{LTV} or θ^{PTI}, it increases the nominal interest rate, and hence tightens the PTI limit through $\hat{r}^*$. As a result, the relative price-rent responses mirror those of Figure 6(a), with a large response in the Benchmark Economy, and minimal change in the LTV Economy or PTI Economy.[36] The magnitude of the response, however, is an order of magnitude larger. This is because, beginning from an annualized steady state value of $\hat{r}^* = 10.4\%$, a 1 *percentage point* increase corresponds to 10 *percent* change in $\hat{r}^*$. As a result, the impact of inflation on house prices is strongly amplified in the joint presence of LTV and PTI constraints, despite having a small impact under either constraint alone.

Next, Panel (b) shows the effect of a permanent 1pp increase in ω, which acts as a secular shock to housing demand. This type of shock is important as house prices display large swings not easily explained by fundamentals, and are widely thought to be influenced by market sentiment. Here, the Benchmark Economy stands out for having the *smallest* change in price-rent ratios, well below the near-unit rises in both the LTV Economy and PTI Economy. The Benchmark Economy response is instead close to $(1 - s_H^{both}) = 0.645$, as predicted by (28), reflecting that the demand of jointly constrained households is unique in not reacting to this type of preference shock. This can also be interpreted through $\hat{F}_t^{LTV}$, which falls as rising house prices endogenously relax LTV limits, leading more households to become PTI-constrained. For the debt response, (29) shows that since neither PTI-only nor jointly constrained households increase borrowing in response to ω, LTI rises by approximately $s_M^{LTV} = 0.564$, roughly half the response in the LTV Economy.

These results show that the effect of combining PTI and LTV limits depends crucially on the type of shock: those that directly affect PTI limits, such as inflation shocks, are amplified, while those that directly affect LTV limits, such as housing preference shocks, are dampened. These disparate effects underscore the importance of a complete theory of how PTI and LTV limits interact for understanding the mortgage credit channel.

5.3 Robustness

To assess the robustness of my findings, I develop a life cycle extension in Appendix D that replaces the representative borrower in the baseline model with an overlapping gen-

[36]The small decline in the price-rent ratio in the LTV Economy occurs because higher inflation makes mortgage payments more frontloaded in real terms (the "tilt" effect of Modigliani (1976)). Because borrowers are effectively impatient, and prefer to repay debt as late as possible, more frontloaded mortgages are less valued by borrowers, reducing the shadow value of credit μ_t. Since the price of housing partially reflects the marginal value of collateral $\mu_t \theta^{LTV}$ in the LTV economy, this reduces house prices. This effect on house prices does not appear in the PTI Economy, where borrowers also value mortgage credit less, but where they do not use housing to collateralize mortgages, so that μ_t does not affect h_t^{PTI}.

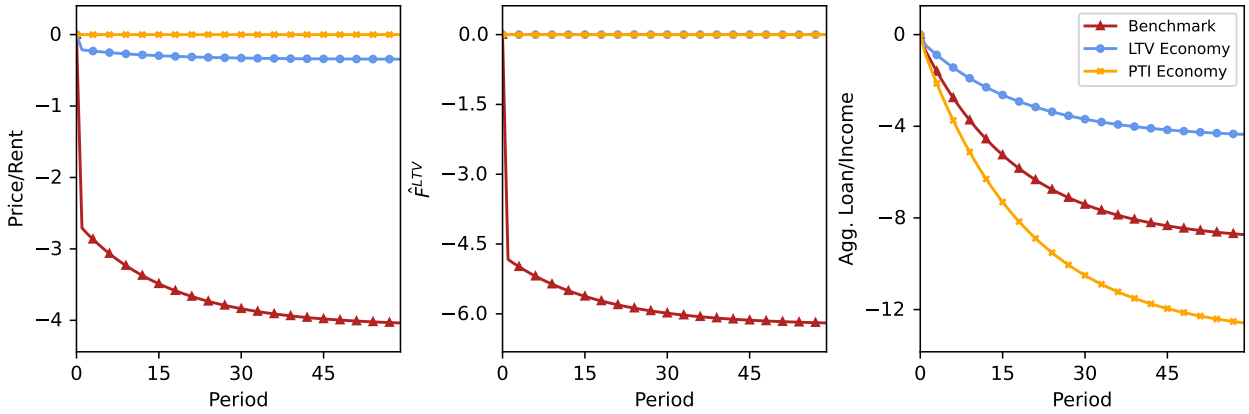

(a) High Inflation Experiment

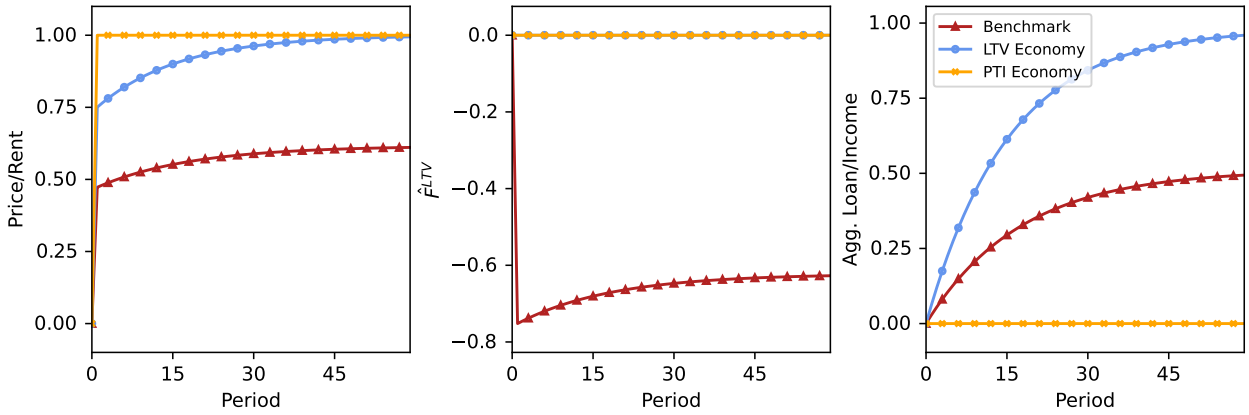

(b) High ω Experiment

Figure 7: Responses, Endogenous Changes to Credit Limits

Note: The figure displays model responses to a permanent increase of 0.0025 (1% annual) to $\log \pi$ (Panel (a)) or a permanent increase of 0.01 in ω (Panel (b)). The responses are shown for the Benchmark Economy, the LTV Economy, and the PTI Economy. A value of 1 represents a 1% increase relative to steady state, except for $\hat{F}^{LTV}$, which is measured in percentage point deviations from steady state. The price-rent ratio is $p_t/(u^h_{B,t}/u^c_{B,t})$, where the denominator is the implied price of rental services. The aggregate loan-income ratio is $M_t/\bar{Y}$.

erations perpetual youth setting following Blanchard (1985). This extension accounts for the deep connections between age and debt limits, due to the need to accumulate a down payment to satisfy the LTV constraint, and the effects of age variation in income on the PTI constraint. The implementation features five age groups, each spanning roughly a decade, with relative income shares calibrated to match the 1998 SCF. Importantly, there are no transfers or risk-sharing arrangements across age groups, so each cohort must self-finance housing purchases from its own accumulated resources.

Although the extension is parsimonious, adding no parameters beyond the income

profile and aging probability, it captures the strong links between age, housing choices, and constraint status. Appendix Figure D.1 shows that the extended model matches the age profile of housing wealth and value-to-income ratios well. Young households are the most likely to be LTV-constrained, consistent with the first-time buyer estimates in Table 1. While these households have the lowest incomes, they choose small houses even relative to income, as their low consumption increases the opportunity cost of a larger down payment. Households in their 50s, while having the highest incomes, are the most likely to be PTI-constrained, as their accumulated home equity provides ample down payments for large homes. This variation generates substantial heterogeneity by age group in the responses of housing and credit demand to the shocks studied above.

Despite this within-borrower heterogeneity, the *aggregate* dynamics of house prices and credit in the life-cycle model remain nearly identical to those of the Benchmark model, as shown in Appendix Figures D.3 and D.5. To understand this result, recall from equation (17) that the housing demand of jointly-constrained households — the key to the model's dynamics — satisfies a corner solution that equalizes the LTV and PTI constraints. The resulting marginal demand elasticities depend on the structure of the constraints themselves, but not on other blocks of the model such as preferences, income expectations, or the life cycle.

To formalize this intuition, Appendix D.2 generalizes the log-linear approximations (28) and (29) to allow for a flexible form of borrower heterogeneity nesting the life-cycle case. These formulas show that, to first order, the response to shocks affecting PTI limits — whether through credit standards or endogenous variables — depends only on the aggregate housing share of jointly-constrained borrowers. While other model features will influence the shares of households in each constraint group, these shares can also be estimated in the data, and can serve as calibration targets. Thus, a wide set of models should have approximately identical responses to shocks affecting the PTI limit, provided each is correctly calibrated to match the empirical shares limited by each constraint.

6 Application: The 2000s Housing Boom

Having established how the constraint structure shapes the model's dynamics, I now use the model to quantify the effect of the dramatic expansion of PTI limits during the 2000s boom, as implied by Figure 2. To do so, I compute three sets of model responses: (i) to changes in credit standards, (ii) to a broader set of shocks that collectively explain the

entire boom, and (iii) to these same shocks under alternative macroprudential policies.

Table 4: Results: Boom Experiments

Experiment	Price-Rent	(Of Actual)	LTI	(Of Actual)
Data	53%	–	70%	–
Figure 8(a): Credit Liberalization Experiments				
1. LTV Liberalized	−3%	(−6%)	14%	(20%)
2. PTI Liberalized	17%	(33%)	18%	(26%)
3. Both Liberalized	34%	(64%)	64%	(91%)
Figure 8(b): Full Boom + Macroprudential Policy Counterfactual				
4. Full Boom	53%	(100%)	61%	(86%)
5. Macropru. Counterfactual	25%	(47%)	29%	(41%)
Appendix Figure A.5: Robustness				
6. Full Boom (LTV Lib)	53%	(100%)	50%	(70%)
7. Macropru. Counterfactual (LTV Lib)	64%	(121%)	46%	(65%)
8. LTV Model, LTV Liberalized	17%	(31%)	41%	(58%)
9. LTV Model, Full Boom	53%	(100%)	95%	(134%)
10. LTV Model, Macropru. Counterfactual	27%	(51%)	31%	(44%)

Note: This table compares the peak responses of the price-rent ratio and aggregate loan-to-income ratios in Figure 8a and 8b to those in the data. The "Data" row represents the peak responses for each series in the data, from 1998:Q1 to 2006:Q1 for the price-rent ratio, and from 1998:Q1 to 2007:Q3 for the loan-to-income ratio. The remaining rows represent total growth in the model from the start of the experiment to the peak of the model boom in 2006:Q4. The columns labeled "Price-Rent" and "LTI" represent growth in the price-rent ratio $p_t/(u^h_{B,t}/u^c_{B,t})$ and loan-to-income (LTI) ratio $M_t/\bar{Y}$, while the "(Of Actual)" columns divide these values by their respective values in the Data row.

Each experiment computes nonlinear perfect foresight paths, as in Section 5, with results summarized in Table 4. Starting from steady state, households learn at $t = 1$ (1998 Q1) of a permanent change to a set of parameters, then unexpectedly learn at $t = 36$ that these parameters have permanently reverted to their baseline values, marking the onset of the bust. This timing implies a boom through 2006 Q4, selected as a compromise between the peaks of price-rent ratios (2006 Q1) and LTI ratios (2007 Q3). For each experiment, I report the rise in price-rent and LTI ratios over the model boom period, compared to their peak increases in the data (53% and 70%, respectively).

These analyses extend those of Section 5 in two ways. First, the shock magnitudes are taken from the data, which matters not only for reproducing this historical episode, but also because the proportional responses can vary with shock size due to nonlinearities. Second, applying multiple shocks simultaneously reveals interactions between them that

prove central to the implications for macroprudential policy.

Credit Liberalization Experiments. To begin, Figure 8(a) shows responses to changes in credit standards. The PTI Liberalized experiment increases θ^{PTI} from 0.36 to 0.70, while the LTV Liberalized experiment increases θ^{LTV} from 0.80 to 0.99, with each change reversed at $t = 36$. The PTI Liberalized experiment is motivated by Figure 2, which shows no bunching in PTI ratios up to the truncation point of 65%, implying an effective PTI limit substantially higher than 65% if imposed at all. The LTV Liberalized experiment, in contrast, is not motivated by an empirical moment, as Figure 2 and Appendix Figure B.18 do not imply large changes in LTV limits over this period, consistent with Adelino, McCartney, and Schoar (2025). Instead, this experiment models a near-full relaxation of the LTV limit, to give this mechanism the best possible chance to make a quantitatively important contribution to the boom.

Figure 8(a) shows that the qualitative findings from Figure 6 carry over to these larger shocks. The PTI Liberalization experiment explains 33% of the observed rise in price-rent ratios, confirming the relaxation of PTI limits as an important driver of the boom. Still, the log price-rent ratio rises only around 60% as much per unit of change in $\log \theta^{PTI}$ as in Figure 6(a), showing the importance of a nonlinear solution given this large shock size. This attenuation reflects that as the PTI limit is loosening, the share of jointly constrained borrowers is falling, reducing the marginal strength of further changes in θ^{PTI} in (28).

In contrast, a large LTV liberalization slightly decreases house prices, as lower demand from jointly constrained households dominates higher demand from LTV-only households. This result is entirely due to the PTI limit, as a similar liberalization in the LTV Economy explains 31% of the observed rise in price-rent ratios (see Appendix Figure A.5).

Turning to credit volumes, we observe a slightly larger increase following the PTI liberalization than following the LTV liberalization, explaining 26% and 20% of the observed rise, respectively. Although Figure 6 showed that a marginal change in $\log \theta^{LTV}$ has a larger per-unit effect on aggregate LTI than a marginal change in $\log \theta^{PTI}$, the near-doubling of θ^{PTI} from 0.36 to 0.7 overwhelms the smaller change in θ^{LTV} from 0.8 to 0.99. That θ^{PTI} can increase by such a large proportion, while θ^{LTV} is naturally bounded above by unity, is an important consideration for macroprudential policy.

Next, the "Both Liberalized" experiment applies both liberalizations simultaneously. Although the empirical evidence for such an LTV liberalization is weak, as discussed above, this experiment reveals important interactions between changes in LTV and PTI

standards. Relaxing both limits generates a much larger boom than loosening PTI limits alone, explaining 64% and 91% of the observed rises in the price-rent and LTI ratios. This shows that changes in LTV standards are greatly amplified under loose PTI limits, pointing to an additional macroprudential benefit of PTI limits.

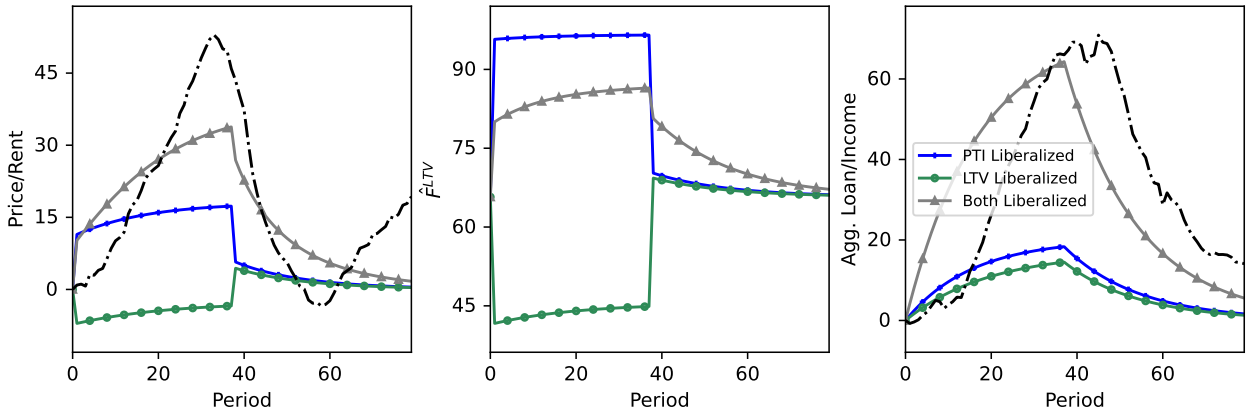

(a) Credit Liberalization Experiments

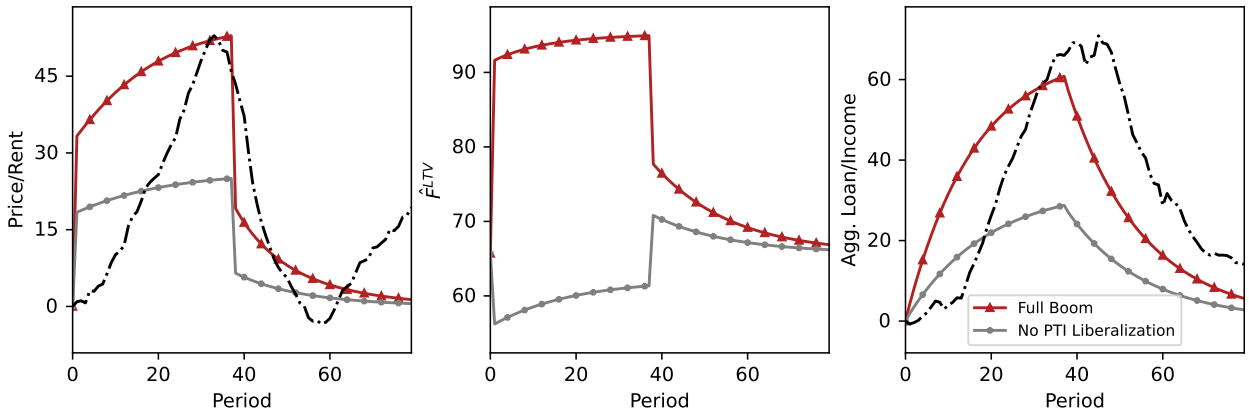

(b) Macroprudential Policy Counterfactuals

Figure 8: Housing Boom Experiments

Note: Paths represent responses of the Benchmark model to various experiments (see text for definitions). A value of 1 represents a 1% increase relative to steady state, except for $\hat{F}_t^{LTV}$, which is measured in percentage points. The price-rent ratio is $p_t/(u^h_{B,t}/u^c_{B,t})$, where the denominator is the implied price of rental services. The aggregate loan-income ratio is $M_t/\bar{Y}$. Where applicable, the black dash-dot lines display the corresponding paths in the data.

Full Boom Experiment. Assessing PTI limits as a macroprudential tool requires understanding their influence on the full set of shocks that drove the housing boom. In particular, this era featured falling mortgage rates and what are widely considered overoptimistic expectations. To incorporate these factors, I augment the PTI Liberalization experiment

with permanent changes to log inflation ($\log \pi$) and the mortgage spread (η) of -0.0021 and -0.0025, to match average mortgage rates and 10-year inflation expectations over the 2004–2006 period.[37] I then increase the ownership utility bonus ω from zero to 0.1980 so that the combined experiment reproduces the peak rise in the price-rent ratio.

The resulting "Full Boom" responses in Figure 8(b) display a good overall fit of the data. The rise in the aggregate LTI ratio, an untargeted moment, closely matches its empirical counterpart, while the increase in $\hat{F}_t^{LTV}$ to nearly 100% is consistent with the evidence in Figure 2 that virtually no households appear PTI-constrained during the boom. The main discrepancy is that the model's price-rent ratio increases too rapidly — a standard feature of rational expectations models in which asset valuations react immediately to news, compared to the more inertial dynamics observed in the data (Guren, 2018).

To assess the macroprudential effect of PTI limits, I compute a "No PTI Liberalization" counterfactual that applies the same shocks to π, η, and ω, but holds θ^{PTI} at its steady state value.[38] Maintaining the historical PTI limit reduces the rises in price-rent ratios and aggregate LTI ratios by 53% and 45% of their observed values, respectively.[39] These reductions far exceed the direct effects of PTI liberalization in Figure 8(a) (33% and 26%), because loose PTI limits also amplify the "overoptimism" shock to ω. While the presence of jointly constrained borrowers limits the effect of changes in ω, relaxing PTI limits by such a large amount drives their share close to zero, removing this dampening force.

These macroprudential effects are specific to PTI limits. To show this, Appendix Figure A.5(a) considers an analogous experiment in which LTV limits are liberalized, interest rates fall, and ω rises to match the peak increase in the price-rent ratio. Unlike the results for the PTI limit in Figure 8, removing the LTV liberalization from the same experiment does not meaningfully dampen the boom. Instead, the peak increase in price-rent ratios is even larger, while the rise in LTI is barely reduced (see rows 6 and 7 of Table 4). These macroprudential effects of LTV and PTI limits differ because house prices rise much faster than incomes during housing booms. As a result, households have enough collateral to satisfy LTV limits holding θ^{LTV} fixed, but struggle to finance large mortgages without an increase in θ^{PTI}. These findings also depend on the simultaneous imposition of LTV and PTI limits, as Appendix Figure A.5(b) shows that a binding LTV cap would strongly dampen a similar boom in the LTV Economy (see rows 8–10 of Table 4).

[37] Incorporating this decline in rates alongside the PTI liberalization augments the rise in the price-rent and loan-income ratios, explaining 57% and 52% of the observed rises, respectively.

[38] This exercise implicitly assumes that other factors, such as the level of overoptimism, are the same across experiments. Further work is needed to study how credit standards affect house price expectations.

[39] These statistics can be computed by subtracting Row 5 from Row 4 in Table 4.

7 Conclusion

In this paper, I studied how the credit constraint structure influences housing and mortgage markets. I showed empirically that a substantial minority of households are constrained by PTI limits in the post-bust period, but that these limits were virtually unenforced during the 2000s boom. To model this, I developed a general equilibrium framework in which heterogeneous households are constrained by LTV limits, PTI limits, or both. The presence of jointly constrained households amplifies shocks to PTI standards and inflation, but dampens shocks to LTV standards and housing preferences.

I found that a PTI liberalization was central to house price dynamics over the boom period, both through its direct contribution, and through its amplification of other forces. PTI liberalization explains between 33% and 53% of the observed increase in price-rent ratios, depending on whether PTI limits are relaxed in isolation or alongside other boom shocks. These results imply that restricting PTI ratios, rather than LTV ratios, would have been the more effective macroprudential policy for limiting the 2000s housing boom.

References

Adelino, Manuel, W Ben McCartney, and Antoinette Schoar. 2025. "The Missing LTV Cycle." *American Economic Review: Insights* Forthcoming.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino. 2016. "Loan Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class." *The Review of Financial Studies* 29 (7):1635–1670.

———. 2025. "Credit Supply and House Prices: Evidence From Mortgage Market Segmentation." *Journal of Financial Economics* 163:103958.

Akinci, Ozge and Jane Olmstead-Rumsey. 2018. "How Effective Are Macroprudential Policies? An Empirical Investigation." *Journal of Financial Intermediation* 33:33–57.

Aladangady, Aditya. 2014. "Homeowner Balance Sheets and Monetary Policy." Tech. rep., FEDS Working Paper.

Albanesi, Stefania, Giacomo DeGiorgi, and Jaromir Nosal. 2022. "Credit Growth and the Financial Crisis: A New Narrative." *Journal of Monetary Economics* 132:118–139.

Allen, Jason and Daniel L Greenwald. 2022. "Managing a Housing Boom." Working Paper.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra. 2019. "Regional Heterogeneity and the Refinancing Channel of Monetary Policy." *The Quarterly Journal of Economics* 134 (1):109–183.

Blanchard, Olivier J. 1985. "Debt, Deficits, and Finite Horizons." *Journal of Political Economy* 93 (2):223–247.

Boar, Corina, Denis Gorea, and Virgiliu Midrigan. 2022. "Liquidity Constraints in the US Housing Market." *The Review of Economic Studies* 89 (3):1120–1154.

Bosshardt, Joshua, Marco Di Maggio, Ali Kakhbod, and Amir Kermani. 2024. "The Credit Supply Channel of Monetary Policy Tightening and Its Distributional Impacts." *Journal of Financial Economics* 160:103914.

Campbell, Jeffrey R and Zvi Hercowitz. 2005. "The Role of Collateralized Household Debt in Macroeconomic Stabilization." Tech. rep., National Bureau of Economic Research.

Campbell, John Y and Joao F Cocco. 2015. "A Model of Mortgage Default." *The Journal of Finance* 70 (4):1495–1554.

Chatterjee, Satyajit and Burcu Eyigungor. 2015. "A Quantitative Analysis of the US Housing and Mortgage Markets and the Foreclosure Crisis." *Review of Economic Dynamics* 18 (2):165–184.

Chernozhukov, Victor, Iván Fernández-Val, and Blaise Melly. 2013. "Inference on Counterfactual Distributions." *Econometrica* 81 (6):2205–2268.

Chow, Gregory C and An-loh Lin. 1971. "Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series." *The Review of Economics and Statistics* :372–375.

Corbae, Dean and Erwan Quintin. 2015. "Leverage and the Foreclosure Crisis." *Journal of Political Economy* 123 (1):1–65.

Davis, Morris A and François Ortalo-Magné. 2011. "Household Expenditures, Wages, Rents." *Review of Economic Dynamics* 14 (2):248–261.

DeFusco, Anthony A, Stephanie Johnson, and John Mondragon. 2020. "Regulating Household Leverage." *The Review of Economic Studies* 87 (2):914–958.

Di Maggio, Marco and Amir Kermani. 2017. "Credit-Induced Boom and Bust." *The Review of Financial Studies* 30 (11):3711–3758.

Elenev, Vadim, Tim Landvoigt, and Stijn Van Nieuwerburgh. 2016. "Phasing Out the GSEs." *Journal of Monetary Economics* 81:111–132.

Favara, Giovanni and Jean Imbs. 2015. "Credit Supply and the Price of Housing." *American Economic Review* 105 (3):958–992.

Favilukis, Jack, Sydney C Ludvigson, and Stijn Van Nieuwerburgh. 2017. "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium." *Journal of Political Economy* 125 (1):140–223.

Foote, Christopher L, Lara Loewenstein, and Paul Willen. 2018. "Technological Innovation in Mortgage Underwriting and the Growth in Credit: 1985–2015." FRB of Cleveland Working Paper.

Foote, Christopher L, Lara Loewenstein, and Paul S Willen. 2021. "Cross-Sectional Patterns of Mortgage Debt During the Housing Boom: Evidence and Implications." *The Review of Economic Studies* 88 (1):229–259.

Ghent, Andra. 2012. "Infrequent Housing Adjustment, Limited Participation, and Monetary Policy." *Journal of Money, Credit and Banking* 44 (5):931–955.

Goodman, Laurie and Karan Kaul. 2017. "Sixty Years of Private Mortgage Insurance in the United States." *Urban Institute, August* .

Greenwald, Daniel L. 2018. "The Mortgage Credit Channel of Macroeconomic Transmission." Working Paper.

Greenwald, Daniel L and Adam Guren. 2025. "Do Credit Conditions Move House Prices?" *American Economic Review* 115 (10):3559–3596.

Guler, Bulent. 2015. "Innovations in Information Technology and the Mortgage Market." *Review of Economic Dynamics* 18 (3):456–483.

Guren, Adam M. 2018. "House Price Momentum and Strategic Complementarity." *Journal of Political Economy* 126 (3):1172–1218.

Higgins, Brian E. 2024. "Mortgage Borrowing Limits and House Prices: Evidence From a Policy Change in Ireland." ECB Working Paper, No. 2909.

Iacoviello, Matteo. 2005. "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle." *American Economic Review* :739–764.

Iacoviello, Matteo and Stefano Neri. 2010. "Housing Market Spillovers: Evidence From an Estimated DSGE Model." *American Economic Journal: Macroeconomics* 2 (2):125–164.

Iacoviello, Matteo and Marina Pavan. 2013. "Housing and Debt Over the Life Cycle and Over the Business Cycle." *Journal of Monetary Economics* 60 (2):221–238.

Johnson, Stephanie. 2020. "Mortgage Leverage and House Prices." Working Paper, Rice University.

Johnson, Stephanie and Nitzan Tzur-Ilan. 2025. "Financial Technology and the 1990s Housing Boom." FRB of Dallas Working Paper.

Juillard, Michel, Douglas Laxton, Peter McAdam, and Hope Pioro. 1998. "An Algorithm Competition: First-Order Iterations Versus Newton-Based Techniques." *Journal of Economic Dynamics and Control* 22 (8):1291–1318.

Jung, Hosung and Jieun Lee. 2017. "The Effects of Macroprudential Policies on House Prices: Evidence From an Event Study Using Korean Real Transaction Data." *Journal of Financial Stability* 31:167–185.

Justiniano, Alejandro, Giorgio E Primiceri, and Andrea Tambalotti. 2019. "Credit Supply and the Housing Boom." *Journal of Political Economy* 127 (3):1317–1350.

Kaplan, Greg, Kurt Mitman, and Giovanni L Violante. 2020. "The Housing Boom and Bust: Model Meets Evidence." *Journal of Political Economy* 128 (9):3285–3345.

Kermani, Amir. 2012. "Cheap Credit, Collateral and the Boom-Bust Cycle." Unpublished Paper, Haas School of Business, University of California, Berkeley.

Kuttner, Kenneth N and Ilhyock Shim. 2016. "Can Non-Interest Rate Policies Stabilize Housing Markets? Evidence From a Panel of 57 Economies." *Journal of Financial Stability* 26:31–44.

Landvoigt, Tim. 2015. "Financial Intermediation, Credit Risk, and Credit Supply During the Housing Boom." Working Paper, UT Austin, McCombs School of Business.

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider. 2015. "The Housing Market(s) of San Diego." *American Economic Review* 105 (4):1371–1407.

Liu, Zheng, Pengfei Wang, and Tao Zha. 2013. "Land-Price Dynamics and Macroeconomic Fluctuations." *Econometrica* 81 (3):1147–1184.

Mian, Atif and Amir Sufi. 2014. "House Price Gains and US Household Spending From 2002 to 2006." Tech. rep., National Bureau of Economic Research.

Missouri Census Data Center. 1990. "Geocorr 1990: Geographic Correspondence Engine." `https://mcdc.missouri.edu/applications/geocorr1990.html`. ZIP code to PUMA crosswalk generated using the Geocorr 1990 geographic correspondence engine. Accessed May 22, 2026.

Modigliani, Franco. 1976. "Some Economic Implications of the Indexing of Financial Assets With Special Reference to Mortgages." In *The New Inflation and Monetary Policy*. Springer, 90–141.

Monacelli, Tommaso. 2008. "Optimal Monetary Policy With Collateralized Household Debt and Borrowing Constraints." In *Asset Prices and Monetary Policy*. University of Chicago Press, 103–146.

Peydró, José-Luis, Francesc Rodriguez-Tous, Jagdish Tripathy, and Arzu Uluc. 2024. "Macroprudential Policy, Mortgage Cycles, and Distributional Effects: Evidence From the United Kingdom." *The Review of Financial Studies* 37 (3):727–760.

Piazzesi, Monika, Martin Schneider, and Selale Tuzel. 2007. "Housing, Consumption and Asset Pricing." *Journal of Financial Economics* 83 (3):531–569.

Rognlie, Matthew, Andrei Shleifer, and Alp Simsek. 2018. "Investment Hangover and the Great Recession." *American Economic Journal: Macroeconomics* 10 (2):113–153.

Wong, Arlene. 2015. "Population Aging and the Transmission of Monetary Policy to Consumption." Working Paper. Northwestern University.

Appendix

This appendix is structured as follows. Section A provides solutions, derivations, and extensions of the structural model. Section B provides additional results and details for the empirical exercises. Section C provides additional detail and structure for the simple numerical example. Section D presents the life-cycle extension.

A Model Appendix

A.1 Model Solution

This section supplements Section 4.1 by deriving the full set of optimality conditions for the model.

A.1.1 Borrower Optimality

Each borrower solves the problem

$$V_{i,0} = \max \sum_{t=0}^{\infty} \sum_{s^t} \beta_B^t \Big[\log c_{i,t}(s^t) + e_i^h (1+\omega) \xi \log h_{i,t-1}(s^t) \Big] \phi_t(s^t) \tag{A.1}$$

where $c_{i,t}(s^t)$ is nondurable consumption after history s^t, $h_{i,t}(s^t)$ is housing consumption after history s^t, assumed equal to the stock of owned housing, e_i^h is a randomly drawn preference for housing, specified in (5), and $\phi_t(s^t)$ is the probability of history s^t. In the remainder of the text, I suppress the s^t notation, but all time-varying states should be considered dependent on the history to that date. Borrowers can trade a complete set of contracts contingent on both the aggregate states and the draws of e_i.

The borrower is subject to the intertemporal budget constraint

$$\sum_{t=0}^{\infty} \sum_{s^t} q_t^0 \Bigg\{ (1-\tau) e_i^y \bar{Y}_B - c_{i,t} - \underbrace{\pi^{-1} G^{-1} \Big((1-\tau) x_{i,t-1} + \nu m_{i,t-1} \Big)}_{\text{mortgage payment}} - \underbrace{\delta p_t h_{i,t-1}}_{\text{maintenance}} + \mathbb{I}_{i,t} \Bigg[\underbrace{\Big(m_{i,t}^* - (1-\nu) \pi^{-1} G^{-1} m_{i,t-1} \Big)}_{\text{new debt}} - \underbrace{p_t (h_{i,t}^* - h_{i,t-1})}_{\text{new housing}} \Bigg] \Bigg\} \geq 0$$

where $\mathbb{I}_{i,t}$ is an indicator for whether the borrower adjusts this period, $e_i^y Y_{B,t}$ is labor income, with e_i^y specified as in (6), π is the (constant) rate of inflation, G is the gross

growth rate of the economy, used for detrending (see Appendix A.2), r_t^* is the interest rate on new mortgages, p_t is the price of housing, $m_{i,t-1}, x_{i,t-1}$, $h_{i,t-1}$, are start-of-period mortgage balance, mortgage payment, and owned housing, respectively, and $m_{i,t}^*$ and $h_{i,t}^*$ are the new debt balance and owned housing conditional on adjusting. Adjusting borrowers are also subject to the LTV and PTI constraints (7) and (8).

The laws of motion for the state variables in the growing economy (to be detrended in Section A.2) are

$$m_{i,t} = \mathbb{I}_{i,t} m_{i,t}^* + (1 - \mathbb{I}_{i,t})(1-\nu)\pi^{-1}G^{-1}m_{i,t-1} \tag{A.2}$$

$$x_{i,t} = \mathbb{I}_{i,t} r_t^* m_{i,t}^* + (1 - \mathbb{I}_{i,t})(1-\nu)\pi^{-1}G^{-1}x_{i,t-1} \tag{A.3}$$

$$h_{i,t} = \mathbb{I}_{i,t} h_{i,t}^* + (1 - \mathbb{I}_{i,t})h_{i,t-1} \tag{A.4}$$

where the π^{-1} and G^{-1} terms deflate and detrend the variables as needed. Note that $h_{i,t}$ is not detrended or deflated, as it is a real variable and does not grow with economic output. Combining these elements, the Lagrangian for the borrower's problem in the growing (non-detrended) economy can be written

$$\begin{aligned}
\mathcal{L} = \sum_{t=0}^{\infty}\sum_{s^t}\beta^t\Big[\log c_{i,t} + e_i^h(1+\omega)\xi\log h_{i,t-1}\Big]\phi_t \\
+ \lambda_i \sum_{t=0}^{\infty}\sum_{s^t} q_t^0 \Bigg\{ (1-\tau)y_{i,t} - c_{i,t} - \pi^{-1}G^{-1}\Big((1-\tau)x_{i,t-1} + \nu m_{i,t-1}\Big) - \delta p_t h_{i,t-1} \\
+ \mathbb{I}_{i,t}\Big[\Big(m_{i,t}^* - (1-\nu)\pi^{-1}G^{-1}m_{i,t-1}\Big) - p_t(h_{i,t}^* - h_{i,t-1})\Big] \\
+ \mu_{i,t}^{LTV}\Big(\theta^{LTV}p_t h_{i,t}^* - m_{i,t}^*\Big) \\
+ \mu_{i,t}^{PTI}\Big(\frac{\theta^{PTI}y_{i,t}}{\hat{r}_t^*} - m_{i,t}^*\Big) \\
- \Omega_{i,t}^m\Big(\mathbb{I}_{i,t}m_{i,t}^* + (1-\mathbb{I}_{i,t})(1-\nu)\pi^{-1}G^{-1}m_{i,t-1} - m_{i,t}\Big) \\
- \Omega_{i,t}^x\Big(\mathbb{I}_{i,t}r_t^*m_{i,t}^* + (1-\mathbb{I}_{i,t})(1-\nu)\pi^{-1}G^{-1}x_{i,t-1} - x_{i,t}\Big) \\
+ \Omega_{i,t}^h\Big(\mathbb{I}_{i,t}h_{i,t}^* + (1-\mathbb{I}_{i,t})h_{i,t-1} - h_{i,t}\Big)\Bigg\}.
\end{aligned}$$

The optimality conditions are

$$(c_{i,t}): \qquad \beta^t c_{i,t}^{-1}\phi_t = \lambda_i q_t^0 \tag{A.5}$$

$$(m^*_{i,t}): \quad \mu^{LTV}_{i,t} + \mu^{PTI}_{i,t} = 1 - \Omega^m_{i,t} - r^*_t \Omega^x_{i,t} \tag{A.6}$$

$$(h^*_{i,t}): \quad p_t = \Omega^h_{i,t} + \mu^{LTV}_{i,t} \theta^{LTV} p_t \tag{A.7}$$

$$(m_{i,t}): \quad q^0_t \Omega^m_{i,t} = \sum_{s_{t+1}|s^t} q^0_{t+1} \pi^{-1} G^{-1} \left[\nu + (1-\nu) \left(\mathbb{I}_{i,t+1} + (1 - \mathbb{I}_{i,t+1}) \Omega^m_{i,t+1} \right) \right] \tag{A.8}$$

$$(x_{i,t}): \quad q^0_t \Omega^x_{i,t} = \sum_{s_{t+1}|s^t} q^0_{t+1} \pi^{-1} G^{-1} \left[(1-\tau) + (1-\nu)(1 - \mathbb{I}_{i,t+1}) \Omega^x_{i,t+1} \right] \tag{A.9}$$

$$\begin{aligned} (h_{i,t}): \quad q^0_t \Omega^h_{i,t} = \sum_{s_{t+1}|s^t} \Big\{ & \lambda_i^{-1} \beta^{t+1} e^h_i (1+\omega) \xi h^{-1}_{i,t} \phi_{t+1} \\ & + q^0_{t+1} \left[-\delta p_{t+1} + \mathbb{I}_{i,t+1} p_{t+1} + (1 - \mathbb{I}_{i,t+1}) \Omega^h_{i,t+1} \right] \Big\}. \end{aligned} \tag{A.10}$$

Beginning with the consumption condition (A.5), note that the same condition applies for an arbitrary agent $i = 0$. As a result, we can substitute and rearrange to obtain

$$c_{i,t} = \left(\frac{\lambda_0}{\lambda_i} \right) c_{0,t}$$

which implies that all borrower consumption allocations are proportional with time-invariant proportions, and without loss of generality we can write $c_{i,t} = C_{B,t}$. A corollary of these constant consumption shares is that all borrower households have the same detrended SDF

$$\hat{\Lambda}_{B,t+1} = \beta G^{-1} \left(\frac{C_{B,t+1}}{C_{B,t}} \right)^{-1}$$

where the G^{-1} term adjusts for trend growth (see Appendix A.2). While $\hat{\Lambda}_{B,t+1}$ is the SDF actually used by the borrower to discount the future, it will also be useful to define the more familiar non-detrended SDF as

$$\Lambda_{B,t+1} = \beta \left(\frac{C_{B,t+1}}{C_{B,t}} \right)^{-1}$$

so that $\hat{\Lambda}_{B,t+1} = G^{-1} \Lambda_{B,t+1}$. As such, I will frequently use $G^{-1} \Lambda_{B,t+1}$ as the borrower's stochastic discount factor in the derivations below. When the G^{-1} term from the SDF does not appear, it means it has canceled with G terms representing trend growth for the time $t+1$ variables inside the discounted expression.

Rearranging the law of motion conditions (A.8) and (A.9), and applying the relation

$$\frac{q^0_{t+1}}{q^0_t} = \Lambda_{B,t+1}\phi(s_{t+1}|s^t)$$

we obtain

$$\Omega^m_{i,t} = \mathbb{E}_t\left\{\Lambda_{B,t+1}\pi^{-1}G^{-1}\left[\nu + (1-\nu)\left(\mathbb{I}_{i,t+1} + (1-\mathbb{I}_{i,t+1})\Omega^m_{i,t+1}\right)\right]\right\}$$
$$\Omega^x_{i,t} = \mathbb{E}_t\left\{\Lambda_{B,t+1}\pi^{-1}G^{-1}\left[(1-\tau) + (1-\nu)(1-\mathbb{I}_{i,t+1})\Omega^x_{i,t+1}\right]\right\}.$$

It is immediate to guess and verify that $\Omega^m_{i,t}$ and $\Omega^x_{i,t}$ do not depend on i, and can be denoted as $\Omega^m_{B,t}$ and $\Omega^x_{B,t}$. Next, applying the law of iterated expectations to replace $\mathbb{I}_{i,t+1}$ with its probability ρ delivers

$$\Omega^m_{B,t} = \mathbb{E}_t\left\{\Lambda_{B,t+1}\pi^{-1}G^{-1}\left[\nu + (1-\nu)\left(\rho + (1-\rho)\Omega^m_{B,t+1}\right)\right]\right\} \tag{A.11}$$
$$\Omega^x_{B,t} = \mathbb{E}_t\left\{\Lambda_{B,t+1}\pi^{-1}G^{-1}\left[(1-\tau) + (1-\nu)(1-\rho)\Omega^x_{B,t+1}\right]\right\}. \tag{A.12}$$

Applying this result to (A.6) yields

$$\mu^{LTV}_{i,t} + \mu^{PTI}_{i,t} = 1 - \Omega^m_{B,t} - r^*_t\Omega^x_{B,t} \equiv \mu_{B,t}. \tag{A.13}$$

Condition (A.13) implies that the sum of the two multipliers is constant across agents, but the exact division across the two multipliers depends on which constraint is binding.

Turning to the housing law of motion (A.10), we can rearrange and apply the law of iterated expectations to obtain

$$\Omega^h_{i,t} = \mathbb{E}_t\left\{\Lambda_{B,t+1}\left[e^h_i(1+\omega)\xi\left(\frac{C_{B,t+1}}{h_{i,t}}\right) - \delta p_{t+1} + \rho p_{t+1} + (1-\rho)\Omega^h_{i,t+1}\right]\right\}.$$

We now guess and verify that $\Omega^h_{i,t}$ has the functional form

$$\Omega^h_{i,t} = \left(\frac{e^h_i}{h_{i,t}}\right)\Psi_{1,t} + \Psi_{2,t}$$
$$\Psi_{1,t} = \mathbb{E}_t\left\{\Lambda_{B,t+1}\left[(1+\omega)\xi C_{B,t+1} + (1-\rho)\Psi_{1,t+1}\right]\right\}$$
$$\Psi_{2,t} = \mathbb{E}_t\left\{\Lambda_{B,t+1}\left[(\rho-\delta)p_{t+1} + (1-\rho)\Psi_{2,t+1}\right]\right\}.$$

Intuitively, $\Psi_{1,t}$ represents the present value of future service flows from housing, evaluated at a normalized level of housing and consumption, while $\Psi_{2,t}$ represents the present value of future cash flows coming from housing transactions and maintenance. As a result, the total marginal value of housing is the present value of $\Psi_{2,t+1}$ plus the present value of $\Psi_{1,t+1}$ adjusted for the actual level of consumption and housing in the borrower's allocation.

To verify this decomposition, we can substitute to obtain

$$\begin{aligned}
\Omega^h_{i,t} &= \left(\frac{e^h_i}{h_{i,t}}\right)\mathbb{E}_t\left\{\Lambda_{B,t+1}\Big[(1+\omega)\xi C_{B,t+1} + (1-\rho)\Psi_{1,t+1}\Big]\right\} \\
&\qquad + \mathbb{E}_t\left\{\Lambda_{B,t+1}\Big[(\rho-\delta)p_{t+1} + (1-\rho)\Psi_{2,t+1}\Big]\right\} \\
&= \mathbb{E}_t\Bigg\{\Lambda_{B,t+1}\left[e^h_i(1+\omega)\xi\left(\frac{C_{B,t+1}}{h_{i,t}}\right) + (1-\rho)\left(\frac{e^h_i}{h_{i,t}}\right)\Psi_{1,t+1}\right] \\
&\qquad\qquad + (\rho-\delta)p_{t+1} + (1-\rho)\Psi_{2,t+1}\Bigg\} \\
&= \mathbb{E}_t\left\{\Lambda_{B,t+1}\left[e^h_i(1+\omega)\xi\left(\frac{C_{B,t+1}}{h_{i,t}}\right) + (\rho-\delta)p_{t+1} + (1-\rho)\Omega^h_{i,t+1}\right]\right\}
\end{aligned}$$

which verifies the guess.

Last, we can substitute into the expression for newly purchased housing $h^*_{i,t}$ (A.7) to obtain the housing demand relation (12), reproduced here as

$$h_{i,t} = \frac{e^h_i \Psi_{1,t}}{(1-\mu^{LTV}_{i,t}\theta^{LTV})p_t - \Psi_{2,t}}. \tag{A.14}$$

This expression shows that housing demand is increasing in demand fundamentals e^h_i, $\Psi_{1,t}$, $\mu^{LTV}_{i,t}\theta^{LTV}$, and $\Psi_{2,t}$, and decreasing in the house price p_t.

The exact solution to (A.14) depends on which of the constraints is binding. If only the LTV constraint is binding, then $\mu^{PTI}_{i,t} = 0$ and so $\mu^{LTV}_{i,t} = \mu_{B,t}$. Conversely, if only the PTI constraint is binding, then $\mu^{LTV}_{i,t} = 0$. If we denote these "only LTV-constrained" and "only PTI-constrained" solutions as $h^{LTV}_{i,t}$ and $h^{PTI}_{i,t}$, then we obtain

$$\begin{aligned}
h^{LTV}_{i,t} &= e^h_i \frac{\Psi_{1,t}}{(1-\mu_{B,t}\theta^{LTV})p_t - \Psi_{2,t}} = e^h_i h^{LTV}_t \\
h^{PTI}_{i,t} &= e^h_i \frac{\Psi_{1,t}}{p_t - \Psi_{2,t}} = e^h_i h^{PTI}_t
\end{aligned}$$

for h_t^{LTV} and h_t^{PTI} as defined in (16) and (18).

In addition to households who are uniquely limited by one binding constraint, there also exists a positive measure of households where both constraints bind simultaneously, as in Justiniano, Primiceri, and Tambalotti (2019). As in the simple example of Section 2, this occurs because the marginal cost of purchasing an additional unit of housing jumps discontinuously when the borrower becomes PTI-constrained. More formally, recall that a borrower is PTI-constrained if

$$\left(\frac{\theta^{PTI}\alpha y_{i,t}}{\hat{r}_t^*}\right) < \theta^{LTV} p_t h_{i,t}. \tag{A.15}$$

The left side of this expression does not depend on the housing choice $h_{i,t}$, but the right side does. As a result, for borrowers with certain values of e_i^h and e_i^y, we can obtain the case

$$\theta^{LTV} p_t e_i^h h_t^{PTI} < \left(\frac{\theta^{PTI}\alpha e_i^y \bar{Y}_B}{\hat{r}_t^*}\right) < \theta^{LTV} p_t e_i^h h_t^{LTV} \tag{A.16}$$

In this case, we can show that both constraints bind by contradiction. If we assume that only the PTI constraint binds then $h_{i,t} = e_i^h h_t^{PTI}$, which from (A.15) and (A.16) implies that the PTI constraint is slack. Conversely, if we assume that only the LTV constraint binds, then $h_{i,t} = e_i^h h_t^{LTV}$ which from (A.15) and (A.16) implies that the LTV constraint is slack. Since both lead to contradictions, and the case that neither constraint binds can be ruled out by $\mu_{B,t} > 0$, it must be the case that both constraints bind. Since equality of the constraints implies

$$\theta^{LTV} p_t h_{i,t} = \left(\frac{\theta^{PTI} e_i^y \alpha \bar{Y}_B}{\hat{r}_t^*}\right) \tag{A.17}$$

we obtain

$$h_{i,t} = \left(\frac{\theta^{PTI}\alpha \bar{Y}_B}{(\hat{r}_t^*)\theta^{LTV} p_t}\right) e_i^y \equiv e_i^y h_t^{both}$$

where h_t^{both} is defined as in (17).

We next determine how the underlying e_i shocks map into the LTV Only, Both, and

PTI Only regimes. Dividing (A.16) by $\theta^{LTV} p_t e_i^h$ and rearranging now yields

$$h_t^{PTI} < h_t^{both} \left(\frac{e_i^y}{e_i^h} \right) < h_t^{LTV}. \tag{A.18}$$

From (5) and (6) we have

$$\begin{aligned} \frac{e_i^y}{e_i^h} &= \exp\left\{ -\frac{1}{2}(1-\gamma)^2\sigma_e^2 - (1-\gamma)\log e_i + \frac{1}{2}\gamma^2\sigma_e^2 - \gamma \log e_i \right\} \\ &= \exp\left\{ -\frac{1}{2}(1-2\gamma)\sigma_e^2 - \log e_i \right\} \\ &= \frac{K}{e_i} \end{aligned}$$

for

$$K = \exp\left\{ -\frac{1}{2}(1-2\gamma)\sigma_e^2 \right\}. \tag{A.19}$$

Substituting, we find that borrowers are jointly constrained if and only if

$$h_t^{PTI} < h_t^{both} \left(\frac{K}{e_i} \right) < h_t^{LTV}. \tag{A.20}$$

which can be manipulated to become

$$K \left(\frac{h_t^{both}}{h_t^{LTV}} \right) < e_i < K \left(\frac{h_t^{both}}{h_t^{PTI}} \right). \tag{A.21}$$

Defining the left value as $\bar{e}_t^{LTV}$ and the right value as $\bar{e}_t^{PTI}$ completes the derivation.

Computing Housing Demand. Evaluating the housing demand condition (21) requires evaluating three integrals. Under our assumption that e_i is lognormally distributed, we can evaluate these integrals in closed form. Note that if $\log e_{i,t} \sim N(0, \sigma_e^2)$, then

$$\begin{aligned} &\int_{k_1}^{k_2} \exp(a + b \log e)\, d\Gamma(e) \\ &\qquad = \exp\left\{ a + \frac{1}{2} b^2 \sigma^2 \right\} \left[\Phi\left(\frac{\log k_2 - b\sigma^2}{\sigma} \right) - \Phi\left(\frac{\log k_1 - b\sigma^2}{\sigma} \right) \right]. \end{aligned} \tag{A.22}$$

From (5) and (6), we have $\log e_i^h = a_h + b_h \log e_i$ and $\log e_i^y = a_y + b_y \log e_i$ for

$$\begin{aligned} a_h &= -\frac{1}{2}\gamma^2\sigma_e^2 & b_h &= \gamma \\ a_y &= -\frac{1}{2}(1-\gamma)^2\sigma_e^2 & b_y &= -(1-\gamma), \end{aligned}$$

which implies

$$\exp\left\{a_h + \frac{1}{2}b_h^2\sigma^2\right\} = \exp\left\{a_y + \frac{1}{2}b_y^2\sigma^2\right\} = 1.$$

Substituting into (A.22), we obtain

$$\begin{aligned} \int^{\bar{e}_t^{LTV}} e_i^h \, d\Gamma(e) &= \Phi\left(\frac{\log \bar{e}_t^{LTV} - \gamma\sigma^2}{\sigma}\right) \equiv F_t^{LTV} \\ \int_{\bar{e}_t^{LTV}}^{\bar{e}_t^{PTI}} e_i^y \, d\Gamma(e) &= \Phi\left(\frac{\log \bar{e}_t^{PTI} + (1-\gamma)\sigma^2}{\sigma}\right) - \Phi\left(\frac{\log \bar{e}_t^{LTV} + (1-\gamma)\sigma^2}{\sigma}\right) \equiv F_t^{both} \\ \int_{\bar{e}_t^{PTI}} e_i^h \, d\Gamma(e) &= 1 - \Phi\left(\frac{\log \bar{e}_t^{PTI} - \gamma\sigma^2}{\sigma}\right) \equiv F_t^{PTI}. \end{aligned}$$

These terms now allow direct evaluation of (21).

A.1.2 Saver Optimality

The saver's optimality condition for debt, similar to that of the borrower, is given by

$$1 = \Omega_{S,t}^m + \Omega_{S,t}^x (r_t^* - \eta).$$

where $\Omega_{S,t}^m$ and $\Omega_{S,t}^x$ are the marginal continuation *benefits* to the saver of an additional unit of face value and an additional dollar of promised initial payments, respectively. The values of $\Omega_{S,t}^m$ and $\Omega_{S,t}^x$ are defined by

$$\Omega_{S,t}^m = \mathbb{E}_t\left\{\Lambda_{S,t+1}\pi^{-1}\Big[\nu + (1-\nu)\rho + (1-\nu)(1-\rho)\Omega_{S,t+1}^m\Big]\right\} \tag{A.23}$$

$$\Omega_{S,t}^x = \mathbb{E}_t\left\{\Lambda_{S,t+1}\pi^{-1}\Big[1 + (1-\nu)(1-\rho)\Omega_{S,t+1}^x\Big]\right\}. \tag{A.24}$$

Overall, the saver's optimality conditions are equivalent to the terms in the borrower's problem, with the following exceptions: savers are unconstrained ($\mu = 0$), use a different stochastic discount factor, do not optimize over housing, and face a proportional tax

(wedge) on their mortgage payment receipts.

A.2 Detrending

In this section, I derive the detrending adjustments used above to transform from a growing economy to a stationary, detrended economy. Let $Y_t^g = A_t \bar{Y}$, with $A_t = GA_{t-1}$, and let superscript g indicate growing-economy variables, so that e.g., $M_t^g = A_t M_t$.

To begin, I detrend the laws of motion for the mortgage state variables. In the growing economy, these laws of motion are

$$\begin{aligned} M_t^g &= \rho M_t^{g,*} + (1-\rho)\pi^{-1} M_{t-1}^g \\ X_t^g &= \rho r_t^* M_t^{g,*} + (1-\rho)\pi^{-1} X_{t-1}^g. \end{aligned}$$

These expressions reflect that a fraction ρ of mortgages prepay, while the remaining mortgages are deflated by π^{-1}, as these state variables are measured in real terms while the mortgages are specified in nominal currency. Substituting to express these variables in detrended terms (e.g., $M_t^g = A_t M_t$), we obtain

$$\begin{aligned} A_t M_t &= \rho A_t M_t^* + (1-\rho)\pi^{-1} A_{t-1} M_{t-1} \\ A_t X_t &= \rho r_t^* A_t M_t^* + (1-\rho)\pi^{-1} A_{t-1} X_{t-1}. \end{aligned}$$

Dividing through by A_t now yields

$$\begin{aligned} M_t &= \rho M_t^* + (1-\rho)\pi^{-1} G^{-1} M_{t-1} \\ X_t &= \rho r_t^* M_t^* + (1-\rho)\pi^{-1} G^{-1} X_{t-1}, \end{aligned}$$

which agree exactly with (23) and (24).

Turning to the policy functions, we need to check or adjust six functions — $\Psi_{1,t}$, $\Psi_{2,t}$, $\Omega_{B,t}^m$, $\Omega_{B,t}^x$, $\Omega_{S,t}^m$, and $\Omega_{S,t}^x$ — after which the remaining variables (such as p_t) can be derived from these policy functions using the equations in the main text. To begin, note that

$$\Lambda_{B,t+1}^g = \beta_B \left(\frac{C_{B,t+1}^g}{C_{B,t}^g} \right)^{-1} = \beta_B \left(\frac{A_{t+1} C_{B,t+1}}{A_t C_{B,t}} \right)^{-1} = \beta_B G^{-1} \left(\frac{C_{B,t+1}}{C_{B,t}} \right)^{-1} = G^{-1} \Lambda_{B,t+1}$$

for

$$\Lambda_{B,t+1} \equiv \beta_B \left(\frac{C_{B,t+1}}{C_{B,t}} \right)^{-1},$$

which motivates our use of the detrended SDF in the derivations above.

For $\Psi_{1,t}$, the expression in the growing economy is

$$\Psi^g_{1,t} = \mathbb{E}_t \left\{ \Lambda^g_{B,t+1} \Big[(1+\omega)\xi C^g_{B,t+1} + (1-\rho)\Psi^g_{1,t+1} \Big] \right\}.$$

Substituting, we obtain

$$\begin{aligned} A_t \Psi_{1,t} &= \mathbb{E}_t \left\{ G^{-1} \Lambda_{B,t+1} \Big[(1+\omega)\xi A_{t+1} C_{B,t+1} + (1-\rho) A_{t+1} \Psi_{1,t+1} \Big] \right\} \\ &= \mathbb{E}_t \left\{ G^{-1} \Lambda_{B,t+1} A_{t+1} \Big[(1+\omega)\xi C_{B,t+1} + (1-\rho) \Psi_{1,t+1} \Big] \right\}. \end{aligned}$$

Dividing through by A_t and applying $A_{t+1}/A_t = G$ yields (13). Similarly, the expression for $\Psi_{2,t}$ in the growing economy is

$$\Psi^g_{2,t} = \mathbb{E}_t \left\{ \Lambda^g_{B,t+1} \Big[(\rho-\delta) p^g_{t+1} + (1-\rho)\Psi^g_{2,t+1} \Big] \right\}.$$

Substituting, we obtain

$$\begin{aligned} A_t \Psi_{2,t} &= \mathbb{E}_t \left\{ G^{-1} \Lambda_{B,t+1} \Big[(\rho-\delta) A_{t+1} p_{t+1} + (1-\rho) A_{t+1} \Psi_{2,t+1} \Big] \right\} \\ &= \mathbb{E}_t \left\{ G^{-1} \Lambda_{B,t+1} A_{t+1} \Big[(\rho-\delta) p_{t+1} + (1-\rho) \Psi_{2,t+1} \Big] \right\}. \end{aligned}$$

Dividing through by A_t and applying $A_{t+1}/A_t = G$ yields (14).

For $\Omega^m_{B,t}$ and $\Omega^x_{B,t}$, these quantities are already scaled per dollar of mortgage debt, and do not grow with trend output. However, we must still adjust for the effect of growth on the borrower SDF. In the growing economy, we have

$$\begin{aligned} \Omega^m_{B,t} &= \mathbb{E}_t \left\{ \Lambda^g_{B,t+1} \pi^{-1} \Big[\nu + (1-\nu) \Big(\rho + (1-\rho)\Omega^m_{B,t+1} \Big) \Big] \right\} \\ \Omega^x_{B,t} &= \mathbb{E}_t \left\{ \Lambda^g_{B,t+1} \pi^{-1} \Big[(1-\tau) + (1-\nu)(1-\rho)\Omega^x_{B,t+1} \Big] \right\}. \end{aligned}$$

Substituting for $\Lambda^g_{B,t+1}$, we obtain (A.11) and (A.12). Last, we can similarly substitute for

the saver SDF in the growing economy expressions

$$\Omega^m_{S,t} = \mathbb{E}_t \left\{ \Lambda^g_{S,t+1} \pi^{-1} \left[\nu + (1-\nu)\rho + (1-\nu)(1-\rho)\Omega^m_{S,t+1} \right] \right\}$$
$$\Omega^x_{S,t} = \mathbb{E}_t \left\{ \Lambda^g_{S,t+1} \pi^{-1} \left[1 + (1-\nu)(1-\rho)\Omega^x_{S,t+1} \right] \right\}.$$

Since the saver's risk-neutral preferences imply $\Lambda^g_{S,t+1} = \Lambda_{S,t+1} = \beta_S$, we immediately obtain (A.23) and (A.24).

A.3 Policy Functions

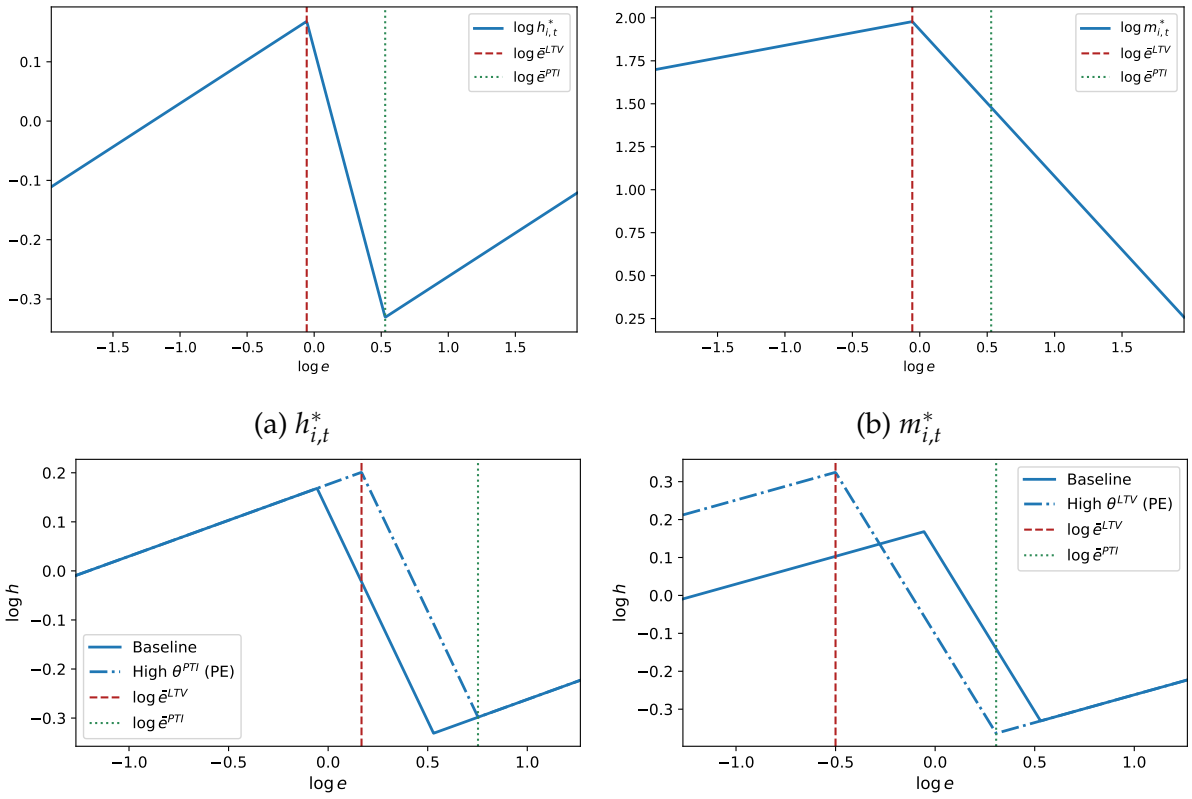

(a) $h^*_{i,t}$

(b) $m^*_{i,t}$

(c) $\theta^{PTI} \uparrow$ (Partial Eqm.)

(d) $\theta^{LTV} \uparrow$ (Partial Eqm.)

Figure A.1: Model Policy Functions

Note: The figure displays model policy functions a function of the idiosyncratic state $\log e_i$ under the baseline calibration (solid line) and alternative credit standards (dash-dotted line). Panel (a) displays the optimal choice of $\log h^*_{i,t}$ in steady state. Panel (b) displays the equilibrium credit allocation $\log m^*_{i,t}$ in steady state. Panel (c) displays the effect of a partial equilibrium increase in the PTI limit ($\theta^{PTI} \uparrow$), while Panel (d) displays the effect of a partial equilibrium increase in the LTV limit ($\theta^{LTV} \uparrow$). The vertical dashed and dotted lines indicate the thresholds $\log \bar{e}^{LTV}$ and $\log \bar{e}^{PTI}$, respectively, under the alternative parameterizations.

For intuition, a numerical example of the optimal policy functions can be visualized

in Figure A.1. To begin, Panel (a) displays the optimal housing allocation $\log h^*_{i,t}$ with respect to $\log e_i$ in the steady state. The red dashed vertical line reflects the transition from being LTV-constrained only to being jointly constrained by both limits, while the green dotted line reflects the transition from being jointly constrained to being PTI-constrained only. In the "LTV Only" and "PTI Only" regions, housing demand is proportional to the housing preference shifter e^h_i, and hence increasing in e_i. However, in the "Both" region, housing demand is proportional to income e^y_i, and hence decreasing in e_i. The parameter γ determines the overall tilt or slope of this curve.

Panel (b) displays the resulting mortgage allocations $\log m^*_{i,t}$, along with vertical lines representing the same transitions. In the "LTV Only" region, the borrower is LTV-constrained, and the debt limit is proportional to the house value, so that $\log m^*_{i,t}$ and $\log h^*_{i,t}$ increase in parallel. In the "Both" region, the borrower is jointly limited by both constraints. This can occur since the borrower's PTI limit is shrinking as his or her income falls over this region, while the LTV limit declines at exactly the same rate, as the optimal house size is proportional to income in this region. Last, in the "PTI Only" region, we see the debt limit continue to tighten, as further declines in the borrower's income reduce the PTI limit further. This occurs despite the rising house value, as the LTV constraint is slack in this region, making additional collateral ineffective at the margin.

Next, Panel (c) shows the partial equilibrium effect of a relaxation of the PTI limit, prior to the response of prices. In this example, this is achieved through an increase in the parameter θ^{PTI}, but could also occur endogenously through a decrease in interest rates or an increase in income. From equation (17), we can see that h^{both} is directly affected by a relaxation of the PTI limit, and increases with θ^{PTI}, as in our simple example in Figure 1, reflected by the upward shift in housing demand in the center of the plot. In addition, both of the cutoff points between regions move to the right, implying an increase in the share of the population in the LTV-only region, and decrease in the share in the PTI-only region. Since this change has strictly increased demand with no change to supply, the house price will subsequently rise to restore market clearing in (21).

Last, Panel (d) shows a similar experiment expanding the LTV limit, again in partial equilibrium, through an increase in θ^{LTV}. From (17) we see that h^{both} decreases, reflecting reduced demand by jointly constrained households, as in Figure 1. Unlike in the simple example, however, h^{LTV} also increases, leading to an upward shift in demand on the left portion of the plot from the borrowers constrained by LTV only. The net effect depends on the balance of the two, and could lead to an increase or decrease in prices in general

equilibrium.

A.4 Deriving the Approximate House Price Formula

This section derives the approximation (25). To begin, rearrange (12) to obtain

$$\left(1 - \mu_{i,t}^{LTV}\theta^{LTV}\right) p_t = \Psi_{1,t}\left(\frac{e_i^h}{h_{i,t}^*}\right) + \Psi_{2,t}.$$

Integrating with respect to e_i now yields

$$\left(1 - \theta^{LTV}\int \mu_{i,t}^{LTV}\, d\Gamma(e)\right) p_t = \Psi_{1,t}\int \frac{e_i^h}{h_{i,t}^*}\, d\Gamma(e) + \Psi_{2,t}. \tag{A.25}$$

I now seek to obtain approximations for both of these integral terms.

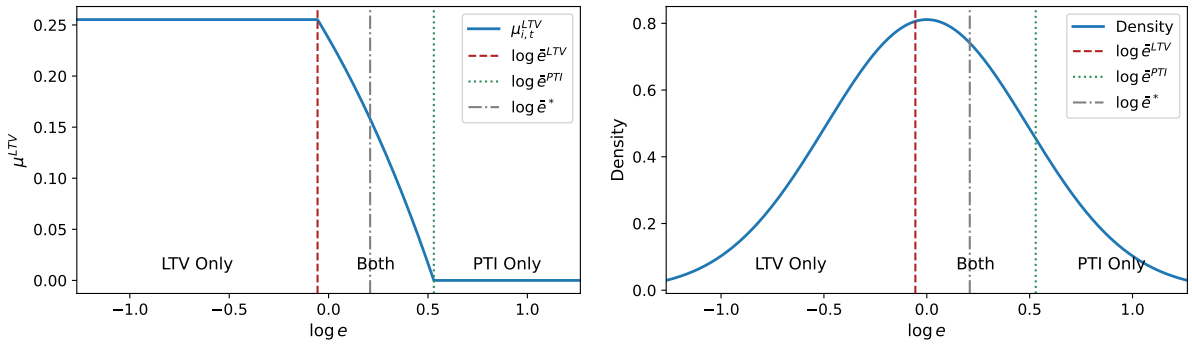

(a) μ^{LTV} (b) Density of $\log e$

Figure A.2: Approximate Pricing Equation, Intuition

Note: Panel (a) plots the LTV multiplier $\mu_{i,t}^{LTV}$ as a function of the household heterogeneity draw $\log e_i$, evaluated at the baseline steady state. The dashed vertical line marks $\log \bar{e}^{LTV}$, the dotted line marks $\log \bar{e}^{PTI}$, and the dash-dotted line marks $\log \bar{e}^{R}$, the threshold used in the ex-ante heterogeneity approximation. Panel (b) plots the density of $\log e_i$ under the same baseline calibration, with the same three thresholds indicated, showing the distribution of the borrower population across the three constraint regimes.

For the integral of $\mu_{i,t}^{LTV}$, recall that $\mu_{i,t}^{LTV}$ is equal to μ_t if a household is LTV-constrained only, to zero if a household is PTI-constrained only, and decreases continuously from μ_t to 0 in the region constrained by both. This can be seen in Appendix Figure A.2(a), which

plots $\mu_{i,t}^{LTV}$ with respect to e_i. Thus, this integral is equal to

$$\begin{aligned}\int \mu_{i,t}^{LTV}\, d\Gamma(e) &= \underbrace{\int^{\bar{e}_t^{LTV}} \mu_{i,t}^{LTV}\, d\Gamma(e)}_{\text{LTV Only}} + \underbrace{\int_{\bar{e}_t^{LTV}}^{\bar{e}_t^{PTI}} \mu_{i,t}^{LTV}\, d\Gamma(e)}_{\text{Both}} + \underbrace{\int_{\bar{e}_t^{PTI}} \mu_{i,t}^{LTV}\, d\Gamma(e)}_{\text{PTI Only}} \\ &= \underbrace{\int^{\bar{e}_t^{LTV}} \mu_t\, d\Gamma(e)}_{\text{LTV Only}} + \underbrace{\int_{\bar{e}_t^{LTV}}^{\bar{e}_t^{PTI}} \mu_{i,t}^{LTV}\, d\Gamma(e)}_{\text{Both}} + \underbrace{\int_{\bar{e}_t^{PTI}} 0\, d\Gamma(e)}_{\text{PTI Only}} \\ &= F_t^{LTV}\mu_t + \int_{\bar{e}_t^{LTV}}^{\bar{e}_t^{PTI}} \mu_{i,t}^{LTV}\, d\Gamma(e).\end{aligned}$$

where $F_t^{LTV} = \Gamma(\bar{e}_t^{LTV})$ is the share of borrowers in the LTV-only region. Since $\mu_{i,t}^{LTV}$ varies from μ_t to 0 over the both-constrained region, this last integral can be approximated using the trapezoid rule to obtain

$$\begin{aligned}\int \mu_{i,t}^{LTV}\, d\Gamma(e) &\simeq F_t^{LTV}\mu_t + \frac{1}{2}(\mu_t + 0)\Big(\Gamma(\bar{e}_t^{PTI}) - \Gamma(\bar{e}_t^{LTV})\Big) \\ &= \left(F_t^{LTV} + \frac{1}{2}F_t^{both}\right)\mu_t \end{aligned} \tag{A.26}$$

where $F_t^{both} \equiv \Gamma(\bar{e}_t^{PTI}) - \Gamma(\bar{e}_t^{LTV})$ is the share of purchasers constrained by both limits.

For the integral of $e_i^h / h_{i,t}^*$, we can take a second-order Taylor approximation of this function around $e_i^h = 1$ and $h_{i,t}^* = \bar{H}_B$ to yield

$$\frac{e_i^h}{h_{i,t}^*} \simeq \frac{1}{\bar{H}_B} + \frac{1}{\bar{H}_B}(e_i^h - 1) - \frac{1}{\bar{H}_B^2}(h_{i,t}^* - \bar{H}_B) - \frac{1}{\bar{H}_B^2}(e_i^h - 1)(h_{i,t}^* - \bar{H}_B) + \frac{1}{\bar{H}_B^3}(h_{i,t}^* - \bar{H}_B)^2.$$

Taking the expectation with respect to e_i, we have $\mathbb{E}[e_i^h] = 1$ and $\mathbb{E}[h_{i,t}^*] = \bar{H}_B$, and so

$$\mathbb{E}\left(\frac{e_i^h}{h_{i,t}^*}\right) \simeq \frac{1}{\bar{H}_B} - \frac{\mathrm{Cov}(e_i^h, h_{i,t}^*)}{\bar{H}_B^2} + \frac{\mathrm{Var}(h_{i,t}^*)}{\bar{H}_B^3}. \tag{A.27}$$

Note that since the first-order terms are already zero, the approximation already holds to first order. However, we can improve by analyzing the second-order terms. In the LTV-only and PTI-only regimes in (15), housing is proportional to the preference shifter e_i^h, so that a useful approximation is $h_{i,t}^* \simeq \bar{H}_B e_i^h$ (ignoring the both-constrained region where

$h^*_{i,t}$ is proportional to e^y_i). Under this approximation,

$$\text{Cov}(e^h_i, h^*_{i,t}) \simeq \bar{H}_B \text{Var}(e^h_i), \qquad \text{Var}(h^*_{i,t}) \simeq \bar{H}^2_B \text{Var}(e^h_i),$$

implying

$$\frac{\text{Cov}(e^h_i, h^*_{i,t})}{\bar{H}^2_B} \simeq \frac{\text{Var}(e^h_i)}{\bar{H}_B}$$
$$\frac{\text{Var}(h^*_{i,t})}{\bar{H}^3_B} \simeq \frac{\text{Var}(e^h_i)}{\bar{H}_B}.$$

As a result, the final two terms of (A.27) largely cancel (and would cancel exactly if $h^*_{i,t} = e^h_i \bar{H}_B$, which occurs in the limit as $\theta^{PTI} \to \infty$), leaving

$$\mathbb{E}_t \left(\frac{e^h_i}{h^*_{i,t}} \right) \simeq \frac{1}{\bar{H}_B}. \tag{A.28}$$

With both integrals approximated, we can now substitute from (A.28) and (A.26) into (A.25) to obtain

$$(1 - \mathcal{C}_t)\, p_t = \Psi_{1,t} \bar{H}_B^{-1} + \Psi_{2,t}. \tag{A.29}$$

for $\mathcal{C}_t = \mu_t \hat{F}^{LTV}_t \theta^{LTV}$. Substituting in for $\Psi_{1,t}$ and $\Psi_{2,t}$ now yields

$$\begin{aligned}
(1 - \mathcal{C}_t)\, p_t &= \mathbb{E}_t \Big\{ \Lambda_{B,t+1} \Big[(1+\omega)\xi C_{B,t+1}/\bar{H}_B + (1-\rho)\Psi_{1,t+1}/\bar{H}_B + (\rho - \delta) p_{t+1} + (1-\rho)\Psi_{2,t+1} \Big] \Big\} \\
&= \mathbb{E}_t \Big\{ \Lambda_{B,t+1} \Big[(1+\omega)\xi C_{B,t+1}/\bar{H}_B + (\rho - \delta) p_{t+1} + (1-\rho) \Big(\Psi_{1,t+1}/\bar{H}_B + \Psi_{2,t+1} \Big) \Big] \Big\} \\
&= \mathbb{E}_t \Big\{ \Lambda_{B,t+1} \Big[(1+\omega)\xi C_{B,t+1}/\bar{H}_B + (\rho - \delta) p_{t+1} + (1-\rho)\,(1 - \mathcal{C}_{t+1})\, p_{t+1} \Big] \Big\} \\
&= \mathbb{E}_t \Big\{ \Lambda_{B,t+1} \Big[(1+\omega)\xi C_{B,t+1}/\bar{H}_B + (1 - \delta - (1-\rho)\mathcal{C}_{t+1}) p_{t+1} \Big] \Big\}
\end{aligned}$$

where the third step follows from (A.29), completing the derivation of (25).

A.5 Simplified Model with Ex-Post Heterogeneity

In this section, I solve a simplified version of the model with ex-post heterogeneity, as developed in the original version of this paper, and show that it satisfies an expression

very close to, and well approximated by, equation (12).

In this version of the problem, each active borrower must choose a housing allocation $h^*_{i,t}$ prior to observing his or her realization of e_i. I further assume that $\gamma = 0$, so that

$$\log e^y_i = -\frac{1}{2}\sigma^2_e - \log e_i$$

and all heterogeneity occurs along the income dimension. In this case, all borrowers will choose a symmetric housing allocation $h^*_{i,t} = H^*_t$, and the problem aggregates to that of a representative borrower who maximizes

$$V_{B,0} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t_B \Big[\log c_{B,t} + (1+\omega)\xi \log h_{B,t-1}\Big]$$

subject to the budget constraint

$$c_{B,t} \leq \underbrace{(1-\tau)\bar{Y}_B}_{\text{labor income}} - \underbrace{\pi^{-1}G^{-1}\big((1-\tau)X_{t-1} + \nu M_{t-1}\big)}_{\text{mortgage payment}} - \underbrace{\delta p_t H_{t-1}}_{\text{maintenance}}$$
$$- \rho \underbrace{p_t(H^*_t - H_{t-1})}_{\text{new housing}} + \rho \underbrace{\big(M^*_t - (1-\nu)\pi^{-1}G^{-1}M_{t-1}\big)}_{\text{new debt}}.$$

and the credit constraint

$$M^*_t \leq \int_i \min\left\{\theta^{LTV} p_t e^h_i H^*_t,\ \frac{\theta^{PTI} e^y_i \alpha \bar{Y}_B}{\hat{r}^*_t}\right\} d\Gamma(e). \tag{A.30}$$

Manipulating (A.30), the LTV limit will bind if and only if

$$\begin{aligned}
\theta^{LTV} p_t e^h_i H^*_t &\leq \frac{\theta^{PTI} e^y_i \alpha \bar{Y}_B}{\hat{r}^*_t} \\
\frac{e^h_i}{e^y_i} &\leq \frac{\theta^{PTI}\alpha \bar{Y}_B}{(\hat{r}^*_t)\theta^{LTV} p_t H^*_t} \\
e_i &\leq K\left(\frac{\theta^{PTI}\alpha \bar{Y}_B}{(\hat{r}^*_t)\theta^{LTV} p_t H^*_t}\right)
\end{aligned}$$

where the last line follows from $e^h_i/e^y_i = e_i/K$. Note that this implies a single threshold value for e_i to determine whether a household is LTV-constrained or PTI-constrained. Thus, in this simpler model with only ex-post heterogeneity we do not observe a positive mass of jointly-constrained households.

If we define

$$\bar{e}_t^* = K\left(\frac{\theta^{PTI}\alpha\bar{Y}_B}{(\hat{r}_t^*)\theta^{LTV}p_tH_t^*}\right) \tag{A.31}$$

then we can rewrite (A.30) as

$$M_t^* \leq \underbrace{\theta^{LTV}p_tH_t^*F_t^{LTV}}_{\text{LTV constrained}} + \underbrace{\frac{\theta^{PTI}\alpha\bar{Y}_B}{\hat{r}_t^*}\int_{\bar{e}_t^*} e_i^y\, d\Gamma(e)}_{\text{PTI constrained}} \tag{A.32}$$

where $F_t^{LTV} = \Gamma(\bar{e}_t^*)$ is the share of borrowers constrained by LTV. Taking the first order condition with respect to H_t^* now yields

$$0 = -p_t + \Omega_{B,t}^h + \mu_t F_t^{LTV}\theta^{LTV}p_t \tag{A.33}$$

where μ_t is the multiplier on (A.32), normalized by ρ, and $\Omega_{B,t}^h$ is the expected marginal value of carrying an additional unit of housing into the next period, defined by

$$\Omega_{B,t}^h = \mathbb{E}_t\left\{\Lambda_{B,t+1}\left[(1+\omega)u_{B,t+1}^h/u_{B,t+1}^c + (\rho-\delta)p_{t+1} + (1-\rho)\Omega_{B,t+1}^h\right]\right\}. \tag{A.34}$$

Rewriting (A.33) as

$$\Omega_t^h = (1-\tilde{\mathcal{C}}_t)p_t$$

for $\tilde{\mathcal{C}}_t = \mu_t F_t^{LTV}\theta^{LTV}$ and substituting into (A.34) and rearranging immediately delivers

$$p_t = \frac{\mathbb{E}_t\left\{\Lambda_{B,t+1}\left[(1+\omega)u_{B,t+1}^h/u_{B,t+1}^c + \left(1-\delta-(1-\rho)\tilde{\mathcal{C}}_{t+1}\right)p_{t+1}\right]\right\}}{1-\tilde{\mathcal{C}}_t} \tag{A.35}$$

This expression is identical to (25) with the exception that we have replaced $\mathcal{C}_t$ with $\tilde{\mathcal{C}}_t$. It is also straightforward to verify that the optimality condition with respect to M_t^* delivers an identical solution for $\Omega_{B,t}^m$, $\Omega_{B,t}^x$, and μ_t. Hence, the last step is to show that F_t^{LTV} in the simplified model with ex-post heterogeneity is close to $\hat{F}_t^{LTV}$ in the full model, in which case $\mathcal{C}_t \simeq \tilde{\mathcal{C}}_t$, and so (25) and (A.35) will align.

To do so, note that (A.31) implies $\bar{e}_t^* = K(h_t^{both}/\bar{H}_B)$, where h_t^{both} is defined as in the baseline model. Further, by running the steps of the derivation of (25) in reverse, after

applying our integral approximations, it is straightforward to verify that

$$p_t \simeq \frac{\Psi_{1,t}\bar{H}_B^{-1} + \Psi_{2,t}}{1 - \mu_t \hat{F}^{LTV}\theta^{LTV}}$$

in the baseline model, which can be rearranged to imply

$$\bar{H}_B \simeq \frac{\Psi_{1,t}}{(1 - \mu_t \hat{F}_t^{LTV}\theta^{LTV})p_t - \Psi_{2,t}} \tag{A.36}$$

Since $0 < \hat{F}_t^{LTV} < 1$, it is clear from comparing (A.36) to (16) and (18) that $h_t^{PTI} < \bar{H}_B < h_t^{LTV}$. Applying this result to (19) and (20) immediately yields

$$\bar{e}_t^{LTV} < \bar{e}_t^* < \bar{e}_t^{PTI}. \tag{A.37}$$

From (A.37), it is clear that, holding the distribution of e_i fixed,

$$\Gamma(\bar{e}_t^{LTV}) < \Gamma(\bar{e}_t^*) < \Gamma(\bar{e}_t^{PTI}).$$

The left hand side of this inequality is F_t^{LTV} in the baseline model, while the right hand side is $F_t^{LTV} + F_t^{both}$ in the baseline model. The center term is F_t^{LTV} in the simple model. Hence, F^{LTV} in the simplified model should fall between F_t^{LTV} and $F_t^{LTV} + F_t^{both}$ in the baseline model, which at the midpoint would equal $\hat{F}_t^{LTV}$. For a quantitative measure, Figure A.2 shows that under this paper's calibration, $\log \bar{e}_t^*$ is extremely close to halfway between $\log \bar{e}_t^{LTV}$ and $\log \bar{e}_t^{PTI}$, and the areas under the curve between $\log \bar{e}_t^{LTV}$ and $\log \bar{e}_t^*$ and between $\log \bar{e}_t^*$ and $\log \bar{e}_t^{PTI}$ are similar.

To summarize, the results in this section show that simpler models with ex-post heterogeneity should deliver model behavior, and in particular house price dynamics, that are very similar to those of this paper's baseline model. As a result, these simpler models, or their extensions in e.g., Greenwald and Guren (2025) or Allen and Greenwald (2022), should deliver accurate results despite lacking jointly-constrained households.

A.6 Steady State Elasticities

To understand the source of the model's elasticities, we can consider the steady state housing demand condition

$$H_B = \underbrace{\int^{\bar{e}^{LTV}} e^h(e_i)\, d\Gamma(e)}_{w_H^{LTV}} h^{LTV} + \underbrace{\int_{\bar{e}^{LTV}}^{\bar{e}^{PTI}} e^y(e_i)\, d\Gamma(e)}_{w_H^{both}} h^{both} + \underbrace{\int_{\bar{e}^{PTI}} e^h(e_i)\, d\Gamma(e)}_{w_H^{PTI}} h^{PTI}$$

where w_H^{LTV}, w_H^{both}, and w_H^{PTI} are the weights that we apply to h^{LTV}, h^{both}, and h^{PTI}, respectively, when aggregating housing demand. First, note that if we differentiate both sides with respect to any variable z, we will obtain

$$\begin{aligned} H_B \frac{\partial \log H_B}{\partial z} = {} & e^h(\bar{e}^{LTV}) f_e(\bar{e}^{LTV}) h^{LTV} \frac{\partial \bar{e}^{LTV}}{\partial z} + w_H^{LTV} h^{LTV} \frac{\partial \log h^{LTV}}{\partial z} \\ & + \left[e^y(\bar{e}^{PTI}) f_e(\bar{e}^{PTI}) \frac{\partial \bar{e}^{PTI}}{\partial z} - e^y(\bar{e}^{LTV}) f_e(\bar{e}^{LTV}) \frac{\partial \bar{e}^{LTV}}{\partial z} \right] h^{both} \\ & + w_H^{both} h^{both} \frac{\partial \log h^{both}}{\partial z} \\ & - e^h(\bar{e}^{PTI}) f_e(\bar{e}^{PTI}) h^{PTI} \frac{\partial \bar{e}^{PTI}}{\partial z} + w_H^{PTI} h^{PTI} \frac{\partial \log h^{PTI}}{\partial z}. \end{aligned}$$

But since

$$\begin{aligned} e^h(\bar{e}^{LTV}) h^{LTV} &= e^y(\bar{e}^{LTV}) h^{both} \\ e^h(\bar{e}^{PTI}) h^{PTI} &= e^y(\bar{e}^{PTI}) h^{both} \end{aligned}$$

all the boundary terms involving the derivatives of $\bar{e}^{LTV}$ or $\bar{e}^{PTI}$ with respect to z cancel, and we are left with

$$H_B \frac{\partial \log H_B}{\partial z} = w_H^{LTV} h^{LTV} \frac{\partial \log h^{LTV}}{\partial z} + w_H^{both} h^{both} \frac{\partial \log h^{both}}{\partial z} + w_H^{PTI} h^{PTI} \frac{\partial \log h^{PTI}}{\partial z}.$$

Dividing through by H_B now yields

$$\frac{\partial \log H_B}{\partial z} = s_H^{LTV} \frac{\partial \log h^{LTV}}{\partial z} + s_H^{both} \frac{\partial \log h^{both}}{\partial z} + s_H^{PTI} \frac{\partial \log h^{PTI}}{\partial z}$$

where e.g., $s_H^{LTV} = w_H^{LTV} h^{LTV} / H_B$ is the steady state share of housing purchased by the sector of the population constrained by LTV only.

To apply this formula to create a log-linear approximation, we can compute the derivatives of the h terms with respect to different variables. To begin, we have

$$h^{LTV} = \frac{\Psi_1}{(1-\mu\theta^{LTV})p - \Psi_2}.$$

But since in steady state we have

$$\Psi_2 = \beta(\rho-\delta)p + \beta(1-\rho)\Psi_2$$

we can solve to obtain

$$\Psi_2 = \left(\frac{\beta(\rho-\delta)}{1-\beta(1-\rho)}\right)p = \psi p$$

for appropriately defined ψ. Substituting, we obtain

$$h^{LTV} = \frac{\Psi_1}{(1-\mu\theta^{LTV}-\psi)p}$$

so that

$$\log h^{LTV} = \log \Psi_1 - \log\left(1-\mu\theta^{LTV}-\psi\right) - \log p.$$

From (13), it is also straightforward to show the steady state solution

$$\Psi_1 = \frac{\beta(1+\omega)\xi C_B}{1-(1-\rho)\beta}$$

and hence that

$$\log \Psi_1 = \log(1+\omega) + \log(\beta\xi C_B) - \log(1-(1-\rho)\beta).$$

Similarly, since

$$h^{PTI} = \frac{\Psi_1}{p - \Psi_2}$$

we can perform symmetric steps to obtain

$$\log h^{PTI} = \log \Psi_1 - \log(1-\psi) - \log p.$$

Last, we have

$$h^{both} = \frac{\theta^{PTI} Y_B}{\hat{r}^* \theta^{LTV} p}.$$

Taking logs, we obtain

$$\log h^{both} = \log \theta^{PTI} + \log Y_B - \log \hat{r}^* - \log \theta^{LTV} - \log p.$$

Substituting into the formula above, and applying the fixed housing supply assumption $\partial \log H_B / \partial z = 0$ for all z, we can approximate around $\log \theta^{LTV}, \log \theta^{PTI}, \log \hat{r}^*$, and ω to obtain

$$\begin{aligned}
0 = s_H^{LTV} &\left[\left(\frac{\mu \theta^{LTV}}{1 - \mu \theta^{LTV} - \psi} \right) d \log \theta^{LTV} - d \log p + d\omega \right] \\
&+ s_H^{both} \left[d \log \theta^{PTI} - d \log \hat{r}^* - d \log \theta^{LTV} - d \log p \right] \\
&+ s_H^{PTI} \left(-d \log p + d\omega \right) \\
d \log p = s_H^{both} &\left[d \log \theta^{PTI} - d \log \hat{r}^* \right] + (1 - s_H^{both}) d\omega \\
&+ \left\{ s_H^{LTV} \left(\frac{\mu \theta^{LTV}}{1 - \mu \theta^{LTV} - \psi} \right) - s_H^{both} \right\} d \log \theta^{LTV}
\end{aligned}$$

Note that the terms including $d\omega$ follow from $\Psi_1 \propto (1 + \omega)$ and $\log(1 + \omega) \simeq \omega$ around the steady state value $\omega = 0$. This expression shows that in steady state, the elasticity of price to the PTI cap θ^{PTI} is s_H^{both}, while the elasticity of price to the effective payment rate $\hat{r}^*$ is $-s_H^{both}$. The elasticity with respect to the LTV cap θ^{LTV} is a more complex expression that balances out the positive effect on housing demand from loosening LTV limits for LTV-constrained households against the negative effect on housing demand from loosening LTV limits for households jointly limited by both constraints.

I next approximate the steady state log debt limit M^*. This value can be expressed as

$$M^* = \theta^{LTV} p h^{LTV} \underbrace{\int^{\bar{e}^{LTV}} e^h(e_i)\, d\Gamma(e)}_{w_M^{LTV}} \quad + \quad \frac{\theta^{PTI} Y_B}{\hat{r}^*} \underbrace{\int_{\bar{e}^{LTV}} e^y(e_i)\, d\Gamma(e)}_{w_M^{both} + w_M^{PTI}}.$$

where w_M^{both} and w_M^{PTI} can be separately defined as

$$w_M^{both} = \int_{\bar{e}^{LTV}}^{\bar{e}^{PTI}} e^y(e_i)\, d\Gamma(e)$$

$$w_M^{PTI} = \int_{\bar{e}^{PTI}} e^y(e_i)\, d\Gamma(e).$$

Applying a symmetric log-linearization to the one for housing demand above now yields

$$\begin{aligned} M^* d\log M^* = {} & w_M^{LTV}\theta^{LTV} p h^{LTV} (d\log\theta^{LTV} + d\log p + d\log h^{LTV}) \\ & + (w_M^{both} + w_M^{PTI}) \left(\frac{\theta^{PTI}\Upsilon_B}{\hat{r}^*}\right) \left(d\log\theta^{PTI} - d\log\hat{r}^*\right) \end{aligned}$$

Dividing through by M^*, and defining

$$\begin{aligned} s_M^{LTV} &= \frac{w_M^{LTV}\theta^{LTV} p h^{LTV}}{M^*} \\ s_M^{both} &= \frac{w_M^{both}\theta^{PTI}\Upsilon_B}{\hat{r}^* M^*} \\ s_M^{PTI} &= \frac{w_M^{PTI}\theta^{PTI}\Upsilon_B}{\hat{r}^* M^*} \end{aligned}$$

to be the shares of new debt allocated to new borrowers in the LTV Only, Both, and PTI Only regions, we obtain

$$\begin{aligned} d\log M^* = {} & s_M^{LTV} \left(d\log\theta^{LTV} + d\log p + d\log h^{LTV}\right) \\ & + (s_M^{both} + s_M^{PTI}) \left(d\log\theta^{PTI} - d\log\hat{r}^*\right). \end{aligned}$$

But since

$$d\log h_t^{LTV} = \left(\frac{\mu\theta^{LTV}}{1-\mu\theta^{LTV}-\psi}\right) d\log\theta^{LTV} - d\log p + d\omega$$

we obtain

$$\begin{aligned} d\log M^* = {} & s_M^{LTV} \left(\frac{1-\psi}{1-\mu\theta^{LTV}-\psi}\right) d\log\theta^{LTV} + s_M^{LTV} d\omega \\ & + (s_M^{both} + s_M^{PTI}) \left(d\log\theta^{PTI} - d\log\hat{r}^*\right). \end{aligned}$$

It is straightforward to show from (23) that $M \propto M^*$ in steady state, so that $d\log M = d\log M^*$, from which (29) immediately follows.

A.7 Calibration Details

Calibration of ν. To calibrate ν, I compute the share of principal that would be repaid in each period under actual 30-year mortgage contracts under the assumption that mortgages randomly prepay with probability ρ. Given a mortgage interest rate expressed as an annual percentage rate (APR), I first derive the monthly interest rate $r_m = APR/12$ and the corresponding monthly prepayment rate $\rho_m = 1 - (1-\rho)^{1/3}$. I then simulate the monthly evolution of a unit-balance loan over $N = 360$ months, calculating the fixed monthly coupon c, interest payments i_t, and principal repayments $p_t = c - i_t$ for each period t. The parameter ν is computed by taking the ratio of the weighted average principal repayment to the weighted average balance, where the weights correspond to the fraction of loans surviving at each age:

$$\nu = 3 \times \frac{\sum_{t=0}^{N-1}(1-\rho_m)^t p_t}{\sum_{t=0}^{N-1}(1-\rho_m)^t b_t} \tag{A.38}$$

where b_t is the remaining balance at month t, and the factor 3 converts from monthly to quarterly frequency. This procedure ensures that the aggregate amortization rate in the model is consistent with the contractual repayment schedule of the underlying mortgage pool, accounting for the endogenous shortening of effective duration caused by prepayments.

Calibration of ρ. To calibrate ρ, I first define a scaled version of the mortgage law of motion (23). If we define $LTV_t = M_t/\text{Value}_{B,t}$, for $\text{Value}_{B,t} \equiv p_t \bar{H}_B$, then we can manipulate (23) to yield

$$LTV_t = \rho LTV_t^* + (1-\rho)(1-\nu)\left(\frac{\text{Value}_{B,t-1}}{\text{Value}_{B,t}}\right) LTV_{t-1}, \tag{A.39}$$

for $LTV_t^* = M_t^*/\text{Value}_{B,t}$. I set ρ to optimize this equation's fit to the aggregate data. To obtain data counterparts of these model objects, I first set M_t equal to aggregate mortgage credit. For $\text{Value}_{B,t}$, we need to adjust the total value of household-owned housing to remove housing held by outright owners. To this end, I define $\text{Value}_{B,t}$ to be equal to $0.7270 \times \text{Value}_t$, where Value_t is the aggregate value of household-owned real estate, and 0.7270 is the share of housing owned by households with a mortgage in the 1998 Survey of Consumer Finances.

I set the new LTV ratio to $LTV_t^* = 0.8$, which, as discussed in Section 4.3, is the average LTV limit on new loans. For simplicity, this approach ignores the presence of the PTI limit

in defining M_t^*. However, given that most PTI-constrained borrowers in the data also appear to be LTV-constrained, and hence will have LTV ratios around the typical benchmark of 80%, this is a reasonable assumption. To the extent that binding PTI limits would limit borrowing, this should cause me to understate ρ, diminishing the importance of credit for house prices in the model. Hence, I consider this to be a conservative calibration approach.

Last, I note that for each value of ρ, I need to compute the implied value of ν, given the procedure described in 4.3.

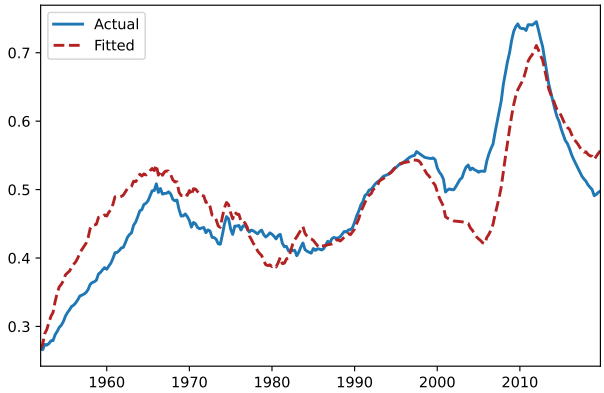

Figure A.3: Calibration of ρ, Comparison of Actual vs. Fitted Aggregate LTV Ratios

Note: This figure compares the fitted values of the aggregate borrower LTV ratio implied by equation (A.39) with its actual data counterpart. The data equivalent is equal to aggregate mortgage credit in the FAUS divided by $\text{Value}_{B,t}$, which is in turn defined as the share of housing owned by households with a mortgage (0.7270) times the aggregate value of household-owned housing in the FAUS. See Table B.1 for detailed data sources.

With these definitions, I initialize $\widehat{LTV}_0$ to match the actual value in 1952:Q1, and then run the model forward to obtain an implied time series $\widehat{LTV}_t$. I then set ρ to minimize the sum of squared errors

$$\sum_{t=1}^{T} \left(\widehat{LTV}_t - LTV_t \right)^2$$

over the sample period 1952:Q1–2019:Q4. The resulting fit is displayed in Appendix Figure A.3, showing that this approach delivers a reasonably close fit over many decades.

Calibration of σ_e. As discussed in Section 4.3, because PTI limits are effectively slack during the boom period, PTI ratios should be proportional to e_i, and hence we should be able to estimate σ_e from the standard deviation of $\log PTI_{i,t} - \log \overline{PTI}_t$, where $\log \overline{PTI}_t$ is the average value of $\log PTI_{i,t}$ at each date.

In practice, there are two complications that need to be overcome. The first is that some borrowers may be unconstrained rather than LTV-constrained, in which case the model has no prediction on what their PTI ratio should be. To address this, I restrict the sample to loans with CLTV ratios in the set $S = \{70, 75, 80, 85, 90, 95, 97\}$ at which we observe large bunched masses in the CLTV distribution.

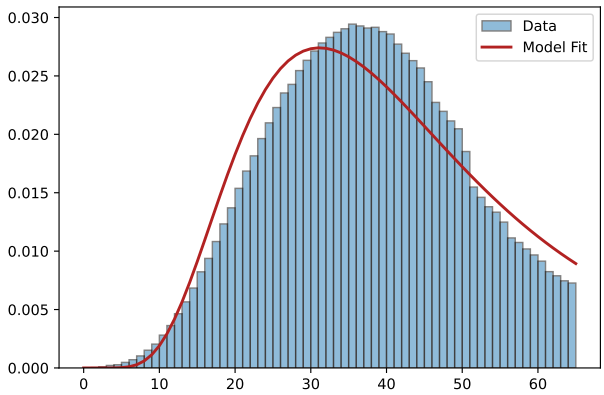

Figure A.4: Calibration of σ_e, Actual vs. Fitted Boom PTI Distributions

Note: This figure plots the distribution of PTI ratios among Fannie Mae purchase loans with CLTV ratios in the set $S = \{70, 75, 80, 85, 90, 95, 97\}$ originated during 2004–2006. The histogram displays the empirical density of PTI ratios (in percent) across this sample. The solid red curve overlays the fitted lognormal density, obtained by jointly estimating the aggregate mean $\log \overline{PTI}_t$ and dispersion σ_e via maximum likelihood under the assumption that the PTI distribution is lognormal and truncated from above at 65%. Source: Fannie Mae Loan Performance dataset.

The second issue is that the data set by construction excludes loans with PTI ratios above 65. Since we do not observe any bunching as we approach 65%, after which the data are truncated by the data provider, it seems reasonable to assume that the distribution continues smoothly through this point. As a result, a simple sample standard deviation over the observed loans understates the true dispersion. To correct for this, I jointly estimate the aggregate mean $\log \overline{PTI}_t$ and the dispersion σ_e using maximum likelihood, treating the observations as draws from a log-normal distribution truncated from above at 65%. The estimation is performed on the sample of Fannie Mae purchase loans with CLTV in the set S originated during 2004–2006, yielding the value of σ_e reported in Table 3. Appendix Figure A.4 overlays the implied lognormal density on the histogram of observed PTI ratios, showing that this approach yields a close approximation for the true distribution.

Calibration of γ. I calibrate γ using the regression

$$\log \text{value}_{i,t} = \alpha_t + \gamma \log PTI_{i,t} + \varepsilon_{i,t} \tag{A.40}$$

where $\log \text{value}_{i,t}$ is log implied house value (obtained by dividing the loan balance by the loan-to-value ratio), $PTI_{i,t}$ is the PTI ratio at origination, and α_t is a time (year-quarter) fixed effect. This regression is estimated on the Fannie Mae purchase loan data for 2000–2007. Results can be found in Table A.1.

Table A.1: Regression to Calibrate γ

Constant	$\log PTI$	N	Adj. R^2
-0.082***	0.147***	3,073,707	0.012
(0.000)	(0.001)		

Note: This table reports the results of the regression (A.40). The regression is estimated on Fannie Mae purchase loan data for the sample 2000–2007. Standard errors in parentheses are heteroskedasticity-robust.

Calibration of α. I calibrate $\alpha = 0.8713$ so that in my "Full Boom" experiment in Section 6, the average PTI ratio under peak house prices matches the average PTI ratio ($\log \overline{PTI}_t$) observed in the 2004–2006 Fannie Mae loan-level data. One issue is that this "average" PTI ratio in the data is the maximum likelihood estimate from the truncated distribution, assuming the true distribution is lognormal. While the relaxed PTI limit of 70% in the model experiment is very high, it does feature some bunching, and hence is not a fully lognormal distribution. To address this, I compute the hypothetical mean if all variables were the same, but the PTI limit did not bind at all, so that $h^*_{i,t} = e^h_i h^{LTV}_t$ for all borrowers. I then equate this hypothetical unconstrained mean in the model to the empirical value for $\log \overline{PTI}_t$ obtained from the data.

A.8 Alternative Calibrations

This section presents the alternative calibration targeting moments from the 2014–2016 post-bust sample. Other than targeting alternative empirical values for the various moments, the calibration procedure is identical to that described in 4.3 with two exceptions: (i) I set the maximum PTI ratio to $\theta^{PTI} = 0.4500$, matching the evidence in Figure 2; and (ii) instead of normalizing ω to zero, I calibrate it so that growth in price-rent ra-

tios between the pre-boom steady state and post-bust steady state in the model is equal to growth in empirical price-rent ratios from the 1994–1996 period to the 2014–2016 period. Appendix Table A.2 presents the parameter values, and for parsimony includes only those parameters that are updated between calibrations (β_B, β_S, ω, θ^{PTI}, κ, η, and π). The negative value of ω reflects that prices were depressed in this period relative to what the model would predict in an environment with such low interest rates, likely due to persistent consequences of the global financial crisis not captured by the model.

Table A.2: Parameter Values, Post-Bust (2014–2016) Calibration

Parameter	Name	Value	Internal	Target/Source
Demographics and Preferences				
Borr. discount factor	β_B	0.9854	Y	PMI premium = 0.8% (ann.)
Saver discount factor	β_S	0.9990	Y	Real rate target
Ownership utility	ω	-0.3931	Y	Price-rent ratio
Mortgages				
Max PTI ratio	θ^{PTI}	0.4500	N	Post-bust limit
PTI offset (taxes, ins.)	κ	0.0050	Y	See text
Term premium (mean)	η	0.0042	N	Avg. mortgage rate
Technology and Government				
Steady state inflation	π	1.0044	N	Avg. infl. expectations, 2014–2016

A.9 Additional Model Results

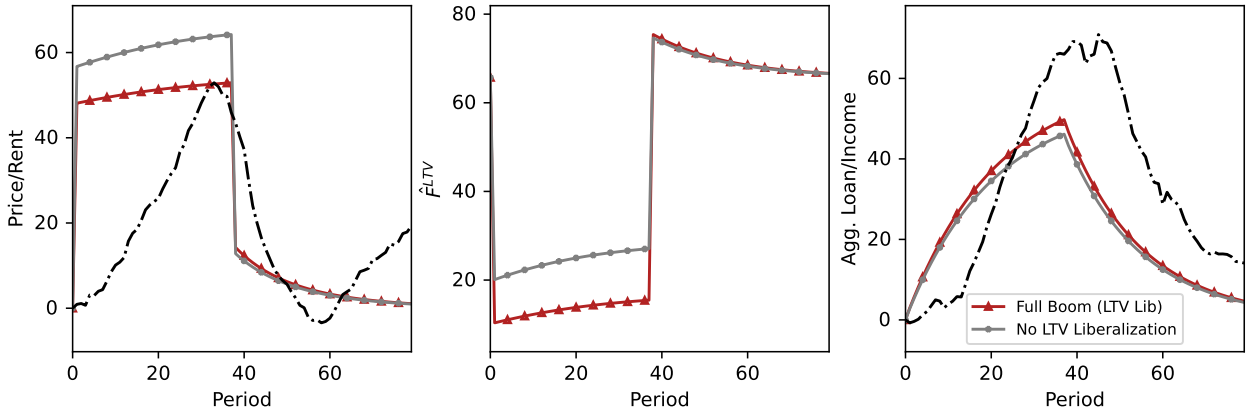

(a) LTV Liberalization, Baseline Economy

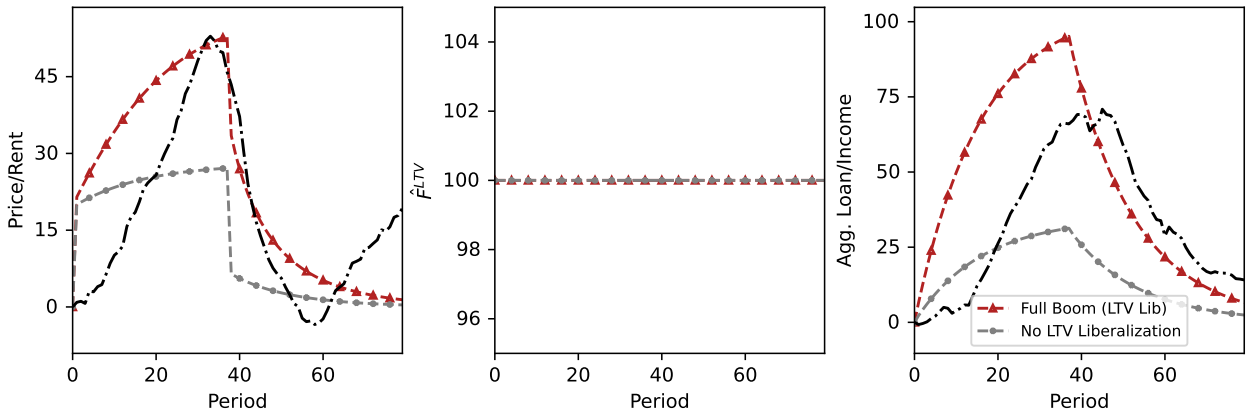

(b) LTV Liberalization, LTV Only Economy

Figure A.5: Macroprudential Policy Counterfactuals, LTV Limit

Note: For the path "Full Boom (LTV Lib)," the LTV limit θ^{LTV} is unexpectedly loosened from 0.80 to 0.99 at time zero, in addition to a fall in interest rates and an increase in ownership utility ω, chosen to exactly match the full increase in price-rent ratios observed during the housing boom. Note that the required change in ω will differ from that of Figure 8b since it is applied alongside a different change in credit standards. After 36 quarters, these values are unexpectedly returned to their steady state values. In the "No LTV Liberalization" counterfactual, the same fall in interest rates and increase in ω are applied, and then reversed after 36 quarters, but the LTV limit θ^{LTV} is maintained at its initial steady state level. Panel (a) shows these paths for the Baseline economy, while Panel (b) shows them for an "LTV Only" economy ($\theta^{PTI} \to \infty$), fitting a separate change of ω in each economy. A value of 1 represents a 1% increase relative to steady state, except for $\hat{F}_t^{LTV}$, which is measured in percentage points. The price-rent ratio is $p_t/(u^h_{B,t}/u^c_{B,t})$, where the denominator is the implied price of rental services. The aggregate loan-income ratio is M_t/Y_t. Where applicable, the black dash-dot lines display the corresponding paths in the data.

B Empirical Appendix

B.1 Data Description

This section describes the various data used in the paper, and provides additional histograms and moments to support the empirical claims of the paper.

B.1.1 Macroeconomic Data

Sources for the various macroeconomic data used in the paper can be found in Table B.1 below. These data were obtained in three ways. The Mortgage Rate and Housing Services data were downloaded from Federal Reserve Economic Data (FRED), provided by the Federal Reserve Bank of St. Louis (`https://fred.stlouisfed.org/`). The Income, Value, Debt, and Gross Income data were obtained as part of the Financial Accounts of the United States data, provided by the Federal Reserve Board of Governors (`https://www.federalreserve.gov/releases/z1/`). Last, the 10Y Inflation Expectations series was obtained from the Inflation Expectations data, provided by the Federal Reserve Bank of Cleveland (`https://www.clevelandfed.org/indicators-and-data/inflation-expectations`).

Table B.1: Data Definitions

Name	Definition	Source	Code
	Main Text + Boom-Bust Experiments		
Income	HH + Nonprofit Disposable Income	FAUS	FA156012005.Q
Value	HH + Nonprofit Res. Real Estate	FAUS	FL155035005.Q
Debt	HH + Nonprofit Home Mort. (Liability)	FAUS	FL153165105.Q
Gross Income	HH + Nonprofit Personal Income	FAUS	FA156010001.Q
Mortgage Rate	30Y Conventional Mortgage Rate	BoG	MORTGAGE30US
Housing Services	Owner-Occupied Housing Services	BEA	A2013C1A027NBEA
Inflation Expectations	10Y Inflation Expectations	CleFed	N/A

Note: Data sources are Bureau of Economic Analysis (BEA), Federal Reserve Bank of Cleveland (CleFed), Federal Reserve Board of Governors (BoG), Financial Accounts of the United States (FAUS), Federal Housing Finance Administration (FHFA). Codes are FRED mnemonics with the exception of data from FAUS.

Computation of the Price-Rent Ratio. For the price-rent ratio, the most appropriate series is the ratio of the value of household real estate to the value of imputed owner-occupied housing services. While this numerator series is available at quarterly frequency, the denominator is only available annually. To address this, I follow Greenwald and

Guren (2025) Appendix B.5 in interpolating the annual series to quarterly frequency using the method of Chow and Lin (1971).

B.1.2 Fannie Mae Loan-Level Data

This set is taken from Fannie Mae's Single Family Loan Performance Data.[40] From the Fannie Mae data description:

> The population includes a subset of Fannie Mae's 30-year, fully amortizing, full documentation, single-family, conventional fixed-rate mortgages. This dataset does not include data on adjustable-rate mortgage loans, balloon mortgage loans, interest-only mortgage loans, mortgage loans with prepayment penalties, government-insured mortgage loans, Home Affordable Refinance Program (HARP) mortgage loans, Refi Plus mortgage loans, and non-standard mortgage loans. Certain types of mortgage loans (e.g., mortgage loans with LTVs greater than 97 percent, Alt-A, other mortgage loans with reduced documentation and/or streamlined processing, and programs or variances that are ineligible today) have been excluded in order to make the dataset more reflective of current underwriting guidelines. Also excluded are mortgage loans originated prior to 1999, sold with lender recourse or subject to other third-party risk-sharing arrangements, or were acquired by Fannie Mae on a negotiated bulk basis.

Additional histograms and quantiles from this dataset are displayed in Figures B.16 - B.18 below.

Computing Implied Income. Since the Fannie Mae data do not include observations on income, I compute implied effective income as follows. I begin with the PTI formula

$$PTI_{i,t} = \frac{(r_{i,t} + \nu_{i,t} + \kappa_{i,t})m^*_{i,t} + d_{i,t}}{y_{i,t}},$$

where $m^*_{i,t}$ is the principal balance, $(r_{i,t} + \nu_{i,t} + \kappa_{i,t})m^*_{i,t}$ is the total payment of principal, interest, taxes, and insurance, $d_{i,t}$ are outside debt payments, and $y_{i,t}$ is income. Rearranging, we obtain

$$y_{i,t} = \frac{(r_{i,t} + \nu_{i,t} + \kappa_{i,t})m^*_{i,t} + d_{i,t}}{PTI_{i,t}}$$

[40] `http://www.fanniemae.com/portal/funding-the-market/data/loan-performance-data.html`

The Fannie Mae data set provides information on $r_{i,t}$, $m^*_{i,t}$, and $PTI_{i,t}$, and $\nu_{i,t}$ can be computed from the standard fixed rate mortgage schedule. However, I do not have information on $d_{i,t}$, or the exact value of $\kappa_{i,t}$. To this end, I proxy for income using

$$\hat{y}_{i,t} = \frac{(r_{i,t} + \nu_{i,t} + \hat{\kappa})m^*_{i,t}}{PTI_{i,t}}$$

where $\hat{\kappa} = 0.0175$. While it would be possible to proxy for an average value of $d_{i,t}$, this would effectively impose a floor of the PTI ratio, and ignore that many borrowers do not have large recurring debt payments, leading to a worse fit of the data. Instead, I will allow the parameter α to account for average deviations due to $d_{i,t}$, which appears to lead to a closer fit in practice.

B.1.3 HMDA Data

The Home Mortgage Disclosure Act (HMDA) data are used to construct pre-boom measures of the PTI-constrained share at the county level. The raw HMDA data cover mortgage applications and originations reported by lending institutions and are available on an annual basis. From this universe I retain only owner-occupied, single-family, purchase loan originations (action taken $= 1$, loan purpose $= 1$, occupancy $= 1$), dropping records with missing applicant income or loan amount, invalid state or county codes, or loan types outside the range 1–4 (conventional, FHA, VA, and FSA/RHS).

Removing Second Liens. A key challenge in using the HMDA data for this analysis is that lien status information is only reliably recorded in the data beginning in 2004. Because second liens (e.g., home equity loans and piggyback loans) tend to have systematically different characteristics than first-lien purchase mortgages, their inclusion would distort measures of the PTI distribution for first-lien borrowers. This issue can be visualized in Appendix Figure B.1, which shows histograms of the loan-to-income distributions for various years. As can be seen, for some years, particularly during the housing boom period (e.g., 2006), the distribution is bimodal due to a large mass of second liens with low LTI ratios.

To address this, I employ a machine-learning classifier to identify and remove likely second liens for the full sample period 1990–2016. Specifically, I train a Random Forest classifier using HMDA records from 2004–2007, years for which lien status is observed. The binary outcome variable is an indicator for whether a loan is a second lien. The feature

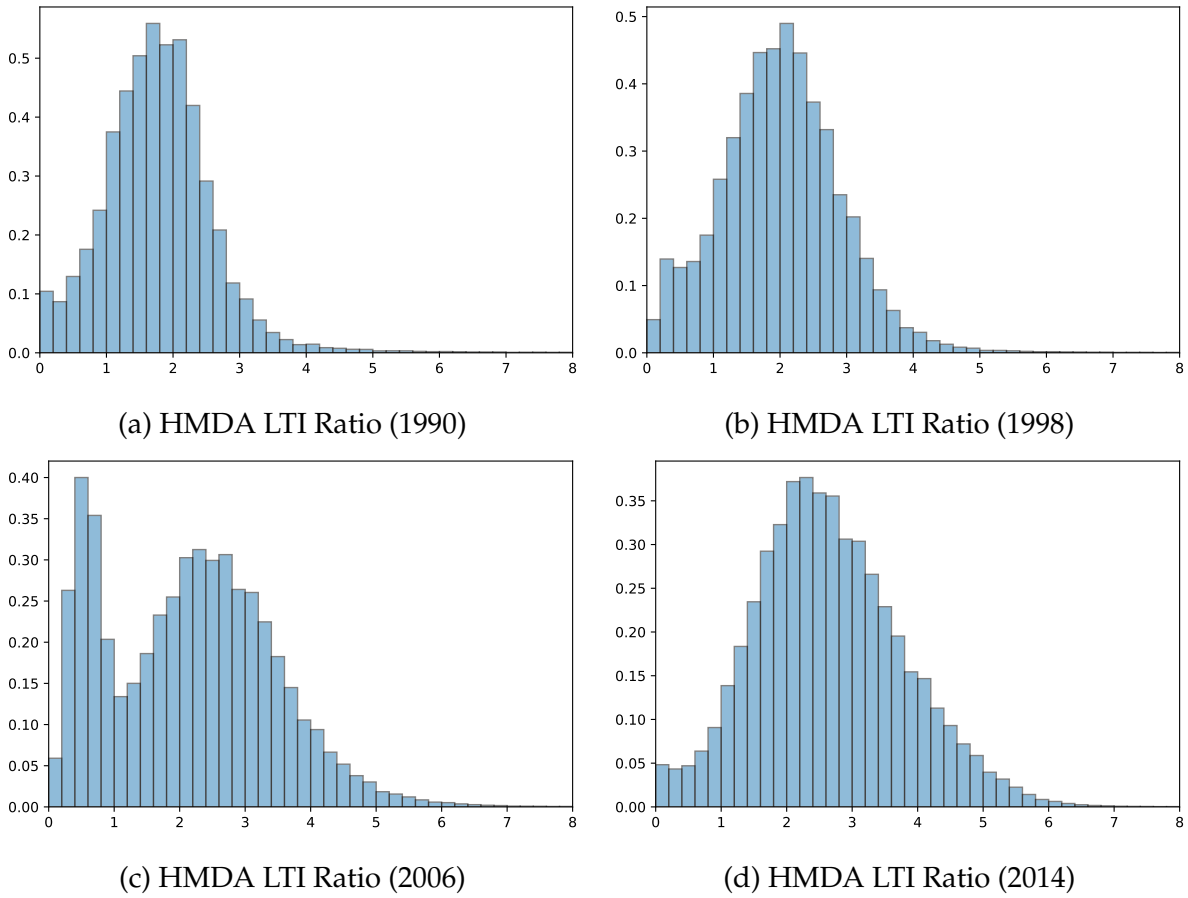

(a) HMDA LTI Ratio (1990)

(b) HMDA LTI Ratio (1998)

(c) HMDA LTI Ratio (2006)

(d) HMDA LTI Ratio (2014)

Figure B.1: HMDA Loan-to-Income Ratios

Note: Each panel displays the empirical density of loan-to-income (LTI) ratios for HMDA owner-occupied, single-family, purchase mortgage originations in the indicated year. The LTI ratio is computed as the loan amount divided by applicant income. The histogram bins LTI values over the range $[0, 8]$ using 40 equally spaced bins, normalized so that the histogram integrates to one. The sample includes all originations satisfying the data restrictions described in Appendix B.1.3, with no restriction on lien status. Source: Home Mortgage Disclosure Act data.

set comprises two continuous variables—the log loan-to-income ratio ($\log \text{LTI}$) and the log loan-to-value proxy ratio ($\log \widehat{LTV}$), where the LTV proxy is computed by dividing the loan balance by the FHFA county-level HPI—and four categorical variables: purchaser type, loan type, an indicator for whether the record carries an edit status flag, and an indicator for whether the loan amount is below \$10,000. The classifier uses 50 decision trees with a maximum depth of 10, trained on a 25% random subsample of the 2004–2007 pooled data, with the remaining 75% held out for out-of-sample performance evaluation. The fitted model is then applied to each year from 1990 to 2016 to generate predicted lien status for all records. Observations predicted to be second liens are excluded from the subsequent analysis.

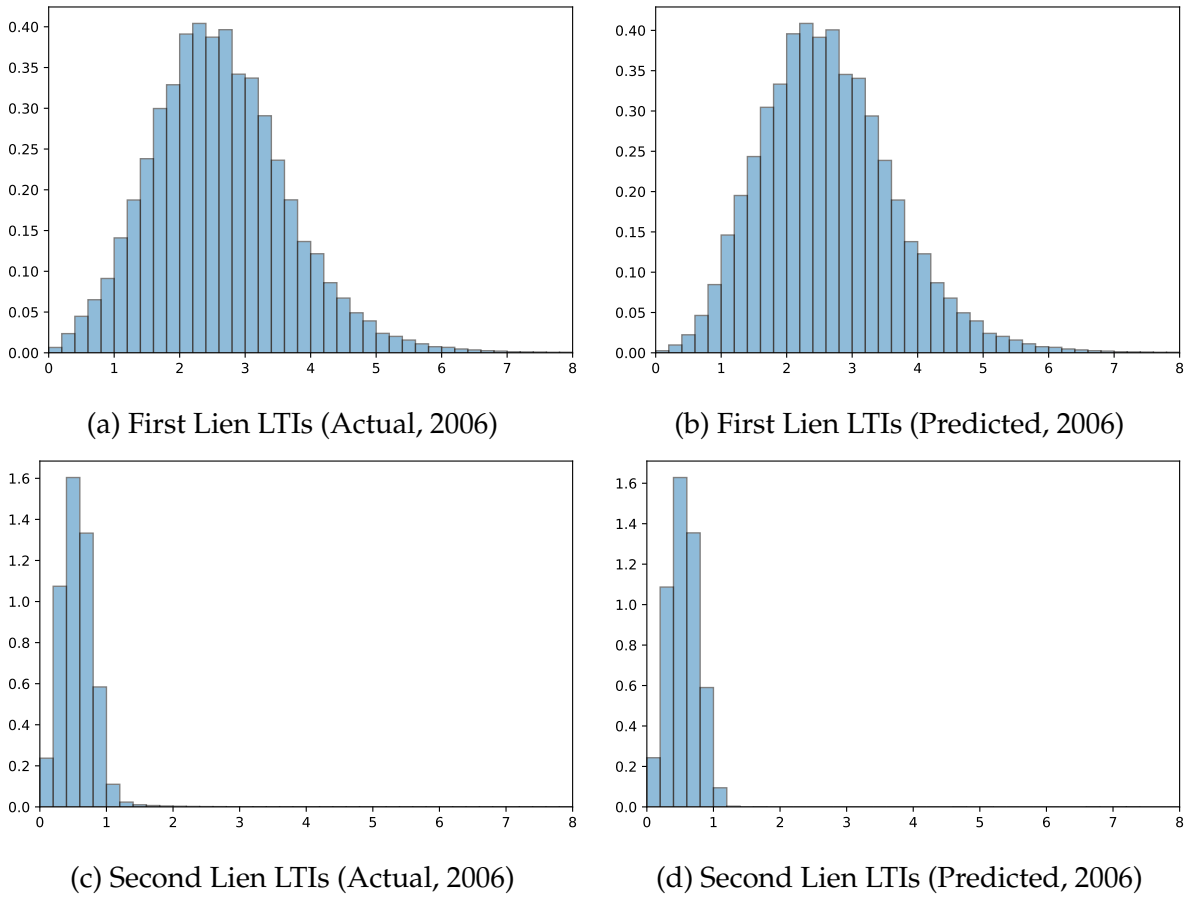

(a) First Lien LTIs (Actual, 2006)

(b) First Lien LTIs (Predicted, 2006)

(c) Second Lien LTIs (Actual, 2006)

(d) Second Lien LTIs (Predicted, 2006)

Figure B.2: HMDA Loan-to-Income Ratios by Lien Status

Note: Each panel displays the empirical density of loan-to-income (LTI) ratios for HMDA owner-occupied, single-family, purchase mortgage originations in 2006, separately by lien status. The LTI ratio is computed as the loan amount divided by applicant income. The histogram bins LTI values over the range $[0, 8]$ using 40 equally spaced bins, normalized so that the histogram integrates to one. Panels (a) and (c) condition on actual recorded lien status (first and second lien, respectively). Panels (b) and (d) condition on lien status as predicted by the Random Forest classifier described in Appendix B.1.3 (predicted first and second lien, respectively). Source: Home Mortgage Disclosure Act data.

Figure B.2 shows the performance of the classifier for 2006, for which we have information on lien status. The predicted distributions are nearly indistinguishable from the actual distributions, showing that the random forest is able to accurately classify the first vs. second liens. Last, Figure B.3 displays the predicted splits into first and second liens in the pre-2004 data, showing that the extracted distributions look very reasonable, and highly consistent with the distributions for the periods for which we have lien status information.

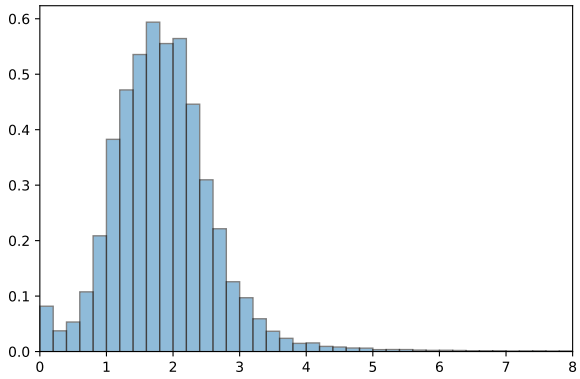

(a) First Lien LTIs (Predicted, 1990)

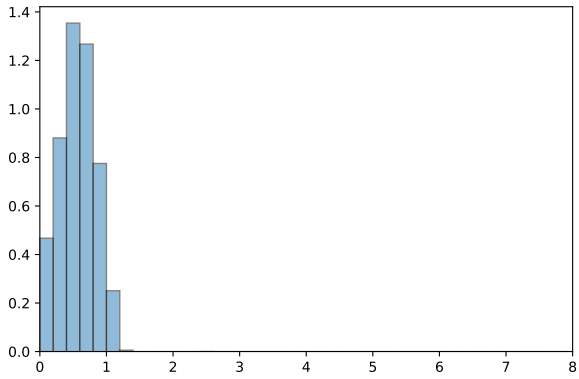

(b) Second Lien LTIs (Predicted, 1990)

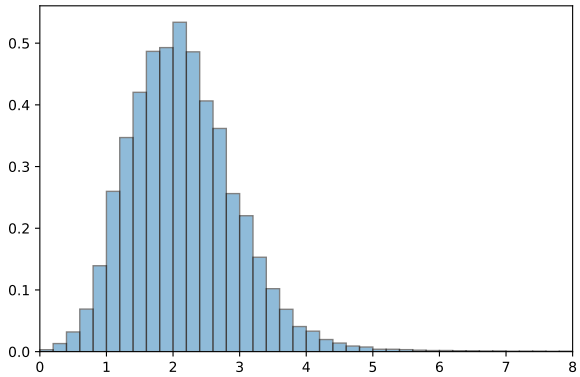

(c) First Lien LTIs (Predicted, 1998)

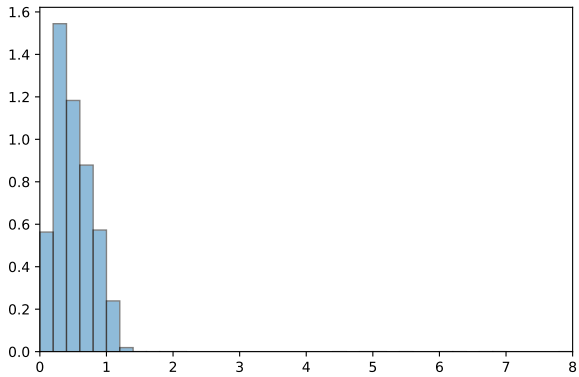

(d) Second Lien LTIs (Predicted, 1998)

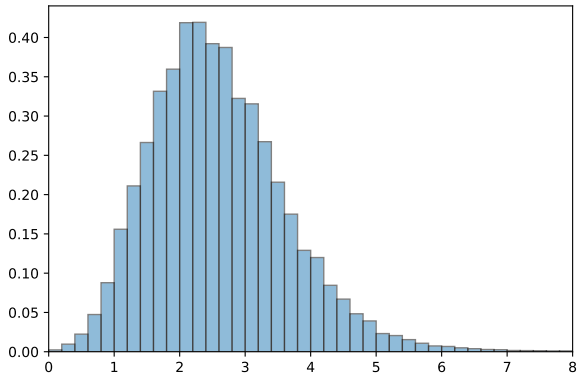

(e) First Lien LTIs (Predicted, 2003)

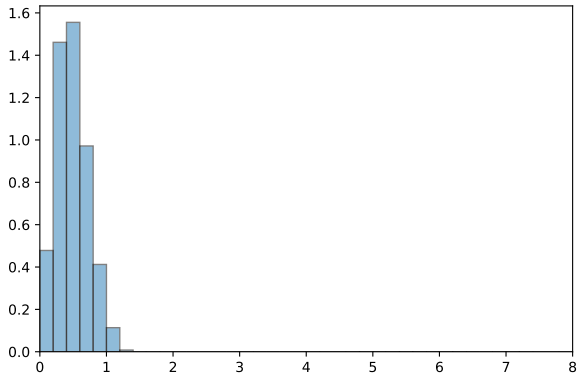

(f) Second Lien LTIs (Predicted, 2003)

Figure B.3: HMDA Loan-to-Income Ratios

Note: Each panel displays the empirical density of loan-to-income (LTI) ratios for HMDA owner-occupied, single-family, purchase mortgage originations in the indicated year, separately by predicted lien status. The LTI ratio is computed as the loan amount divided by applicant income. The histogram bins LTI values over the range $[0, 8]$ using 40 equally spaced bins, normalized so that the histogram integrates to one. Because lien status is not reliably recorded in HMDA prior to 2004, lien status for all panels is assigned using the Random Forest classifier described in Appendix B.1.3. Left panels condition on predicted first-lien originations; right panels condition on predicted second-lien originations. Source: Home Mortgage Disclosure Act data.

Computing the Implied PTI Ratio. Because HMDA data do not include the actual interest rate paid on the loan, or information on taxes, insurance, and non-mortgage debt payments, the true PTI ratio cannot be computed using HMDA data. However, we can proxy for these additional variables to obtain an approximate implied PTI ratio.

To do so, I first approximate the total coupon payment on the loan (the sum of interest, principal, taxes, and insurance). I begin with the Freddie Mac national average 30-year mortgage rate described in Appendix B.1.1. I convert this annual interest rate to a monthly interest rate using

$$r_{monthly,t} = (1 + r_{ann,t})^{1/12} - 1. \tag{B.1}$$

Using this rate, the effective payment rate can be computed as

$$\hat{r}_t = 12 \times r_{monthly,t} \times \left(1 - (1 + r_{monthly,t})^{-360}\right)^{-1} + 0.0175. \tag{B.2}$$

where the 0.0175 proxies for taxes and insurance payments. Using (B.2), we can approximate the PTI ratio as

$$\widehat{PTI}_{i,t} \equiv \frac{\hat{r}_t \times \text{balance}_{i,t}}{\text{income}_{i,t}}$$

where $\text{balance}_{i,t}$ is the initial balance on the loan, and $\text{income}_{i,t}$ is borrower income, both of which are available from HMDA.

B.1.4 Survey of Consumer Finances Data

To obtain calibration targets for the model, I use the Survey of Consumer Finances, obtained from the Federal Reserve Board of Governors.

B.1.5 Census PUMS Data

I use microdata from the 1990 Census of Population and Housing Public Use Microdata Sample (PUMS), accessed at the Public Use Microarea (PUMA) level. The main Census analysis restricts attention to recent movers by using the person-level PUMS extract and keeping only the household reference person (PERNUM $=$ 1) in households that moved into their current residence within the last year (MOVEDIN $=$ 1). I then limit the sample to owner-occupied households that report positive household income

and at least one positive mortgage payment (first or second mortgage). The key variables extracted from the Census are: first mortgage payment (MORTAMT1), second mortgage payment (MORTAMT2), property insurance payments (PROPINSR), property tax payments (PROPTX99, provided as categorical bins), household income (HHINCOME), household weights (HHWT), person number (PERNUM), and recent-mover status (MOVEDIN). Property tax values are converted from their labeled bins to dollar midpoints. Property insurance values of 99,999 or above and household income values of 9,999,999 or above indicate top-coded or missing observations and are set to missing, with the former set to zero for the purpose of computing total housing payments and the latter causing the household to be dropped from the sample.

Computing the PTI Ratio. The housing payment-to-income (PTI) ratio is constructed as the ratio of total monthly housing payments to monthly income. Total monthly housing payments are the sum of the first and second mortgage payments (MORTAMT1 + MORTAMT2), property insurance (PROPINSR), and property taxes (PROPTX99 /12, after converting annual taxes to a monthly basis). Monthly income is computed as annual household income divided by 12. Formally,

$$PTI_i = \frac{(MORTAMT1_i + MORTAMT2_i) + (PROPINSR_i + PROPTX99_i/12)}{HHINCOME_i/12}.$$

For each PUMA, I then compute the share of mortgaged homeowner households with $PTI_i > 28\%$, weighting by household weights (HHWT). To obtain PUMA-level house price and boom-era income information, I aggregate FHFA house prices indices and IRS data (see Appendix B.1.6 for details on these sources) from the ZIP-5 level to the PUMA level using the Geocorr 1990 crosswalk (Missouri Census Data Center, 1990). These PUMA-level house price index and income series are merged with the PTI $> 28\%$ shares to form the dataset used in the Census-based cross-sectional analysis described in Section 3.3.

B.1.6 Geographic-Level Data

I use geographic-level data at the three-digit ZIP (ZIP3) and county levels for cross-sectional regressions in Sections 3.2 and 3.3.

IRS Data The IRS Statistics of Income (SOI) data contain annual tabulations of individual income tax returns at the 5-digit ZIP code level. The data are available for 1998, 2001,

and annually from 2004 onward. Each annual file is distributed as a collection of state-level spreadsheets that are concatenated to form a national dataset. The key variables extracted from each file are total wages and salaries, and the number of returns. Given these variables, I sum over 5-digit ZIP codes to obtain sums of wages and salaries and the number of returns at the ZIP3 level. I then divide total wages and salaries by the total number of returns to obtain a ZIP3-level average for wage and salary income for each period.

FHFA Data. For local house price indices, I use the Federal Housing Finance Agency (FHFA) All-Transactions Indexes, which can be directly downloaded at the County and ZIP3 levels.

Census and American Community Survey Data. I obtain data on median gross rent at the county and census tract levels from the IPUMS National Historical Geographic Information System (NHGIS). For rents at the start of the housing boom, I use Table NH063A from the 2000 Decennial Census (NHGIS ID: ds151). For rents at the end of the housing boom, I use Table B25064 from the 2009 American Community Survey: 5-Year Data (NHGIS ID: ds195) covering the years 2005–2009. In addition, I obtain household counts from Table NP015A of the 2000 Decennial Census (NHGIS ID: ds146).

Crosswalks. To aggregate from the census tract or ZIP level to the PUMA or ZIP3 level, I use the "Geocorr" crosswalks from the Missouri Census Data Center.

B.2 Estimating the Fraction Constrained

This section provides details on the methodology used to measure the share of borrowers constrained by LTV and PTI limits, as described in Section 3.2, as well as additional detail on the results.

B.2.1 Estimation Procedure

The estimation proceeds in seven steps.

Step 1: Collapse Fannie Mae Data to ZIP3 Level. For the first step, I collapse the loan-level data to ZIP3 level by averaging individual indicators for $PTI \leq j$ for $j \in \{1, 2, \ldots, 64\}$.

Since all loans have PTI ratio less than or equal to 64 in this top-coded data set, the final indicator will be the total number of loans for a given ZIP3 and year. Dividing the summed indicators by this total yields shares of loans with PTI ratio less than or equal to that value (i.e., the CDF), which I denote $F_{i,t}^{PTI \leq j}$, where j is the PTI ratio value, i is the ZIP3, and t is the year. Note that $F_{i,t}^{PTI \geq 29}$ can be computed as $1 - F_{i,t}^{28}$.

For the conditional regressions determining whether a borrower is jointly constrained by PTI and LTV, I also compute the following loan-level indicators: (i) an LTV-limit indicator for whether the combined LTV ratio falls in the set of standard limits $S = \{70, 75, 80, 85, 90, 95, 97\}$,

$$LTVLimit_{i,t} = \mathbf{1}_{LTV_{i,t} \in S}$$

and (ii) a control indicator

$$Control_{i,t} = \frac{1}{2}\left(\mathbf{1}_{\{LTV_{i,t}+1 \in S\}} + \mathbf{1}_{\{LTV_{i,t}-1 \in S\}}\right).$$

When summed over a geographic area, the sum of $LTVLimit_{i,t}$ reflects the share of loans actually at LTV bunching bins, while $Control_{i,t}$ reflects a proxy for the share of loans that would be in those bins, absent LTV constraints, by averaging mass over the neighboring bins. I also compute the joint shares of loans with $PTI \leq j$ and each of these indicators, obtained by multiplying $\mathbf{1}_{\{PTI_{i,t} \leq j\}}$ by $LTVLimit_{i,t}$ and $Control_{i,t}$ and averaging at the ZIP3 level, which I denote $F_{i,t}^{PTI \leq j, LTV}$ and $F_{i,t}^{PTI \leq j, Control}$, respectively. Dividing these joint shares by the marginal shares $F_{i,t}^{LTV}$ and $F_{i,t}^{Control}$ yields conditional shares, denoted $F_{i,t}^{PTI \leq j | LTV}$ and $F_{i,t}^{PTI \leq j | Control}$. I also note that given any probability that $PTI \leq j$, it is straightforward to compute the probability that $PTI \geq j$ as the complement, and I will at times convert between the two forms below.

Step 2: Construct Predictor. For the second step, I construct the predictor variable $\log e_{i,t}^*$ at the ZIP3-year level. From (2), this requires three inputs: ZIP3 level house prices ($p_{i,t}$), ZIP3 level income ($y_{i,t}$), and the aggregate effective payment rate ($\hat{r}_t^*$). For house prices I use the FHFA ZIP3 House Price Index, described in Appendix B.1.6. For income, I use average wage and salary income from the IRS Statistics of Income, described in Appendix B.1.6. For the effective payment rate $\hat{r}_t^*$ I follow the procedure described in equations (B.1) and (B.2), using the same Freddie Mac mortgage rate series. Combining these sources yields $\log e_{i,t}^*$ at the ZIP3-year level, which I merge into the data set constructed in the first step.

Step 3: Estimate Distributional Regressions. For the third step, for each PTI threshold $j \in \{1, 2, \ldots, 64\}$ I estimate the regression

$$F_{i,t}^{PTI \leq j} = \alpha_i + \beta_1^j \log e_{i,t}^* + \beta_2^j F_{i,t}^{PTI \geq 29} + \varepsilon_{i,t}^j. \tag{B.3}$$

where α_i is a ZIP3 fixed effect. I weight the regression by the number of loans in each ZIP3-year observation, and cluster the standard errors by ZIP3. This yields fitted values $\widehat{F}_{i,t}^j$ for each period. To compute estimates within each percentage point bin, we can compute

$$\hat{f}_{i,t}^{PTI=j} = \widehat{F}_{i,t}^{PTI \leq j} - \widehat{F}_{i,t}^{PTI \leq j-1}$$

for $F_{i,t}^{PTI \leq 0} = 0$. These values, averaged over the post-bust sample using loan counts as weights, comprise the counterfactual distribution histogram in Figure 3.

Step 4: Estimate Constraint Regressions. To estimate the fraction of borrowers bound by each constraint, I run the regression

$$F_{i,t}^{PTI \geq j} = \alpha_i + \beta_1^j \log e_{i,t}^* + \beta_2^j F_{i,t}^{PTI \geq 29} + \varepsilon_{i,t}^j. \tag{B.4}$$

where $F_{i,t}^{PTI \geq j}$ is the share of loans with PTI ratio $PTI \geq j$ for $j \in \{46, 51\}$. Trivially, we can compute these as $F_{i,t}^{\geq 46} = 1 - F_{i,t}^{45}$ and $F_{i,t}^{\geq 51} = 1 - F_{i,t}^{50}$.

After estimating these regressions on the boom (pre-2008) sample, I aggregate the observed values $F^{PTI \geq j}$, fitted values $\widehat{F}_{i,t}^{PTI \geq j}$ and predictor variables by taking averages over ZIP3 values for each year weighted by loan count. I denote these aggregate values as $F_t^{PTI \geq j}$, $\widehat{F}_t^{PTI \geq j}$ and $x_t' = (\log e_t^*, F_t^{PTI \geq 29})$, respectively. I use the variance-covariance matrix of the β^j estimates, clustered at the ZIP3 level, to compute the variance of the fitted values as

$$\text{Var}(\widehat{F}_t^{PTI \geq j}) = \tilde{x}_t' \text{Cov}(\beta^j) \tilde{x}_t$$

where $\tilde{x}_t$ are the values of the regressors with fixed effects removed (or the intercept, depending on the specification). It is then straightforward to compute the estimate of the missing mass as

$$\tilde{F}_t^{PTI \geq j} = \widehat{F}_t^{PTI \geq j} - F_t^{PTI \geq j}$$

and its variance as

$$\text{Var}(\tilde{F}_t^{PTI \geq j}) = \text{Var}(\widehat{F}_t^{PTI \geq j}) + \text{Var}(\varepsilon_t^j)$$

where $\text{Var}(\varepsilon_t^j)$ is computed as the sample variance of ε_t^j, which is in turn computed as the aggregated value of $\varepsilon_{i,t}^j$ for each t, weighted by loan count. The estimates, which are used to create Figure 4, are displayed in Tables B.2 and B.3.

Step 5: Adjust for Households with PTI Between 46 and 50. To go from the missing mass above 45 to the full share constrained by PTI, I need to account for the loans with $PTI \in [46, 50]$. From the histograms, it is clear that only a small subset of loans are allowed these higher PTI limits, while the rest must obey $PTI \leq 45$. However, it is not clear which, if any, of these borrowers are PTI constrained. A borrower with $PTI \in [46, 50]$ region is not constrained if the borrower would have otherwise chosen this value, but is constrained if the borrower would have otherwise chosen $PTI \geq 51$. As such, I compute the total fraction constrained by PTI as

$$\tilde{F}_t = \underbrace{\tilde{F}_t^{PTI \geq 46}}_{\text{missing mass } PTI \geq 46} + \zeta \underbrace{(F_t^{PTI \geq 46} - F_t^{PTI \geq 51})}_{\text{actual mass} \in [46,50]} \tag{B.5}$$

where ζ represents the fraction of borrowers in the $[46, 50]$ region who would have chosen $PTI \geq 51$ if possible. For my baseline estimate, I assume that assignment to the higher $PTI \leq 50$ limit is random, in which case the probability that a borrower in this region would have chosen $PTI \geq 51$ absent the limit is simply the share of loans with $PTI \geq 46$ that have $PTI \geq 51$ in the counterfactual:

$$\zeta = \frac{\widehat{F}_t^{PTI \geq 51}}{\widehat{F}_t^{PTI \geq 46}}.$$

Plausible alternative assumptions for ζ lie on the interval $[0, 1]$, generating the minimum and maximum values quoted in the main text.

Step 6: Estimate Joint Constraints. To compute the fractions that are jointly constrained by both LTV and PTI, I re-run (B.4) on $F_{i,t}^{PTI \geq j|LTV}$ and $F_{i,t}^{PTI \geq j|Control}$, as described in Step 1, aggregate over ZIP3 areas i, and compute the conditional missing mass for each as

$$\begin{aligned} \tilde{F}_t^{PTI \geq j|LTV} &= \widehat{F}_t^{PTI \geq j|LTV} - F_t^{PTI \geq j|LTV} \\ \tilde{F}_t^{PTI \geq j|Control} &= \widehat{F}_t^{PTI \geq j|Control} - F_t^{PTI \geq j|Control}. \end{aligned}$$

Note that these still represent conditional probabilities within the "LTV" and "Control" groups, respectively. I then adjust these shares for the population with $PTI \in [46, 50]$ as above to obtain the conditional fraction constrained by PTI within each group. For instance, I define

$$\tilde{F}_t^{PTI|k} = \tilde{F}_t^{PTI \geq 46|k} + \zeta^k (F_t^{PTI \geq 46|k} - F_t^{PTI \geq 51|k})$$
$$\zeta^k = \frac{\widehat{F}_t^{PTI \geq 51|k}}{\widehat{F}_t^{PTI \geq 46|k}}$$

for $k \in \{LTV, Control\}$. Last, to translate from conditional probabilities of being PTI constrained within each group to an overall share of borrowers that are jointly constrained, I compute the final adjusted joint missing mass as:

$$\tilde{F}_t^{both} = \tilde{F}_t^{PTI|LTV} \times F_t^{LTV} - \tilde{F}_t^{PTI|Control} \times F_t^{Control}$$

where F_t^{LTV} is the share of loans in the specified LTV bins, and $F_t^{Control}$ is the share of loans in the adjacent bins, so that weighting by these variables translates the conditional probabilities back to joint probabilities.

Step 7: Estimate Joint Constraint Regressions. To create analogues to Figure 4 for the interacted constraints, I run the regression

$$F_{i,t}^{m,k} = \alpha_i^{m,k} + \beta_1^{m,k} \log e_{i,t}^* + \beta_2^{m,k} F_{i,t}^{PTI \geq 29} + \varepsilon_{i,t}^{m,k} \tag{B.6}$$

for $m \in \{PTI \leq 45, PTI \geq 46\}$, and $k \in \{\text{LTV Constrained, Not LTV Constrained}\}$. For each m, the joint share that is LTV constrained is computed as

$$F_{i,t}^{m,\text{LTV Constrained}} = F_{i,t}^{m,LTV} - F_{i,t}^{m,Control}$$

using the definitions in Step 1, above, and define

$$F_{i,t}^{m,\text{Not LTV Constrained}} = F_{i,t}^{m} - F_{i,t}^{m,\text{LTV Constrained}}.$$

Aside from this, the regression is otherwise identical to (B.4). Standard errors are computed as in Step 4, and the results are shown in Appendix Figures B.6 and B.7.

B.2.2 Estimation Results

This section contains additional figures and tables associated with Section 3.2. Figure B.4 displays a version of Figure 3 in which we use only a single predictor $\log e^*_{i,t}$ and omit $F^{PTI\geq 29}_{i,t}$. The figure shows that the resulting counterfactual distribution has too much mass in the left portion of the distribution that should not be affected at the margin by PTI constraints, demonstrating the importance of including both predictors. Appendix Tables B.2–B.5 contain full sets of outputs for the constraint regressions used to create Figure 4 and Table 1. Appendix Figure B.5 provides an equivalent to Figure 4 for the $PTI \geq 51$ share, while Appendix Figures B.6 and B.7 provide additional plots for the joint shares (constrained by both, PTI only, LTV only).

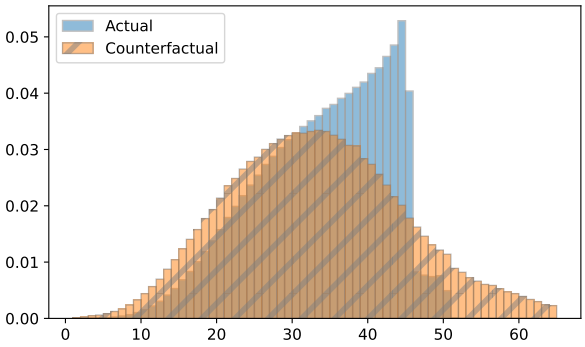

Figure B.4: Actual vs. Counterfactual PTI Distributions, $\log e^*_{i,t}$ as Predictor Only

Note: The figure compares the actual distribution of PTI ratios for purchase loans in the Fannie Mae data for the post-bust sample (2014–2016) against a counterfactual fitted distribution, computed using the distribution regression methodology described in Section 3.2. Unlike in Figure 3, this figure uses only $\log e^*_{i,t}$ as a predictor, while omitting $F^{PTI\geq 29}_{i,t}$.

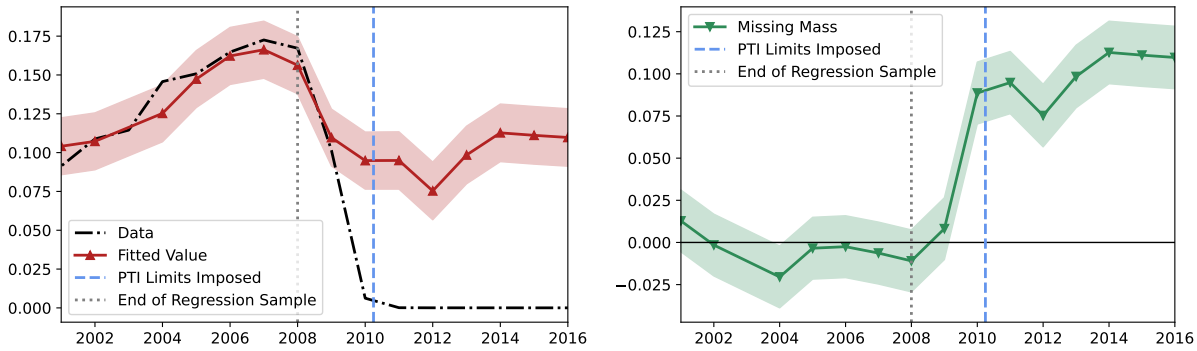

(a) Share PTI ≥ 51

(b) Share PTI ≥ 51

Figure B.5: High PTI Shares, Model vs. Data

Note: The figure displays the time series of the share of purchase loans with PTI ratios ≥ 51. Panel (a) compares the actual share in the data (dash-dotted line) to the counterfactual fitted values (solid line with markers) from the distribution regression. Panel (b) displays the estimated missing mass, computed as the difference between the fitted and actual shares. Shaded regions denote 95% confidence intervals accounting for uncertainty in the regression estimates. The vertical dotted line indicates the end of the regression sample (2007), and the vertical dashed line indicates the imposition of PTI limits (2010).

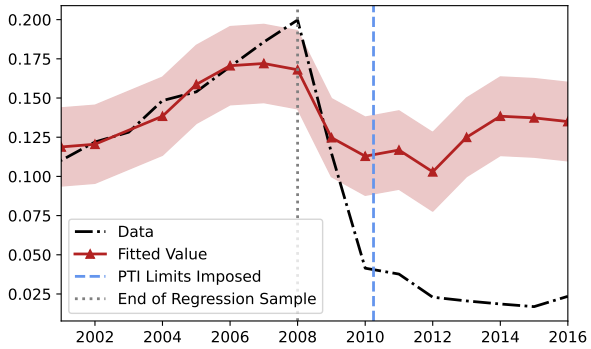

(a) $PTI \geq 46$, LTV Constrained (Purchase)

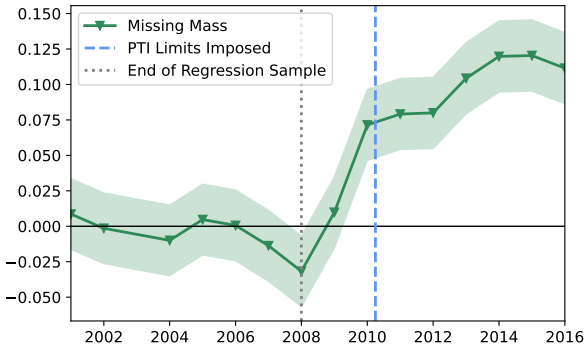

(b) Excess Both Share (Purchase)

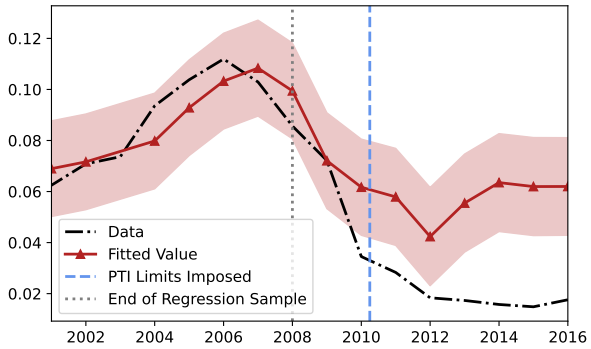

(c) $PTI \geq 46$, Not LTV Constrained (Purchase)

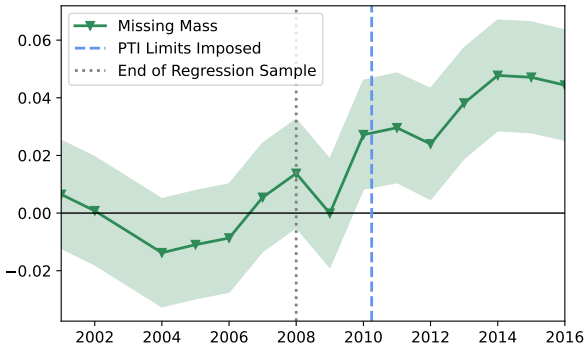

(d) Excess PTI Only Share (Purchase)

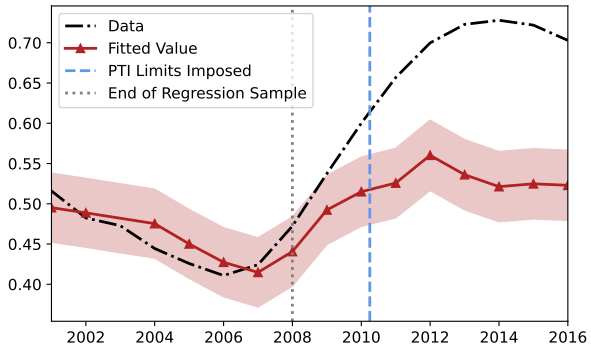

(e) $PTI \leq 45$, LTV Constrained (Purchase)

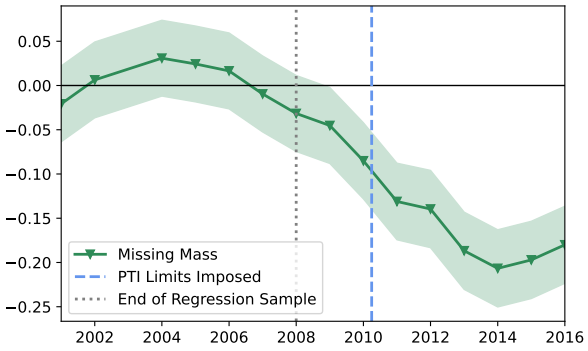

(f) Excess LTV Only Share (Purchase)

Figure B.6: Joint Shares, Model vs. Data

Note: Left panels compare the data to fitted values from regression (B.4), applied to joint probabilities, for the subsample of purchase loans. The right panels plot the differences between actual and fitted values. The vertical dotted line marks the end of the regression sample and the vertical dashed line marks the introduction of formal PTI limits. Shaded bands represent 95% confidence intervals accounting for both parameter and residual uncertainty. Source: Fannie Mae Single Family Loan Performance Dataset.

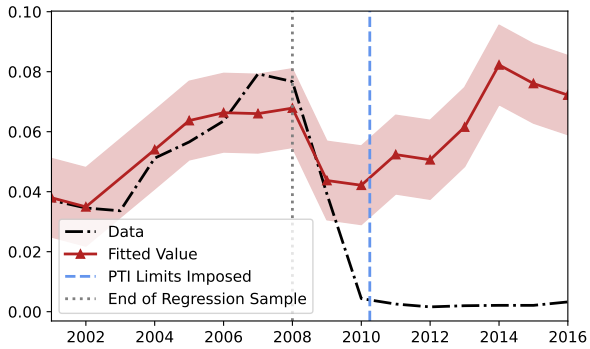

(a) $PTI \geq 46$, LTV Constrained (Cash-Out)

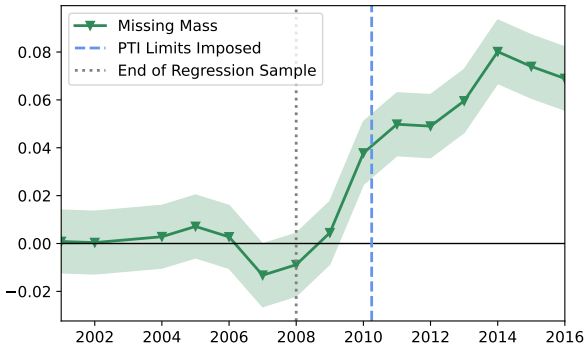

(b) Excess Both Share (Cash-Out)

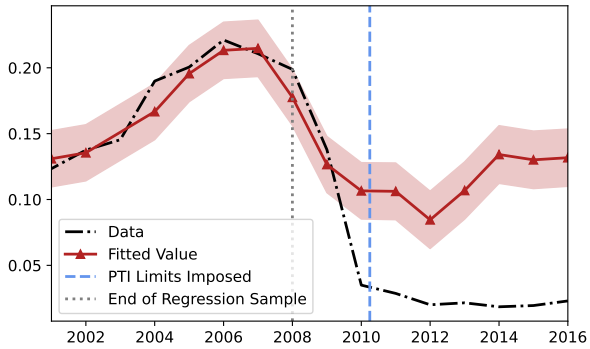

(c) $PTI \geq 46$, Not LTV Constrained (Cash-Out)

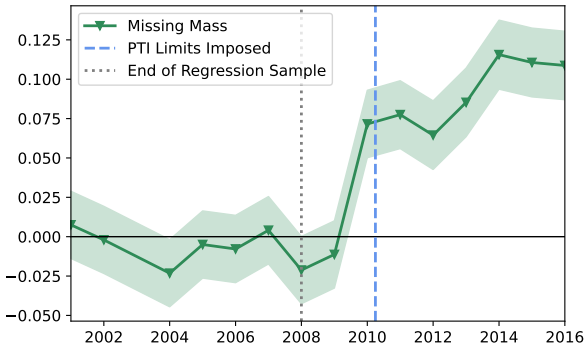

(d) Excess PTI Only Share (Cash-Out)

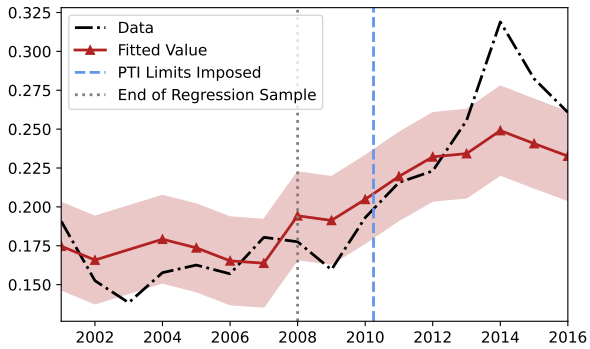

(e) $PTI \leq 45$, LTV Constrained (Cash-Out)

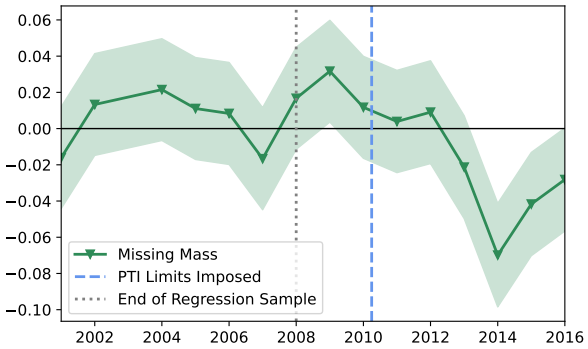

(f) Excess LTV Only Share (Cash-Out)

Figure B.7: Joint Shares, Model vs. Data

Note: Left panels compare the data to fitted values from regression (B.4), applied to joint probabilities, for the subsample of cash-out refinance loans. The right panels plot the differences between actual and fitted values. The vertical dotted line marks the end of the regression sample and the vertical dashed line marks the introduction of formal PTI limits. Shaded bands represent 95% confidence intervals accounting for both parameter and residual uncertainty. Source: Fannie Mae Single Family Loan Performance Dataset.

Table B.2: PTI Share ≥ 46 Regression, Boom Period

	(1) PTI ≥ 46	(2) PTI ≥ 46	(3) PTI ≥ 46	(4) PTI ≥ 46	(5) PTI ≥ 46
$\log(e^*)$	0.118***	0.058***	0.078***	0.309***	
	(0.006)	(0.005)	(0.006)	(0.007)	
PTI ≥ 29	0.653***	0.608***	0.457***		0.792***
	(0.013)	(0.014)	(0.012)		(0.011)
N	4947	4957	4947	4947	8103
Adj. R^2	0.829	0.718	0.894	0.640	0.735
Adj. R^2 (Within)	0.796	0.718	0.491	0.569	0.667
ZIP-3 FE	✓		✓	✓	✓
Time FE			✓		

Note: This table reports the estimates of regression (1) where the dependent variable $F^j_{i,t}$ is the share of purchase loans in 3-digit ZIP code i with PTI ratios ≥ 46 at time t. The predictors are the local PTI fundamental $\log e^*_{i,t}$ (a function of local house prices, mortgage rates, and local wages) and the share of loans with PTI ratios ≥ 29, $F^{PTI\geq 29}_{i,t}$. The regression includes ZIP3 fixed effects and is estimated on Fannie Mae purchase loan data for 2000–2007, the "boom" period in which PTI limits were effectively slack. The regression is weighted using the number of loans in each ZIP3 at time t. Standard errors in parentheses are clustered by ZIP3. The fitted values from this regression are used to construct counterfactual PTI distributions for the post-bust period, as described in Section 3.2.

Table B.3: PTI Share ≥ 51 Regression, Boom Period

	(1) PTI ≥ 51	(2) PTI ≥ 51	(3) PTI ≥ 51	(4) PTI ≥ 51	(5) PTI ≥ 51
$\log(e^*)$	0.080***	0.046***	0.052***	0.205***	
	(0.005)	(0.004)	(0.005)	(0.005)	
PTI ≥ 29	0.426***	0.362***	0.289***		0.515***
	(0.010)	(0.011)	(0.011)		(0.009)
N	4947	4957	4947	4947	8103
Adj. R^2	0.744	0.596	0.835	0.574	0.637
Adj. R^2 (Within)	0.705	0.596	0.349	0.509	0.568
ZIP-3 FE	✓		✓	✓	✓
Time FE			✓		

Note: This table reports the estimates of regression (1) where the dependent variable $F^j_{i,t}$ is the share of purchase loans in 3-digit ZIP code i with PTI ratios ≥ 51 at time t. The predictors, sample, fixed effects, weighting, and clustering are identical to those in Table B.2. The fitted values are used to estimate the share constrained in the $[46, 50]$ PTI interval, i.e. the share of those borrowers who would have exceeded a PTI of 50 absent constraints, as described in Section 3.2.

Table B.4: PTI Share ≥ 46 Conditional Regressions, Boom Period

	(1) PTI ≥ 46	(2) PTI ≥ 46 (LTV)	(3) PTI ≥ 46 (Control)
$\log(e^*)$	0.118***	0.152***	0.156***
	(0.006)	(0.008)	(0.012)
PTI ≥ 29	0.653***	0.662***	0.717***
	(0.013)	(0.014)	(0.029)
N	4947	4943	4820
Adj. R^2	0.829	0.800	0.495
Adj. R^2 (Within)	0.796	0.763	0.462
ZIP-3 FE	✓	✓	✓
Time FE			

Note: This table reports estimates of regression (1) using three dependent variables: (1) the unconditional share of loans with PTI ≥ 46; (2) the share of loans with PTI ≥ 46 *among* those at a standard LTV limit (i.e. the joint share divided by the share at an LTV limit); and (3) the share of loans with PTI ≥ 46 *among* those in LTV "control" bins (near but not at a standard LTV limit). Columns (2) and (3) decompose the high-PTI share by LTV constraint status, showing whether borrowers at and away from LTV limits are similarly affected by PTI constraints. The specification, sample, and standard errors are identical to those in Table B.2.

Table B.5: PTI Share ≥ 51 Conditional Regressions, Boom Period

	(1) PTI ≥ 51	(2) PTI ≥ 51 (LTV)	(3) PTI ≥ 51 (Control)
$\log(e^*)$	0.080***	0.101***	0.091***
	(0.005)	(0.006)	(0.010)
PTI ≥ 29	0.426***	0.425***	0.488***
	(0.010)	(0.011)	(0.023)
N	4947	4943	4820
Adj. R^2	0.744	0.707	0.362
Adj. R^2 (Within)	0.705	0.666	0.344
ZIP-3 FE	✓	✓	✓
Time FE			

Note: This table is identical to Table B.4 but uses a PTI threshold of 51 instead of 46, corresponding to the threshold used in Table B.3.

B.3 Links Between PTI Ratios and House Price Growth

This section provides additional results and detail for the analysis in Section 3.3. To begin, Table B.6 displays regression results for the fit lines in Figure 5.

Table B.6: House Price Growth vs. High-PTI Share, Census Data

	PUMA	State
Constant	0.107***	0.043
	(0.010)	(0.064)
Share PTI $>$ 28	1.482***	1.880***
	(0.029)	(0.238)
N	1,708	50
Adj. R^2	0.573	0.757

Note: This table reports estimates of weighted least squares regressions at the PUMA and state levels. The dependent variable is log growth in the FHFA house price index from 1998 to 2006. The PUMA-level house price index is constructed by collapsing the FHFA all-transactions 5-digit ZIP code house price indices to the PUMA level using the 1990 crosswalk. The main independent variable is the 1990 Census PUMS share of households with housing payment-to-income ratios above 28% among households with positive mortgage payments and income that moved within the last year, where housing payments include first and second mortgage payments, property insurance, and property taxes. Column (1) reports PUMA-level estimates; column (2) reports state-level estimates. Regressions are weighted using summed 1990 Census PUMS household weights over all households by PUMA or state. Standard errors in parentheses are heteroskedasticity-robust.

I next provide additional detail on the regressions incorporating additional controls in Table 2. These regressions are motivated by the natural concern that the initial share with PTI $>$ 28 is correlated with demand factors unrelated to credit standards, such as local income or income growth. To this end, Table 2 repeats the PUMA-level cross-sectional regression including additional PUMA-level controls. To set a baseline, column (1) reproduces the original result. Column (2) adds a control for average income in the 1990 Census PUMS, which has a minimal effect on the coefficient of interest. Column (3) adds a control for the contemporaneous *change* in log wage income from the same 1998–2006 period used to measure house price growth. While income growth unsurprisingly has a positive relationship with contemporaneous house price growth, the main coefficient on the share with PTI $>$ 28 is only modestly diminished and remains highly significant.

Next, column (4) adds a control for pre-period log house price growth in the FHFA data, designed to capture the well-known autocorrelation in house price growth (Guren, 2018). For this variable, the pre-period is chosen as 1993–1998, since granular FHFA cov-

erage drops off rapidly prior to 1993. While this variable has a highly significant negative impact on boom house price growth, consistent with the typical finding of long-run reversal, it has only a modest effect on the main coefficient.

Last, column (5) adds contemporaneous boom rental growth, measured as the log difference between median gross rent in the 2000 Decennial Census and the 2005–2009 5-year American Community Survey. This provides a powerful control for any factors affecting housing preferences, fundamentals, or supply elasticity, that would affect the price of housing services (rent) but are not driven by a relaxation of PTI limits. At the same time I note that this variable is likely an overcontrol, as boom house price growth may cause boom rental growth through the user cost of housing, and thus may absorb valid variation. Column (5) shows that incorporating this factor indeed noticeably changes the coefficients on some of the other controls, particularly pre-period house price growth. However, while adding this control slightly attenuates the main coefficient as expected, it still retains nearly 85% of its value compared to the completely uncontrolled specification in column (1), and remains highly significant.

To summarize, while many of these controls are associated with house price growth, and often in intuitive ways, the main results survive the inclusion of a powerful set of controls for a variety of demand factors with only modest attenuation in the main coefficient. The results remain highly significant in all specifications, with heteroskedasticity-robust t-statistics consistently in excess of 19. I conclude that the core relationship between the 1990 high-PTI share and house price growth from 1998–2006 is highly robust to controls for non-credit fundamentals.

For further robustness, I use two alternative data sources for the share of high-PTI loans. First, I repeat the exercise using loans from the 1990–1995 HMDA data. While these data do not include a number of components of the PTI ratio, such as the exact interest rate, property taxes, or property insurance payments, they are measured at origination, and effectively cover the universe of new mortgages over this period, and thus provide a valuable robustness check. I impute these missing components using national averages or standard benchmarks and filter out likely second liens (see Appendix B.1.3). I then re-estimate each cross-sectional regression, this time at the county level at which the HMDA data are geographically coded.

The results are displayed in Appendix Figure B.8 and Appendix Table B.7. These results show that the strong relationship between pre-period PTI ratios and boom house price growth is preserved. Appendix Table B.7 repeats each specification from Table 2,

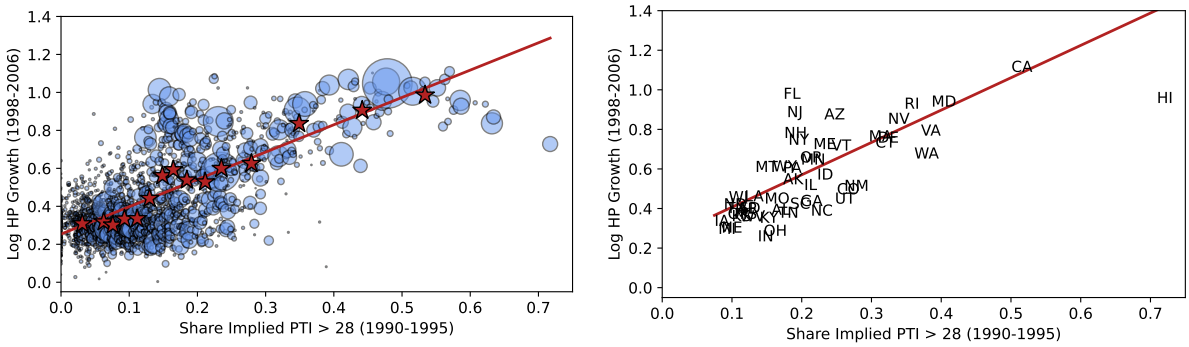

(a) HMDA: County Data

(b) HMDA: State Data

Figure B.8: High-PTI Share vs. House Price Growth

Note: Each circle represents one county in Panel (a) and each pair of letters represents one state in Panel (b). The y-axis variable is log growth in the FHFA house price index from 1998 to 2006. The x-axis variable is the share of purchase loans with implied PTI ratios exceeding 28% in the HMDA data for years 1990–1995 (see Appendix B.1.3). The line in each panel reflects fitted values from a regression of the y-axis variable on the x-axis variable and a constant, weighted by the number of households in the 2000 Census. The size of each circle is proportional to the number of households in the 2000 Census. The stars in Panel (a) represent a binscatter of averages, weighted by the number of households in the 2000 Census, using 15 equally sized bins.

with the exception that pre-period income is measured using initial log wage income per return in the IRS data over the 1990–1995 pre-period in place of 1990 Census PUMS income. The table shows that these results are similar in magnitude and predictive power, and are again robust to the inclusion of demand controls.

Second, I repeat the exercise at the ZIP3 level using Fannie Mae data from 2014–2016, a period with stable and binding PTI limits. While these data are taken from a different time period, they have the advantage that they use exactly measured PTI ratios at origination that account for the actual mortgage interest rate, taxes and insurance, and outside debt payments. Although these data are measured after the boom rather than before, they may still provide a measure of pre-boom incidence of PTI constraints to the extent that the fundamental characteristics leading to geographic variation in PTI ratios are persistent across time.

To this end, I construct an alternative proxy using the share of loans with PTI in excess of 40%, with the higher cutoff being appropriate given the higher PTI limits in this period and the use of the true PTI ratio that includes payments on outside debt. I show in Appendix Figure B.10 that this share aligns closely with the HMDA implied PTI $>$ 28 share over this period. The results for this variable are displayed in Appendix Figure B.9. The figure shows that we once again obtain a remarkably close fit of boom-era house price

Table B.7: House Price Growth vs. Implied High-PTI Share, County-Level Robustness

	(1)	(2)	(3)	(4)	(5)
Constant	0.254***	0.618***	-0.120	0.126	-0.464***
	(0.018)	(0.156)	(0.166)	(0.169)	(0.158)
Share Implied PTI > 28	1.439***	1.495***	1.171***	1.045***	0.956***
	(0.091)	(0.096)	(0.130)	(0.105)	(0.083)
Pre-Period Wage Income		-0.114**	0.036	0.008	0.102**
		(0.050)	(0.049)	(0.047)	(0.042)
Boom Wage Income Growth			1.481***	1.282***	0.816***
			(0.160)	(0.166)	(0.144)
House Price Growth (1993-1998)				-0.565***	-0.245***
				(0.103)	(0.085)
Boom Rental Growth					1.113***
					(0.100)
N	2,017	2,017	2,017	1,974	1,974
Adj. R^2	0.559	0.566	0.648	0.683	0.739

Note: This table reports estimates of weighted least squares regressions at the county level of log growth in the FHFA house price index from 1998 to 2006 on the pre-period share of HMDA mortgage applications with implied PTI ratios above 28%. The implied PTI ratio is constructed from loan balances, national mortgage rates, and borrower income (see Section B.1.3 for details). Pre-Period Wage Income is log wage and salary income per tax return, averaged over 1990–1995 in the IRS data. Boom Wage Income Growth is the change in log wage and salary income per return from 1998 to 2006 in the IRS data. House Price Growth (1993–1998) is the log change in the FHFA house price index from 1993 to 1998. Boom Rental Growth is the log difference in median gross rent between the 2005–2009 5-year American Community Survey and the 2000 Decennial Census. Regressions are weighted by household count in the 2000 Census. Standard errors in parentheses are heteroskedasticity-robust.

growth using this alternative variable.

These results are formalized in Appendix Table B.8, which reproduces Appendix Table B.7 for this case. To control for confounds with the independent variable, it seems more appropriate to control for post-period income and house prices, while leaving boom-era house price growth unchanged. We once again observe a very high R^2 and coefficients that are only modestly affected by controls.

B.4 Characteristics by PTI Ratio

In this section, I document how borrower characteristics vary with the PTI ratio, and connect these results to the literature on cross-sectional patterns of credit growth during the housing boom.

For background, a number of notable papers including Adelino, Schoar, and Severino

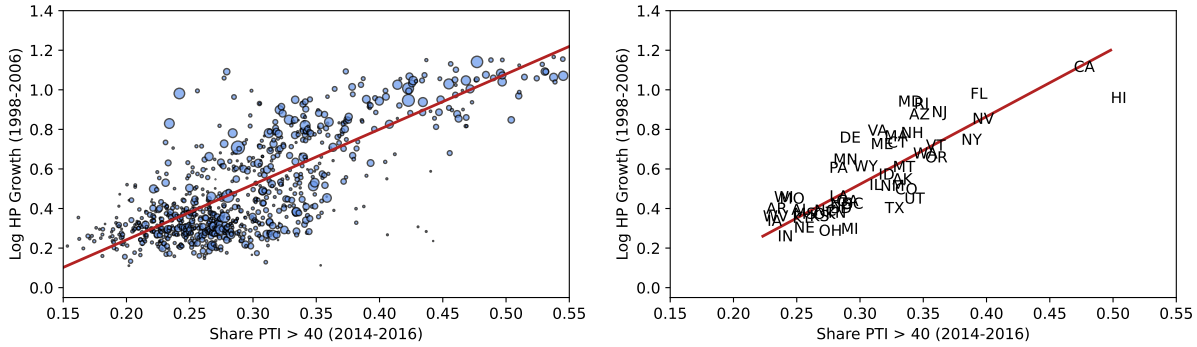

(a) Fannie: ZIP3 Data

(b) Fannie: State Data

Figure B.9: Fannie Mae High-PTI Share vs. House Price Growth, Additional Plots

Note: Each circle represents one ZIP3 in Panel (a) and one state in Panel (b). The y-axis variable is log growth in the FHFA house price index from 1998 to 2006. The x-axis variable is the share of purchase loans with PTI ratios exceeding 40% in the Fannie Mae data for years 2014–2016 (see Appendix B.1.3). The line in each panel reflects fitted values from a regression of the y-axis variable on the x-axis variable and a constant, weighted by the number of households in the 2000 Census. The size of each circle is proportional to the number of households in the 2000 Census. The stars in Panel (a) represent a binscatter of averages, weighted by the number of households in the 2000 Census, using 15 equally sized bins.

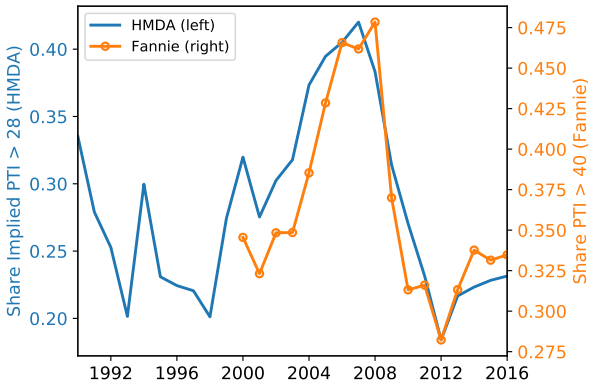

Figure B.10: Fannie vs. HMDA (Implied) High-PTI Share

Note: The figure displays the national time series of the share of HMDA purchase loans with an implied PTI ratio above 28% against the share of Fannie Mae purchase loans with a recorded PTI ratio above 40%. The HMDA data exclude likely second liens as described in Appendix B.1.3. Both series are aggregated to an annual frequency, weighted by household count in the 2000 Census.

(2016), Albanesi, DeGiorgi, and Nosal (2022), and Foote, Loewenstein, and Willen (2021), have documented that growth in mortgage credit during the 2000s housing boom was disproportionately among "prime" borrowers with higher incomes and credit scores. These results point to a broad-based expansion of credit rather than one concentrated among low-income or subprime borrowers.

To test whether this paper's hypothesis of a PTI-based expansion in credit standards

Table B.8: House Price Growth vs. High-PTI Share, Fannie Mae Data

	(1)	(2)	(3)	(4)	(5)
Constant	-0.317***	-0.968***	-0.963***	-0.659***	-1.348***
	(0.035)	(0.186)	(0.163)	(0.119)	(0.112)
Share PTI > 40	2.792***	2.656***	2.434***	2.251***	1.960***
	(0.099)	(0.118)	(0.122)	(0.118)	(0.113)
2014-2016 Wage Income		0.183***	0.145***	0.111***	0.230***
		(0.055)	(0.046)	(0.031)	(0.027)
Boom Wage Income Growth			0.861***	0.693***	0.404***
			(0.126)	(0.110)	(0.079)
House Price Growth (1993-1998)				-0.490***	-0.271***
				(0.099)	(0.088)
Boom Rental Growth					1.198***
					(0.113)
N	867	867	867	849	849
Adj. R^2	0.607	0.634	0.682	0.717	0.770

Note: This table reports estimates of weighted least squares regressions at the 3-digit ZIP code (ZIP3) level. The dependent variable is log growth in the FHFA all-transactions house price index from 1998 to 2006, measured directly at the ZIP3 level. The main independent variable is the share of Fannie Mae purchase loans with recorded PTI ratios above 40%, averaged over the 2014–2016 period. Pre-Period Wage Income is log wage and salary income per tax return, averaged over 2014–2016 in the IRS data, collapsed from ZIP-level to ZIP3 level using households counts from the 2000 Census. Boom Wage Income Growth is the change in log wage and salary income per return from 1998 to 2006 in the same IRS data. House Price Growth (1993–1998) is the log change in the FHFA house price index from 1993 to 1998. Boom Rental Growth is the log difference in median gross rent between the 2005–2009 5-year American Community Survey and the 2000 Decennial Census, aggregated from census tract to ZIP3 level using the 2000 crosswalk. Regressions are weighted by household count in the 2000 Census. Standard errors in parentheses are heteroskedasticity-robust.

fits these patterns, I plot a number of characteristics by PTI limit. I restrict attention to the two main types of constrained borrowers: borrowers constrained by LTV only, and borrowers jointly constrained by LTV and PTI. To proxy for LTV constraints, I restrict attention to the set of borrowers with $LTV_{i,t} \in S = \{70, 75, 80, 85, 90, 95, 97\}$, as defined earlier. To proxy for PTI constraints, I compare borrowers with $PTI_{i,t} < 40$ for the "Low PTI" group, and $PTI_{i,t} \geq 40$ for the "High PTI" group, with the caveat that borrowers with $PTI_{i,t} \geq 40$ are merely more likely to be PTI constrained, and that this is not a strict definition of a PTI-constrained borrower. For the sample, I focus on the 2014–2016 period that featured stable house prices and consistently enforced PTI limits.

Figure B.11a displays the results. For notation, "(Demeaned)" indicates that I have removed the unconditional mean, in logs, for all purchase loans within the same month.

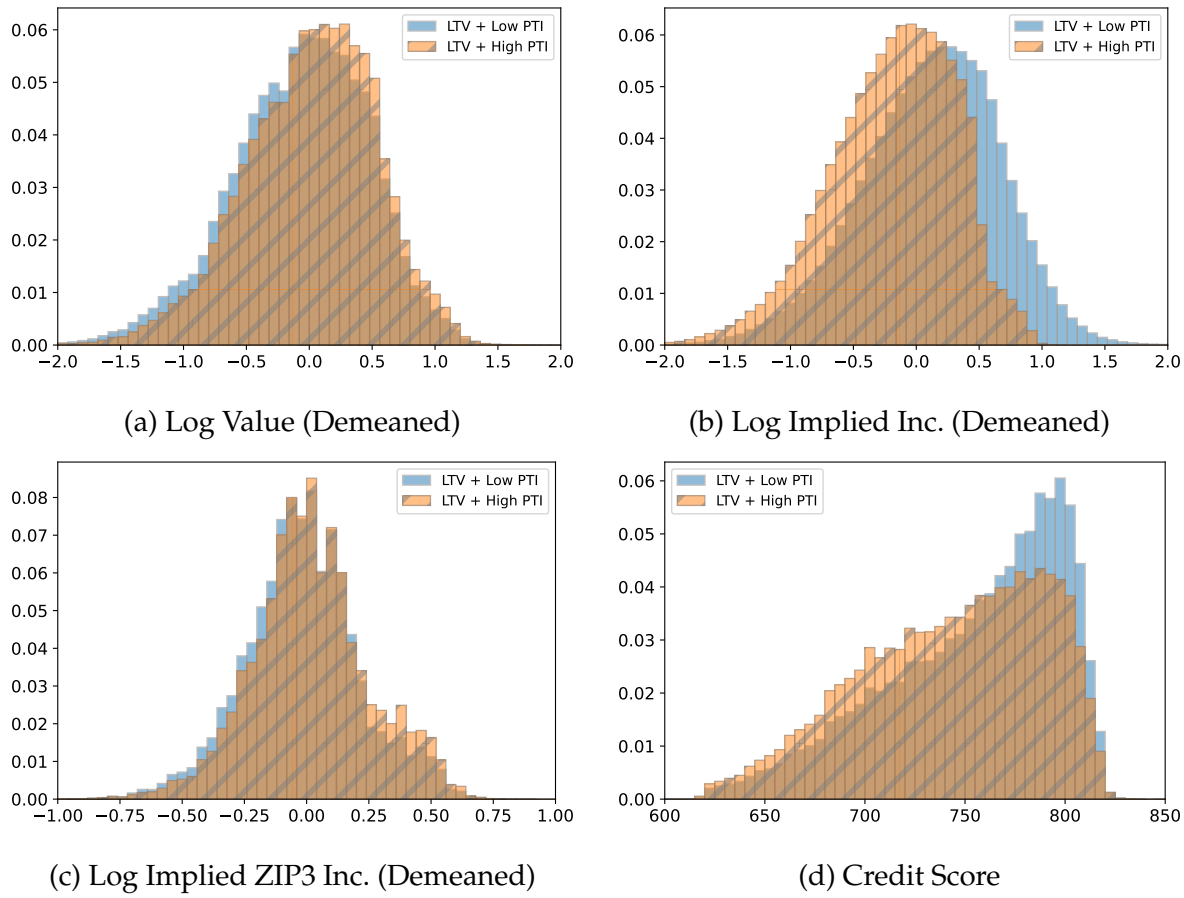

(a) Log Value (Demeaned)

(b) Log Implied Inc. (Demeaned)

(c) Log Implied ZIP3 Inc. (Demeaned)

(d) Credit Score

Figure B.11: Characteristics by PTI Ratio, 2014–2016

Note: Each panel displays overlapping histograms comparing the distributions of borrower characteristics for two groups of Fannie Mae purchase loan originations from 2014–2016: loans at a standard LTV limit (CLTV $\in \{70, 75, 80, 85, 90, 95, 97\}$) with PTI ratios below 40 ("LTV + Low PTI", unhatched bars), and loans simultaneously at a standard LTV limit and with PTI ratios at or above 40 ("LTV + High PTI", hatched bars). Implied house value is computed as the unpaid principal balance divided by the LTV ratio. Implied income is computed from the loan balance, the mortgage coupon rate, and the PTI ratio, see Appendix B.1.2. Panels (a)–(c) display demeaned log variables, where the period mean (computed at the monthly level) has been subtracted prior to plotting. Panel (d) displays credit scores. Source: Fannie Mae Single Family Loan Performance Data.

To begin, Panel (a) compares the log demeaned values of the properties purchased by these groups. This panel shows nearly complete overlap between the two distributions, showing that the two groups are not purchasing notably different house values.

Next, Panel (b) compares log demeaned implied incomes between the two groups, with implied income computed as described in Appendix B.1.2. Here we do see a nontrivial difference, with the high-PTI group having lower incomes on average compared to the low-PTI group. However, there is substantial overlap between the groups, and substantial fractions of both high-PTI borrowers with above-average incomes, and low-PTI borrow-

ers with below-average incomes. I further note that because implied income is computed using the PTI ratio in a way that does not directly account for outside debt obligations, this procedure is mechanically biased toward finding discrepancies in implied income to the extent that borrowers vary in their non-mortgage debt payments. Thus, while there are real differences here, the effect of PTI limits was not overwhelmingly concentrated among low-income borrowers for this period.

To understand geographic variation, Panel (c) displays the distribution of average log implied income for borrowers in that ZIP3, rather than the individual level. This plot thus provides information on whether high-PTI borrowers are concentrated in high-income or low-income geographic areas. This panel shows that the distribution of local borrower incomes looks identical for high-PTI and low-PTI borrowers. This is important evidence, since Adelino, Schoar, and Severino (2016) find that credit growth was not concentrated in low-income ZIP codes over the housing boom period.

Last, Panel (d) compares the distribution of credit scores across these groups. This panel shows that while low-PTI borrowers have somewhat higher credit scores, the two distributions largely overlap, and there is not strong segmentation by credit score.

Summarizing, I find that the incidence of PTI limits, proxied here by high PTI ratios, is broad-based, in the sense that it affects borrowers with similar house values, somewhat but not drastically lower income, similar local average income, and similar credit scores. As a result, I conclude that a relaxation of PTI limits could lead to a broad-based expansion of credit that is not concentrated among low-income, low-credit-score borrowers, consistent with the results of Adelino, Schoar, and Severino (2016).

B.5 Additional Empirical Results

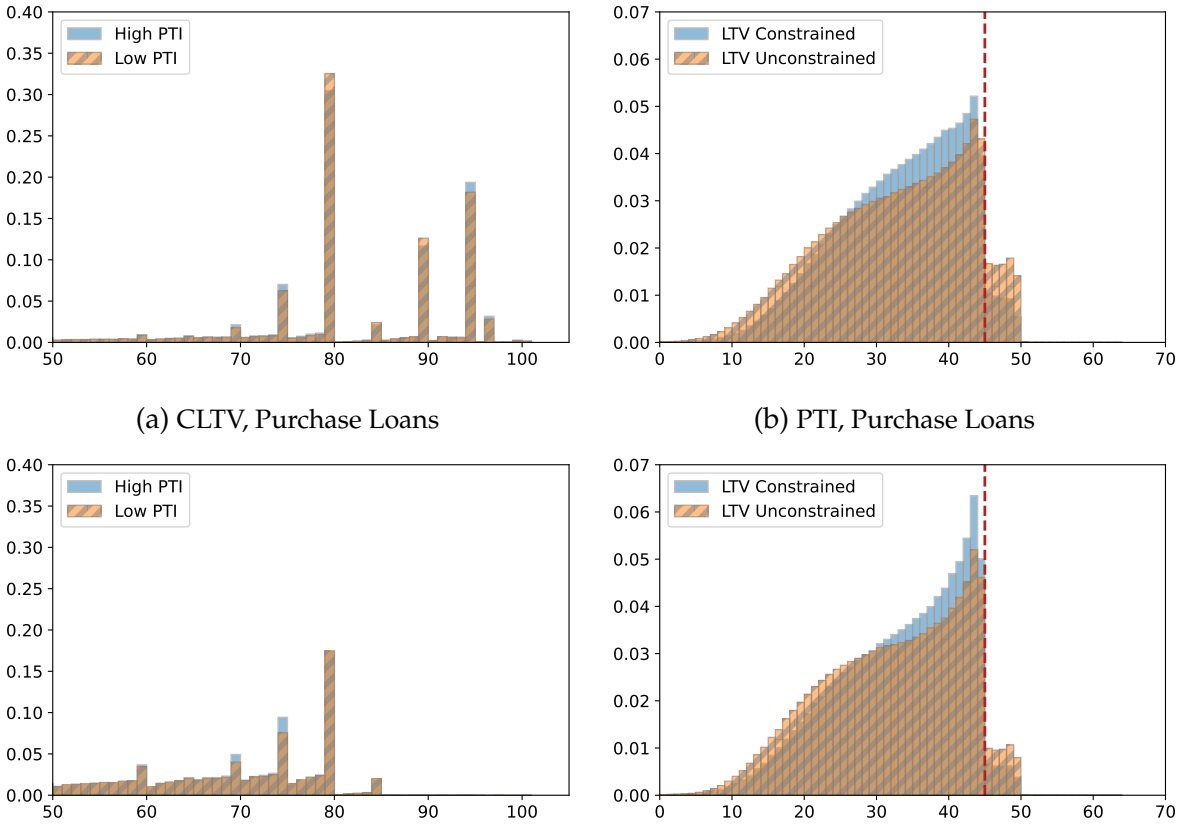

(a) CLTV, Purchase Loans

(b) PTI, Purchase Loans

(c) CLTV, Cash-Out Refinances

(d) PTI, Cash-Out Refinances

Figure B.12: LTV and PTI Ratios, by Constraint Proxy, Post-Boom Period

Note: Each panel displays overlapping histograms comparing the distributions of CLTV (left panels) and PTI ratios (right panels) for two groups of Fannie Mae loan originations from the post-boom period (April 2010 onward). The "LTV Constrained" group (unhatched bars) consists of loans with CLTV ratios in $\{60, 70, 75, 80, 85, 90, 95, 97\}$, while the "LTV Unconstrained" group consists of all remaining loans of that type. Similarly, the "High PTI" group (unhatched bars) consists of loans with PTI ratios of 40 or more, while the "Low PTI" group (hatched bars) consists of all remaining loans. Panels (a)–(b) display purchase loans (first-time and repeat buyers combined); panels (c)–(d) display cash-out refinances. The vertical dashed line marks the typical PTI limit of 45 for this period. Histograms are weighted by original principal balance. Source: Fannie Mae Single Family Loan Performance Data.

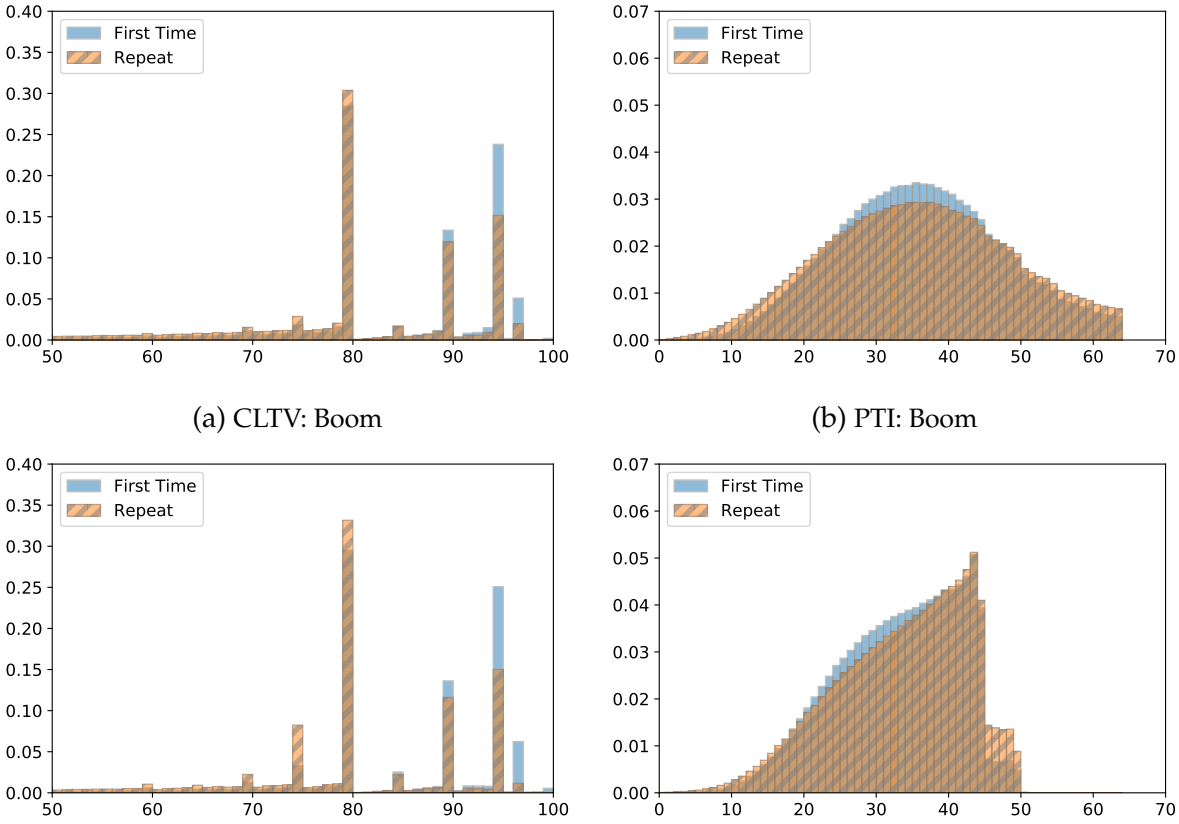

(a) CLTV: Boom

(b) PTI: Boom

(c) CLTV: Post-Boom

(d) PTI: Post-Boom

Figure B.13: LTV and PTI Ratios, First-Time vs. Repeat Buyers

Note: Each panel shows balance-weighted histograms of the combined loan-to-value (CLTV) ratio and payment-to-income (PTI) ratio, comparing first-time buyers ("F" loan purpose) against repeat buyers ("P" loan purpose). The top row uses the boom-era sample (originations through 2008) and the bottom row uses the post-bust sample (originations from April 2010 onward). Left panels show the CLTV distribution; right panels show the PTI distribution. Source: Fannie Mae Single Family Loan Performance Dataset.

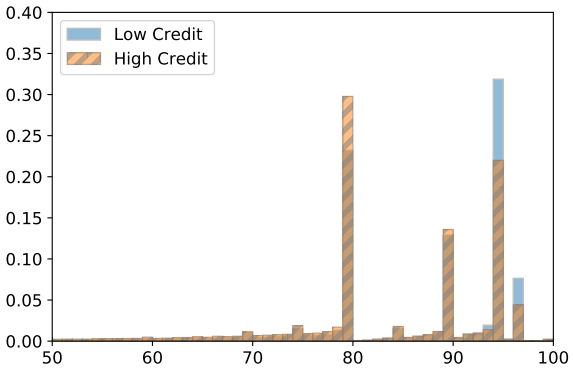

(a) CLTV: First-Time Buyers

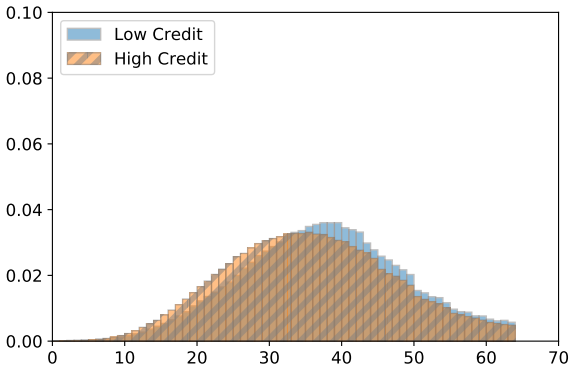

(b) PTI: First-Time Buyers

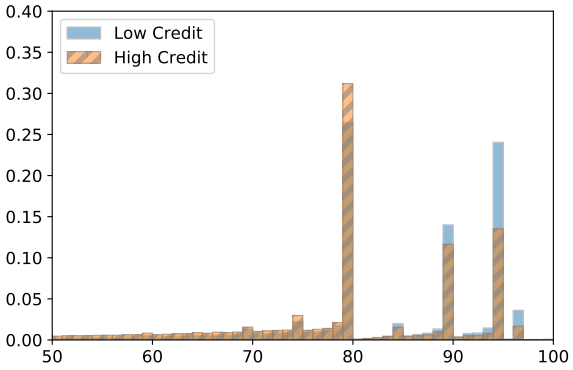

(c) CLTV: Repeat Buyers

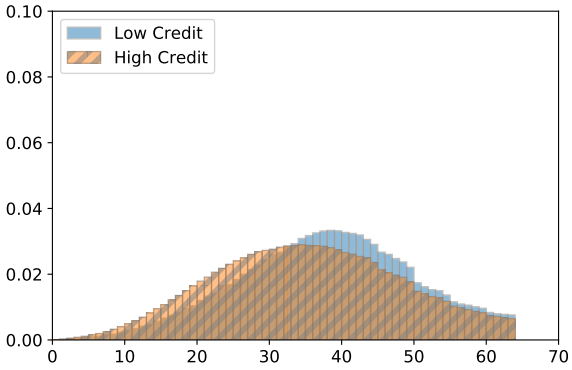

(d) PTI: Repeat Buyers

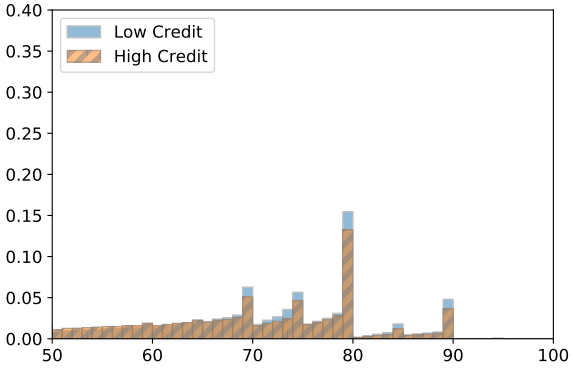

(e) CLTV: Cash-Out Refi

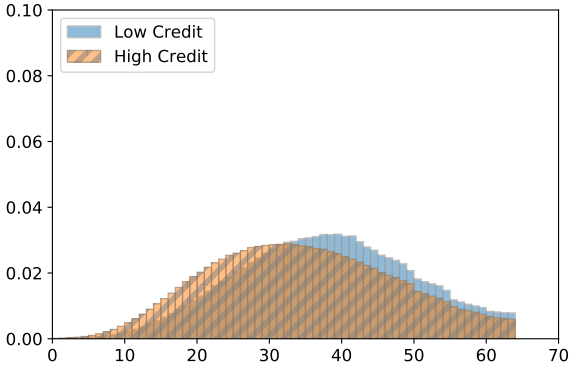

(f) PTI: Cash-Out Refi

Figure B.14: LTV and PTI Ratios by Credit Score (Boom Period)

Note: Each panel shows balance-weighted histograms of the combined loan-to-value (CLTV) ratio and payment-to-income (PTI) ratio, comparing near-prime borrowers (credit scores 620–679) against prime borrowers (credit scores $\geq$ 680), using loans originated over the 2000–2008 sample. Rows correspond to first-time buyers, repeat buyers, and cash-out refinances, respectively. Source: Fannie Mae Single Family Loan Performance Dataset.

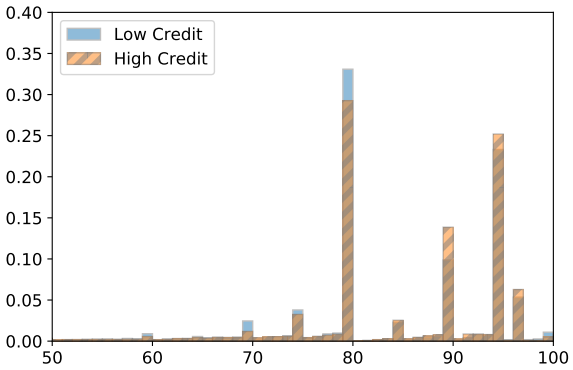

(a) CLTV: First-Time Buyers

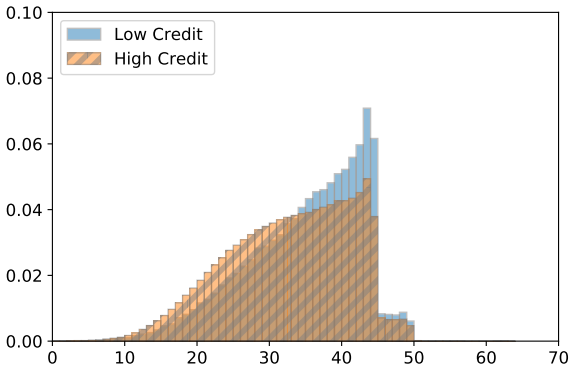

(b) PTI: First-Time Buyers

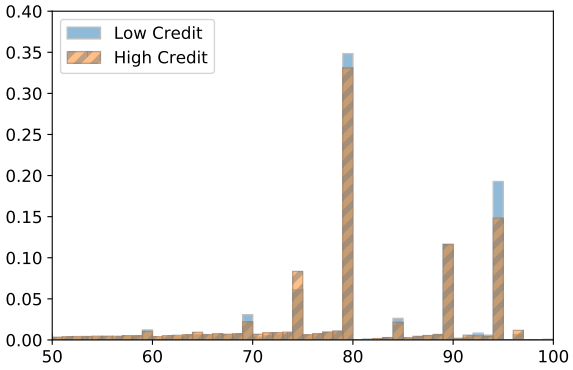

(c) CLTV: Repeat Buyers

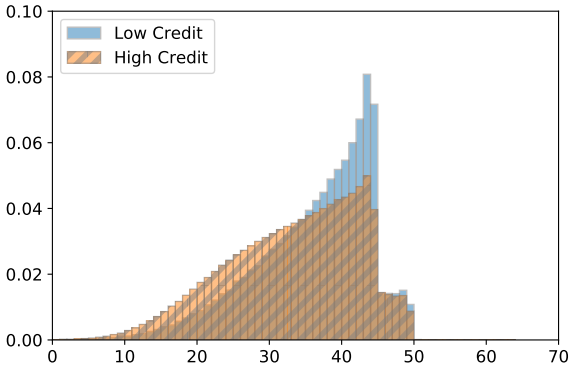

(d) PTI: Repeat Buyers

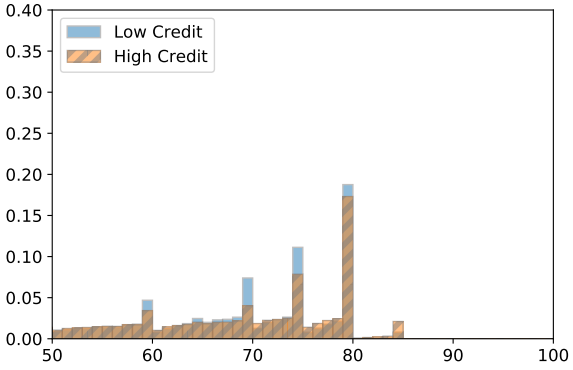

(e) CLTV: Cash-Out Refi

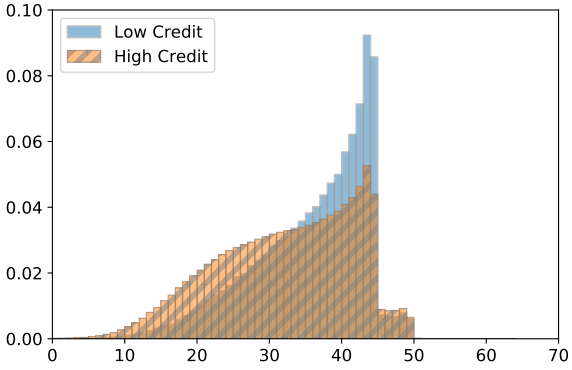

(f) PTI: Cash-Out Refi

Figure B.15: LTV and PTI Ratios by Credit Score (Bust Period)

Note: Each panel shows balance-weighted histograms of the combined loan-to-value (CLTV) ratio and payment-to-income (PTI) ratio, comparing near-prime borrowers (credit scores 620–679) against prime borrowers (credit scores $\geq$ 680), using loans originated over the 2010–2016 sample. Rows correspond to first-time buyers, repeat buyers, and cash-out refinances, respectively. Source: Fannie Mae Single Family Loan Performance Dataset.

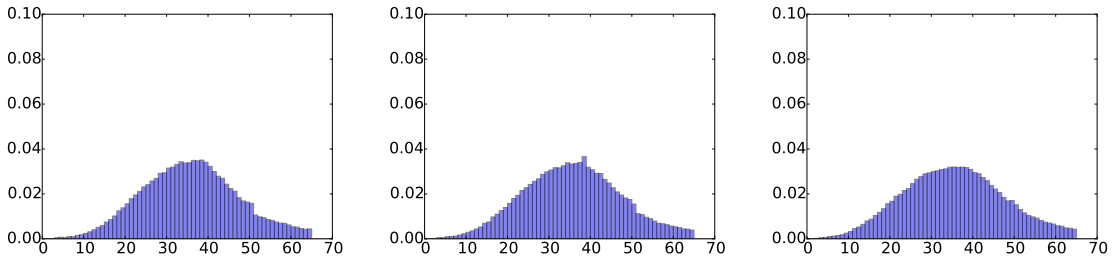

(a) PTI: 2000 Q1

(b) PTI: 2001 Q1

(c) PTI: 2002 Q1

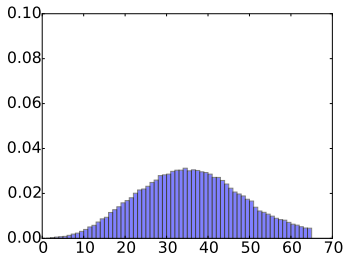

(d) PTI: 2003 Q1

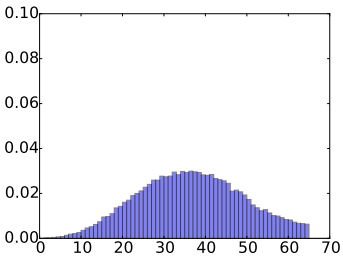

(e) PTI: 2004 Q1

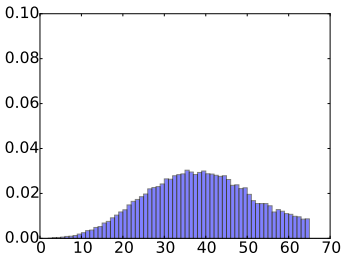

(f) PTI: 2005 Q1

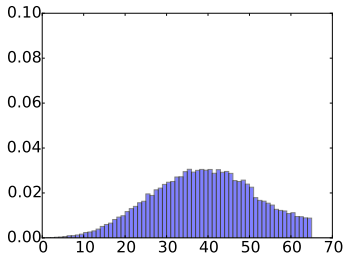

(g) PTI: 2006 Q1

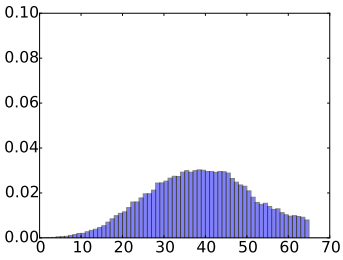

(h) PTI: 2007 Q1

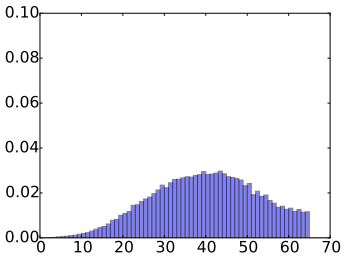

(i) PTI: 2008 Q1

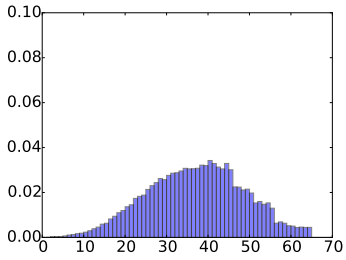

(j) PTI: 2009 Q1

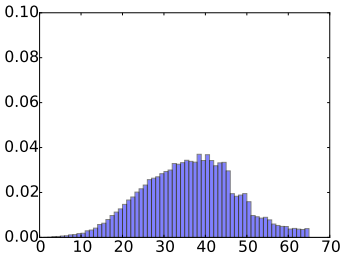

(k) PTI: 2010 Q1

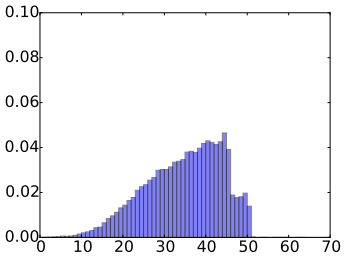

(l) PTI: 2011 Q1

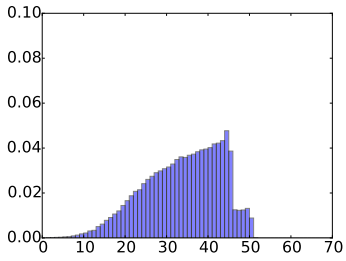

(m) PTI: 2012 Q1

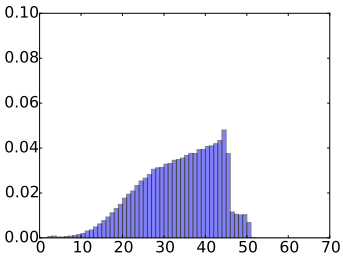

(n) PTI: 2013 Q1

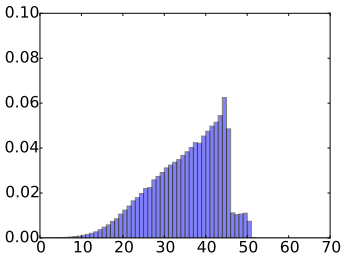

(o) PTI: 2014 Q1

Figure B.16: PTI, Newly Originated FNMA Purchase Loans, Additional Years

Note: Histograms are weighted by loan balance. Source: Fannie Mae Single Family Dataset.

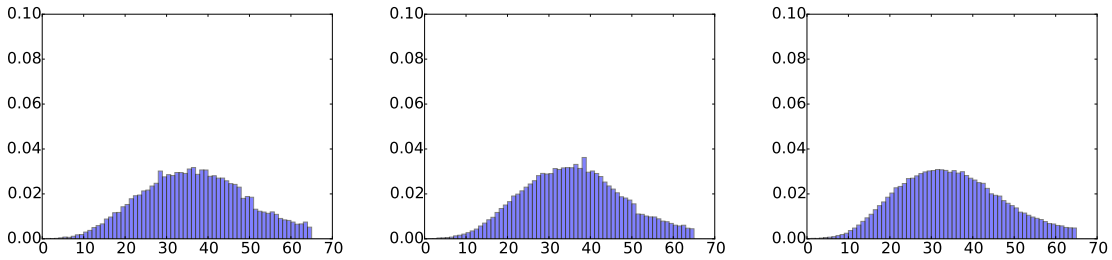

(a) PTI: 2000 Q1 (b) PTI: 2001 Q1 (c) PTI: 2002 Q1

(d) PTI: 2003 Q1 (e) PTI: 2004 Q1 (f) PTI: 2005 Q1

(g) PTI: 2006 Q1 (h) PTI: 2007 Q1 (i) PTI: 2008 Q1

(j) PTI: 2009 Q1 (k) PTI: 2010 Q1 (l) PTI: 2011 Q1

(m) PTI: 2012 Q1 (n) PTI: 2013 Q1 (o) PTI: 2014 Q1

Figure B.17: PTI, Newly Originated FNMA Cash-Out Refi Loans, Additional Years

Note: Histograms are weighted by loan balance. Source: Fannie Mae Single Family Dataset.

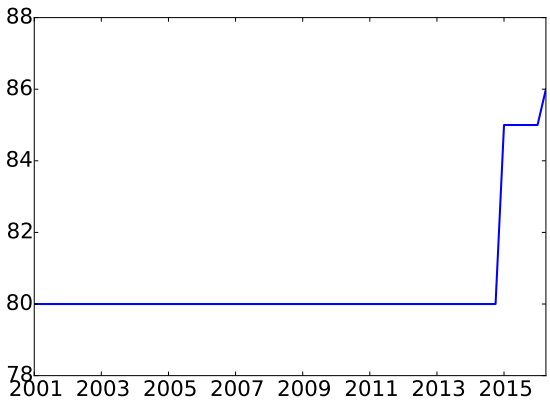

(a) CLTV: 50th Percentile

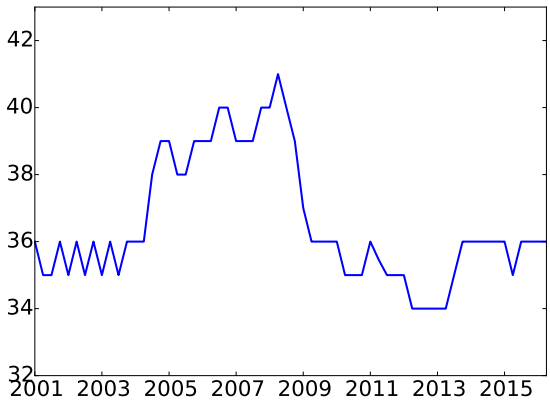

(b) PTI: 50th Percentile

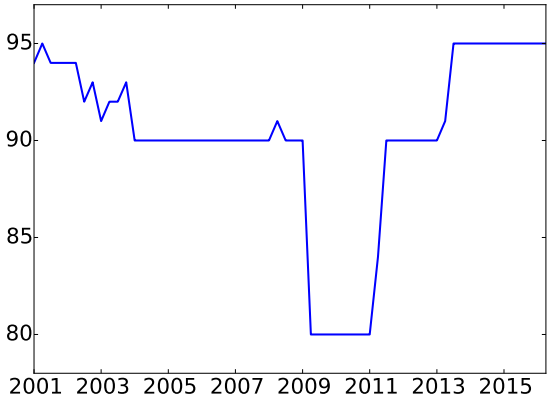

(c) CLTV: 75th Percentile

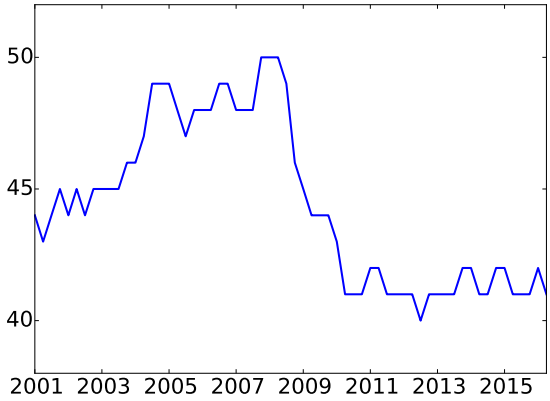

(d) PTI: 75th Percentile

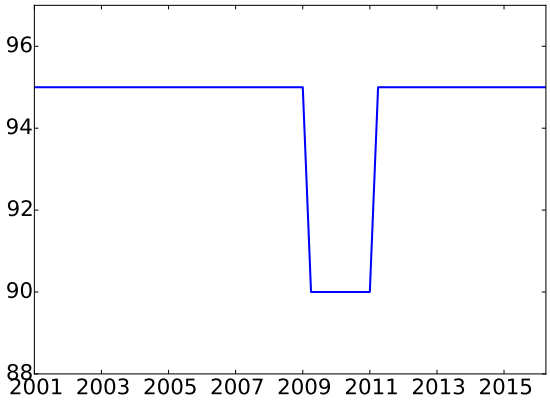

(e) CLTV: 90th Percentile

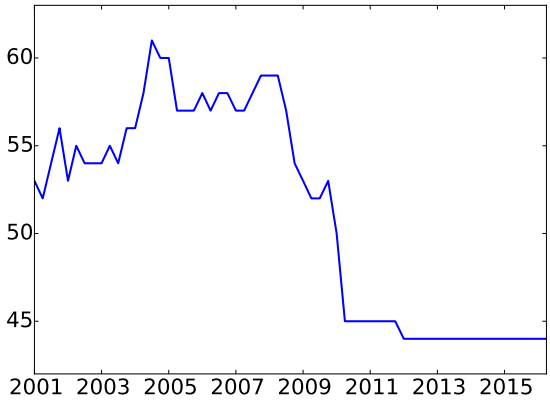

(f) PTI: 90th Percentile

Figure B.18: CLTV and PTI Percentiles, Newly Originated FNMA Purchase Loans

Note: Plots report percentiles weighted by loan balance. Source: Fannie Mae Single Family Dataset.

C Simple Numerical Example: Details

This section provides a quantitative version of the simple example of Section 2, which is inspired by Justiniano, Primiceri, and Tambalotti (2019), and closely follows Allen and Greenwald (2022). The agent makes a one-time choice over housing quantity to maximize

$$V_0 = \max_{h,m} c + \sum_{t=0}^{\infty} \beta^t \xi \log(h) - \sum_{t=1}^{T} \beta^t q m \tag{C.1}$$

subject to the budget constraint

$$c \leq W - (h - m)$$

and the credit constraint

$$m \leq \bar{m}(h) = \min\Big\{\theta^{LTV} h, \quad \frac{\theta^{PTI} y}{q + \kappa}\Big\},$$

where W is total wealth (including labor income), q is the principal and interest payment on the loan, κ reflects taxes and insurance, T is the length of the mortgage contract, and I have normalized the price of housing to unity. The value function above reflects the value of current consumption plus lifetime discounted housing services consumption, net of the discounted costs of repayment, using linear utility for nondurable consumption.[41]

For notation, define

$$\phi = \sum_{t=1}^{T} \beta^t q.$$

Substituting for c and m using the two constraints, and solving the geometric sum in (C.1) yields

$$V_0 = W + \underbrace{\frac{\xi \log(h)}{1 - \beta}}_{\text{PV benefit}} - \underbrace{(h - (1 - \phi)\bar{m}(h))}_{\text{PV cost}}.$$

The "PV benefit" term represents the marginal benefit of housing services, while the "PV

[41] Including κ in the PTI limit but not in (C.1) is a simplifying assumption. Since taxes and insurance are proportional to the house value rather than the loan size, they do not affect the marginal cost of additional debt, and can be omitted from (C.1) without loss. However, under either specification they affect the PTI limit, and by a similar amount, generating the core economics observed here.

cost" term reflects the present value cost of repaying the mortgage. When borrowers are impatient relative to the prevailing interest rate, we will have $\phi < 1$ and so the present value cost of buying \$1 of housing can fall below \$1 today. Differentiating these terms yields marginal benefits and costs

$$MB(h) = \frac{\xi}{(1-\beta)h}$$
$$MC(h) = 1 - (1-\phi)\bar{m}'(h).$$

The discontinuous jump in marginal cost occurs because $\bar{m}'(h) = \theta^{LTV}$ in the LTV-constrained region, then drops discretely to zero once the borrower becomes PTI-constrained. The borrower will choose the corner solution

$$h^{both} = \frac{\theta^{PTI} y}{\theta^{LTV}(q+\kappa)}$$

whenever

$$1 - (1-\phi)\theta^{LTV} < \frac{\xi}{(1-\beta)h^{both}} < 1 \tag{C.2}$$

where the left and right terms correspond to $MC(h)$ in the LTV-constrained and PTI-constrained areas of the state space, respectively, and the center term is $MB(h^{both})$.

We can calibrate this example at annual frequency by setting $\beta = 0.9$, $y = 63.02$, $\theta^{LTV} = 0.8$, $\theta^{PTI} = 0.28$, $\xi = 21.1$, and $\kappa = 0.0175$, computing q as the annual principal and interest on a loan with interest rate 6%, and scaling c, h, and m to be measured in thousands of dollars. It is easily checked that condition (C.2) holds for all the experiments of Section 2, verifying that the borrower indeed follows the corner solution as pictured.

D Life Cycle Extension

In this section I develop an extension to the model featuring a life cycle.

D.1 Model Structure.

In this extension, borrowers have an age group $j \in \{1, \ldots, J\}$, and transition to the next age group with probability ϕ. Households receiving the aging shock at age J die and are replaced by newborn households of age 0. As is standard in overlapping generations models, I assume that borrower households are able to trade a complete set of contingent

securities that allow them to fully share risk among borrower households of the same age, but not across ages. For the oldest households, this means that the housing assets and mortgage liabilities of dying households are reallocated proportionally to surviving households in the same group, as in Blanchard (1985).

I assume that borrowers of age j have income

$$y_{i,t} = e_i^y \tilde{y}_B^j \bar{Y}_B$$

for some relative income parameters $\tilde{y}_B^j$, capturing the life-cycle path of income. The budget constraint becomes

$$\begin{aligned} c_{B,t}^j \leq \underbrace{(1-\tau)\tilde{y}_B^j \bar{Y}_B}_{\text{labor income}} - \underbrace{\pi^{-1}G^{-1}\big((1-\tau)X_{t-1}^j + \nu M_{t-1}^j\big)}_{\text{mortgage payment}} - \underbrace{\delta p_t H_{t-1}^j}_{\text{maintenance}} \\ - \rho \underbrace{p_t(H_t^{j,*} - H_{t-1}^j)}_{\text{new housing}} + \rho \underbrace{\big(M_t^{j,*} - (1-\nu)\pi^{-1}G^{-1}M_{t-1}^j\big)}_{\text{new debt}} \end{aligned} \tag{D.1}$$

where the laws of motion for the state variables become

$$\begin{aligned} M_t^j = (1-\phi) \underbrace{\left\{\rho M_t^{j,*} + (1-\rho)(1-\nu)(\pi G)^{-1} M_{t-1}^j\right\}}_{\text{non-aging households}} \\ + \phi \underbrace{\left\{\rho M_t^{j-1,*} + (1-\rho)(1-\nu)(\pi G)^{-1} M_{t-1}^{j-1}\right\}}_{\text{aging households}} \\ X_t^j = (1-\phi) \underbrace{\left\{\rho r_t^* M_t^{j,*} + (1-\rho)(1-\nu)(\pi G)^{-1} X_{t-1}^j\right\}}_{\text{non-aging households}} \\ + \phi \underbrace{\left\{\rho r_t^* M_t^{j-1,*} + (1-\rho)(1-\nu)(\pi G)^{-1} X_{t-1}^{j-1}\right\}}_{\text{aging households}} \end{aligned}$$

for $1 < j < J$. These laws of motion reflect convex combinations of share $1-\phi$ of households with starting age j that do not receive the aging shock and remain at age j, and share ϕ of households with starting age $j-1$ that receive the aging shock and transition to age j.

For $j = J$, the oldest households, the assumption that assets and liabilities of dying

households are shared among the surviving households implies

$$M_t^J = \left\{ \rho M_t^{J,*} + (1-\rho)(1-\nu)(\pi G)^{-1} M_{t-1}^J \right\}$$
$$X_t^J = \left\{ \rho r_t^* M_t^{J,*} + (1-\rho)(1-\nu)(\pi G)^{-1} X_{t-1}^J \right\}$$

since all assets and liabilities are now retained within the group regardless of the aging probability.

For $j = 1$, there are no younger households to age up into this group. Instead, a share ϕ of households with $j = 1$ are newborn households entering the market. Since these newborn households require housing, I assume that these households enter the market "active," which matches the relatively high LTV among young households. Thus, the total share of active households with $j = 1$ is $(1-\phi)\rho + \phi$, corresponding to the share ϕ of newborn households that are active by construction, and the share ρ of non-newborn households that are randomly chosen to be active. This translates to the $j = 1$ laws of motion

$$M_t^1 = \Big[\phi + (1-\phi)\rho \Big] M_t^{1,*} + (1-\phi)(1-\rho)(1-\nu)(\pi G)^{-1} M_{t-1}^1$$
$$X_t^1 = \Big[\phi + (1-\phi)\rho \Big] r_t^* M_t^{1,*} + (1-\phi)(1-\rho)(1-\nu)(\pi G)^{-1} X_{t-1}^1.$$

This assumption ensures that first-time purchases must be made using only the resources of age $j = 1$ agents, which are limited by the relatively low incomes of this group.

All borrowers with $j < J$ maximize the value function

$$V_t^j = \log(C_{B,t}^j) + \xi(1+\omega)\log(H_{B,t-1}^j) + \beta \mathbb{E}_t \left\{ (1-\phi) V_{t+1}^j + \phi V_{t+1}^{j+1} \right\}$$

subject to (D.1), and where $V^{J+1} = 0$ upon the agent's death.

The optimal housing allocations become age-dependent but are generally similar to the solution in the baseline model:

$$h_{i,t}^* = \begin{cases} e_i^h h_{j,t}^{LTV} & \text{for } e_i < \bar{e}_{j,t}^{LTV} \\ e_i^y h_{j,t}^{both} & \text{for } \bar{e}_{j,t}^{LTV} \leq e_i \leq \bar{e}_{j,t}^{PTI} \\ e_i^h h_{j,t}^{PTI} & \text{for } e_i \geq \bar{e}_{j,t}^{PTI}, \end{cases} \tag{D.2}$$

for

$$h_{j,t}^{LTV} = \frac{\Psi_{1,t}^{j}}{(1-\mu_t\theta^{LTV})p_t - \Psi_{2,t}^{j}} \tag{D.3}$$

$$h_{j,t}^{both} = \frac{\theta^{PTI}\alpha\tilde{y}_B^j\bar{Y}_B}{\hat{r}_t^*\theta^{LTV}p_t} \tag{D.4}$$

$$h_{j,t}^{PTI} = \frac{\Psi_{1,t}^{j}}{p_t - \Psi_{2,t}^{j}} \tag{D.5}$$

and where the cutoffs $(\bar{e}_{j,t}^{LTV}, \bar{e}_{j,t}^{PTI})$ become

$$\bar{e}_{j,t}^{LTV} = K\left(\frac{h_{j,t}^{both}}{h_{j,t}^{LTV}}\right) \tag{D.6}$$

$$\bar{e}_{j,t}^{PTI} = K\left(\frac{h_{j,t}^{both}}{h_{j,t}^{PTI}}\right). \tag{D.7}$$

The functions $\Psi_{1,t}$ and $\Psi_{2,t}$ now become age-specific, given by

$$\begin{aligned}\Psi_{1,t}^{j} = \left(\frac{\beta}{u'(C_{B,t}^{j})}\right)\mathbb{E}_t\Big\{&(1-\phi)u'(C_{B,t+1}^{j})\left[(1+\omega)\xi C_{B,t+1}^{j} + (1-\rho)\Psi_{1,t+1}^{j}\right]\\ &+\phi u'(C_{B,t+1}^{j'})\left[(1+\omega)\xi C_{B,t+1}^{j'} + \Psi_{1,t+1}^{j'}\right]\Big\}\\ \Psi_{2,t}^{j} = \left(\frac{\beta}{u'(C_{B,t}^{j})}\right)\mathbb{E}_t\Big\{&(1-\phi)u'(C_{B,t+1}^{j})\left[(\rho-\delta)p_{t+1} + (1-\rho)\Psi_{2,t+1}^{j}\right]\\ &+\phi u'(C_{B,t+1}^{j'})\left[(\rho-\delta)p_{t+1} + (1-\rho)\Psi_{2,t+1}^{j'}\right]\Big\}\end{aligned}$$

where $j' = j+1$ for $j < J$, and $j' = J$ for $j = J$. Compared to (13) and (14), these expressions require two types of changes. First, since the Ψ functions are age-dependent, the continuation value differs depending on whether the borrower ages. Second, in place of a single borrower SDF, the marginal utility in the next period also depends on whether the borrower ages, and thus the two marginal utility terms must be handled separately.

The house price is then determined by market clearing according to

$$\bar{H}_B = \frac{1}{J}\sum_{j=1}^{J}\left\{h_{j,t}^{LTV}\int^{\bar{e}_{j,t}^{LTV}} e_i^h\, d\Gamma(e) + h_{j,t}^{both}\int_{\bar{e}_{j,t}^{LTV}}^{\bar{e}_{j,t}^{PTI}} e_i^y\, d\Gamma(e) + h_{j,t}^{PTI}\int_{\bar{e}_{j,t}^{PTI}} e_i^h\, d\Gamma(e)\right\}$$

where the initial scaling by $1/J$ aggregates over age groups of equal size.

On the mortgage side, borrower optimality implies

$$\Omega_{B,t}^{m,j} = \left(\frac{\beta}{u'(C_{B,t}^j)}\right)\mathbb{E}_t\Big\{(1-\phi)u'(C_{B,t+1}^j)\pi^{-1}G^{-1}\left[\nu+(1-\nu)\left(\rho+(1-\rho)\Omega_{B,t+1}^{m,j}\right)\right] + \phi u'(C_{B,t+1}^{j'})\pi^{-1}G^{-1}\left[\nu+(1-\nu)\left(\rho+(1-\rho)\Omega_{B,t+1}^{m,j'}\right)\right]\Big\}$$

$$\Omega_{B,t}^{x,j} = \left(\frac{\beta}{u'(C_{B,t}^j)}\right)\mathbb{E}_t\Big\{(1-\phi)u'(C_{B,t+1}^j)\pi^{-1}G^{-1}\left[(1-\tau)+(1-\nu)(1-\rho)\Omega_{B,t+1}^{x,j}\right] + \phi u'(C_{B,t+1}^{j'})\pi^{-1}G^{-1}\left[(1-\tau)+(1-\nu)(1-\rho)\Omega_{B,t+1}^{x,j'}\right]\Big\}$$

where $j' = j+1$ for $j < J$, and $j' = J$ for $j = J$. Age-specific debt limit multipliers are then defined by

$$\mu_{j,t} = 1 - \Omega_{B,t}^{m,j} - r_t^*\Omega_{B,t}^{x,j}.$$

The remaining model structure remains identical to the baseline in Section 4.

D.2 Approximation Formulas.

In this section I rederive the approximations (28) and (29) for the life-cycle case. These derivations are important for intuition, as they make it straightforward to understand why the life-cycle model and original model have nearly identical dynamics.

To begin, I rewrite the aggregate housing demand condition as

$$H_B = \sum_{j=1}^{J}\Bigg\{\underbrace{\int^{\bar{e}_j^{LTV}} e^h(e_i)\, d\Gamma(e)}_{w_H^{LTV,j}} h_j^{LTV} + \underbrace{\int_{\bar{e}_j^{LTV}}^{\bar{e}_j^{PTI}} e^y(e_i)\, d\Gamma(e)}_{w_H^{both,j}} h_j^{both} + \underbrace{\int_{\bar{e}_j^{PTI}} e^h(e_i)\, d\Gamma(e)}_{w_H^{PTI,j}} h_j^{PTI}\Bigg\}$$

Following the approach in Appendix A.6, we can log-linearize this expression with re-

spect to an arbitrary variable z, divide through by H_B, and rearrange to obtain

$$d\log p = \bar{s}_H^{both}\left[d\log\theta^{PTI} - d\log\theta^{LTV} - d\log\hat{r}^*\right] + \left\{\frac{1}{J}\sum_{j=1}^{J} s_H^{LTV,j}\left(\frac{\mu_j\theta^{LTV}}{1-\mu_j\theta^{LTV}-\psi_j}\right)\right\} d\log\theta^{LTV} + (1-\bar{s}_H^{both})d\omega. \tag{D.8}$$

where e.g., $s_H^{LTV,j} = w_H^{LTV,j}h_j^{LTV}/H_B$, and where

$$\bar{s}_H^{both} = \frac{1}{J}\sum_{j=1}^{J} s_H^{both,j}$$

is the aggregate share of jointly-constrained households. Equation (D.8) is the life-cycle version of (28), and shows that up to first order, the responses to shocks affecting PTI limits and to housing preferences are identical to the baseline model holding the aggregate share of jointly-constrained households fixed.

Note that this approximation can be applied beyond this specific life-cycle model. The derivation above assumes only that we are subdividing the population into j separate groups, but does not impose any structure on the membership of those groups aside from assuming that they are of equal size (which could be easily relaxed using weighted averages) and that borrowers in each group face the same credit limits and interest rates. As a result, a wide variety of models or extensions should deliver the same responses to shocks affecting PTI limits, up to first order, provided they have the same jointly-constrained share. Since this share is in turn calibrated directly from the distribution observed in the data, this provides confidence that the results are highly robust and not dependent on specific model assumptions.

For intuition behind this result, recall that jointly-constrained households are at a corner solution in which they choose housing to equalize their LTV and PTI limits. As long as they remain jointly constrained, their housing demand is therefore a highly mechanical function of their credit limits

$$h_{i,t}^* = \frac{\theta^{PTI}\hat{y}_{i,t}}{\hat{r}_t^*\theta^{LTV}p_t}$$

where $\hat{y}_{i,t} = \alpha\bar{Y}_B e_i^y$ is the borrower's effective income. Taking logs of this expression yields

$$\log h_{i,t}^* = \log\theta^{PTI} + \log\hat{y}_{i,t} - \log\hat{r}_t^* - \log\theta^{LTV} - \log p_t.$$

This relationship predicts an elasticity of 1 for the housing demand of a jointly-constrained household with respect to PTI standards θ^{PTI} and income $\hat{y}_{i,t}$, and an elasticity of -1 with respect to the effective payment rate $\hat{r}_t$, LTV standards θ^{LTV}, and house prices p_t. These individual demand elasticities follow directly from the corner solution equalizing the LTV and PTI constraints, and do not depend on other blocks of the model.

To map to aggregate housing demand, we can simply multiply these elasticities by the share of jointly-constrained households $\bar{s}_H^{both}$. These elasticities aggregate housing demand elasticities in turn translate directly to those of house prices due to my model's fixed housing supply and Cobb-Douglas consumption aggregator, which corresponds to an elasticity of substitution between housing and nondurables of one. While in principle this elasticity of substitution could take a different value in a more general CES aggregator, the extreme stability of observed housing expenditure shares over time implies an elasticity very close to one (Piazzesi, Schneider, and Tuzel, 2007).

The response to a change in the LTV limit is less tightly pinned down by the aggregate shares at each constraint. In (D.8), this response also depends on group-specific objects such as the steady-state shadow value of debt μ_j and the marginal value of future housing transactions relative to price ψ_j. Thus, unlike the response to shocks affecting PTI limits, the LTV response in the life-cycle extension need not exactly match the Benchmark model even when the aggregate LTV-constrained share is the same. In practice, however, I show below that the difference is quantitatively small.

Similarly, we can update the derivation of (29) to obtain

$$\begin{aligned} d\log M^* = &\left[\frac{1}{J}\sum_{j=1}^{J} s_M^{LTV,j}\left(\frac{1-\psi_j}{1-\mu_j\theta^{LTV}-\psi_j}\right)\right] d\log\theta^{LTV} + \bar{s}_M^{LTV} d\omega \\ &+ (\bar{s}_M^{both} + \bar{s}_M^{PTI})\left(d\log\theta^{PTI} - d\log\hat{r}^*\right). \end{aligned} \tag{D.9}$$

where e.g.,

$$\bar{s}_M^{both} = \frac{1}{J}\sum_{j=1}^{J} s_M^{both,j}.$$

This debt approximation delivers a similar conclusion to that for house prices above. Responses to changes in the PTI limit, whether exogenous or induced by changes in $\hat{r}^*$, are approximately invariant across the Benchmark and life-cycle models, as long as the aggregate mortgage shares of jointly-constrained and PTI-only borrowers are the same. The

response to the housing preference shifter ω is similarly pinned down by the aggregate mortgage share of LTV-only borrowers. In contrast, the response to changes in LTV limits may again differ across models because it depends on the group-specific values μ_j and ψ_j, but I show below that this difference is quantitatively small between the baseline model and life-cycle extension.

D.3 Calibration

For the age structure, I choose five age groups: (1) under 30, (2) 30–39, (3) 40–49, (4) 50–59, and (5) 60 and over. These generally correspond to 10-year age groups, extended at the lower and upper bounds. I correspondingly set $\phi = 1/40$, mapping to one age transition on average every 10 years (40Q), and an expected adult lifespan of 50 years.[42] For the age-specific endowments, I set the values $\tilde{y}^j_{B,t}$ so that the relative income compared to the average borrower in the model exactly matches the relative average income for mortgagors within the corresponding age group compared to average income among mortgagors in the 1998 SCF. The remaining non-life-cycle parameters are calibrated as in the benchmark model in Section 4.3.

Table D.1: Parameter Values, Life Cycle Calibration

Parameter	Name	Value	Internal	Target/Source
		Life Cycle		
Aging probability	ϕ	0.025	N	10-year average transitions
Relative income, age 1	$\tilde{y}^1_B$	0.6231	N	SCF mortgagor income by age
Relative income, age 2	$\tilde{y}^2_B$	0.9981	N	SCF mortgagor income by age
Relative income, age 3	$\tilde{y}^3_B$	1.1955	N	SCF mortgagor income by age
Relative income, age 4	$\tilde{y}^4_B$	1.2619	N	SCF mortgagor income by age
Relative income, age 5	$\tilde{y}^5_B$	0.9213	N	SCF mortgagor income by age

D.4 Life Cycle Profiles

To validate the model's life-cycle implications, Figure D.1 compares several statistics by age group in the model's steady state to those for mortgagors in the 1998 SCF.

[42] In reality, households at age 60 can expect to live longer than 10 years on average. However, this specification avoids addressing the persistent accumulation of assets late in life that is challenging for standard models to explain. I leave a complete treatment addressing this pattern to future work.

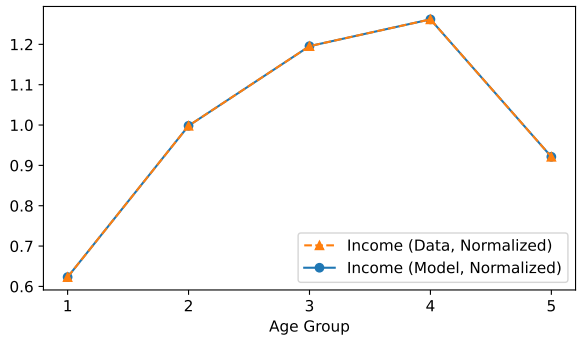

(a) Income

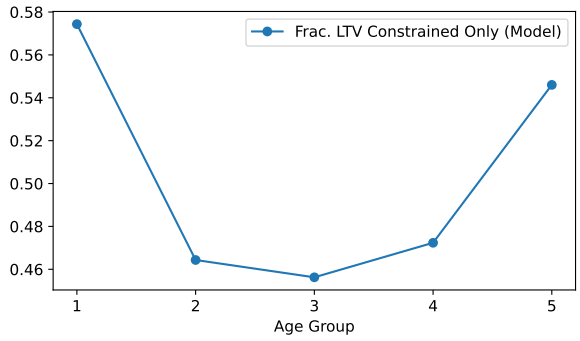

(b) Frac. Constrained by LTV Only

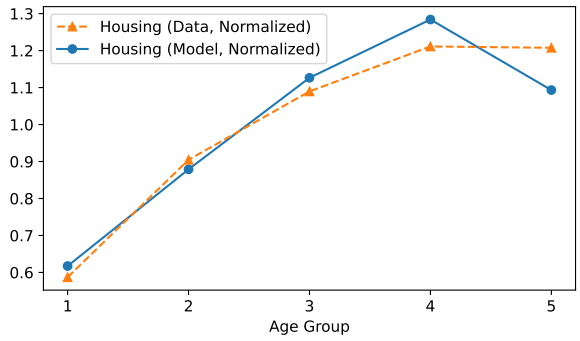

(c) House Value

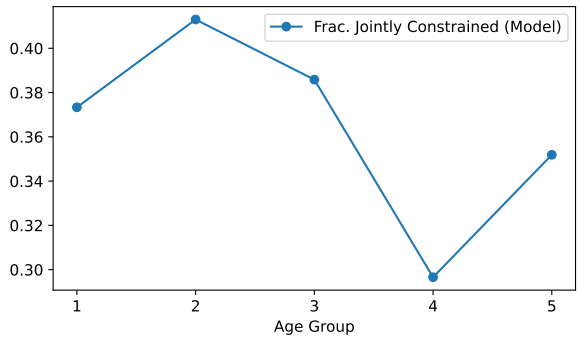

(d) Frac. Jointly Constrained

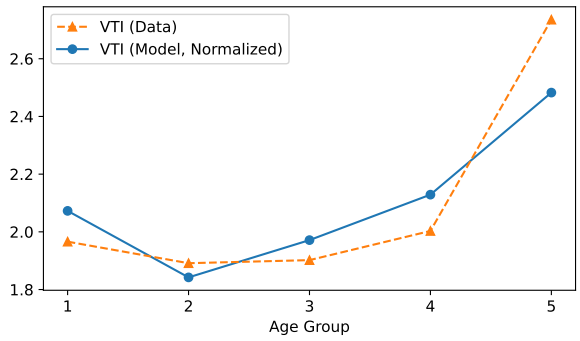

(e) Value-to-Income

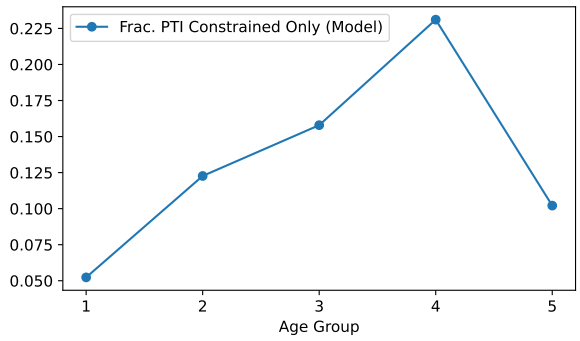

(f) Frac. Constrained by PTI Only

Figure D.1: Life Cycle Profiles, Model vs. Data

Note: Each panel displays averages for each age group in the model. Panels (a), (c), and (e) also display averages from the data, obtained from the subsample of mortgagors (households with a house and a positive mortgage balance) in the 1998 Survey of Consumer Finances. In panels (a) and (c), each series is divided by a constant to achieve a mean of unity across age groups. In Panel (e), the model series is normalized by multiplying by a constant to equalize its mean across age groups with the corresponding average in the data series.

The left column compares the implied model moments to the data. To begin, Panel (a) displays income by age group, with both series normalized by dividing by a constant so that they have mean unity. The plot shows a typical hump-shaped income profile, which

peaks in middle age, and is substantially lower for the youngest and oldest agents. By construction, the model is able to perfectly match these moments.

Next, Panel (c) displays the value of housing by age group, with both series again normalized to have mean unity. Although the extension adds no housing-related parameters, and these moments are untargeted, the model provides a good fit of the overall pattern, with housing increasing over the life cycle before falling for the very oldest group.

Taking the ratio of house value to income yields the value-to-income ratio, plotted in Panel (e). This ratio is the key to determining which constraint is binding. The model successfully matches the increasing VTI ratio profile observed in the data. While both young and old households have relatively low incomes, the effect is not symmetric, with much higher VTI ratios for the old than for the young. The model reproduces this pattern since the youngest households do not begin with housing but instead must accumulate it during a period when their marginal utility of consumption is high, leading them to choose homes of relatively modest size. While older households also receive low incomes, they still possess the housing that they accumulated over their lifetime, and downsize only slowly, leading to high average VTI ratios.

The right column of Figure D.1 displays the shares of households by age and constraint type in the model. Panel (b) shows that the youngest households are the most likely to be LTV-constrained. This is consistent with their low VTI ratios, due to the high opportunity cost of making a larger down payment, and is consistent with the high LTV-only shares for first-time buyers in Table 1. The share constrained by PTI increases sharply in middle age (groups 2–4). While these households have higher incomes, they have accumulated home equity over time, and find the required down payments implied by LTV limits easier to meet than their PTI limits as housing purchases peak. As households enter the oldest age group the LTV-only share increases, as they downsize in line with their smaller incomes.

D.5 Responses to Shocks

I next reproduce responses to the changes in credit limits previously studied in Figure 6 for this life cycle extension. To begin, Appendix Figure D.2 displays responses to these shocks by age group. The figure shows that the responses of housing demand, effective collateral demand (represented by $\hat{F}^{LTV}$) and the loan-to-income ratio to these shocks all exhibit economically significant variation by age.

For instance, following a shock to PTI limits in Appendix Figure D.2(a), households

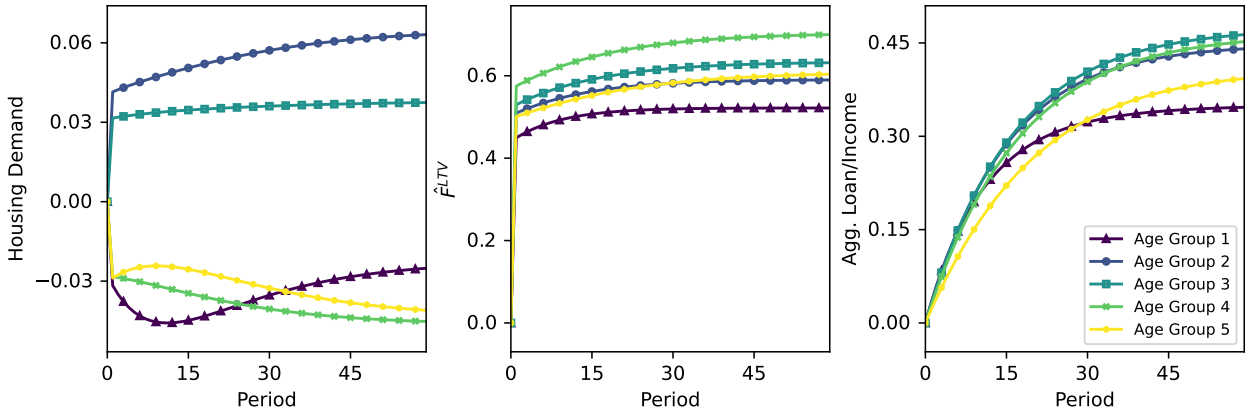

(a) Shock to θ^{PTI}

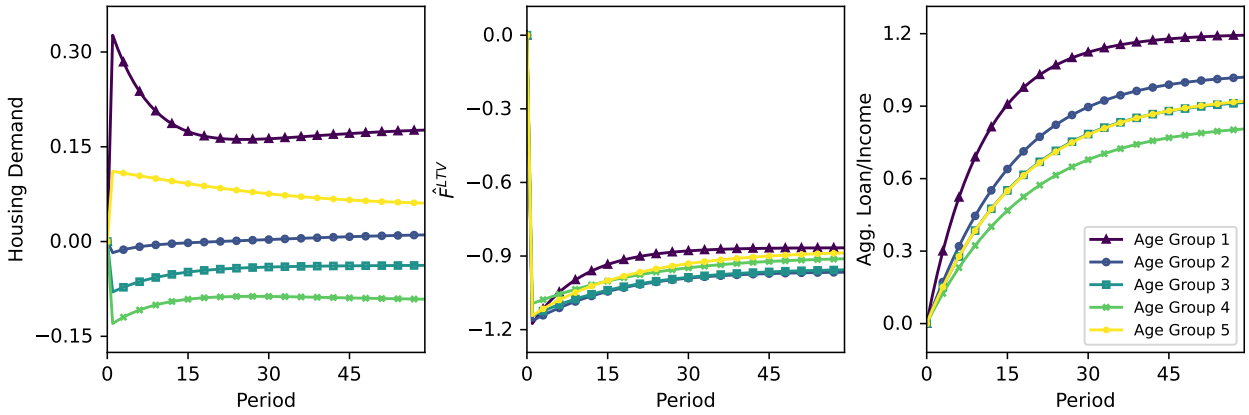

(b) Shock to θ^{LTV}

Figure D.2: Responses by Age, Shocks to Credit Standards

Note: The figure displays model responses to a permanent increase of 0.01 to $\log \theta^{PTI}$ (Panel (a)) or $\log \theta^{LTV}$ (Panel (b)). Each line indicates the response of one age group, as defined above. The responses are shown for the Benchmark Economy and the Life Cycle Economy. A value of 1 represents a 1% increase relative to steady state, except for $\hat{F}^{LTV}$, which is measured in percentage point deviations from steady state. The price-rent ratio is defined as $p_t/(u^h_{B,t}/u^c_{B,t})$, where the denominator is the implied price of rental services. Aggregate LTI is defined as $M_t/\bar{Y}$.

in Age Groups 2 and 3, corresponding to households in their 30s and 40s, see a large relative increase in housing demand, due to their higher fractions of jointly-constrained households. Conversely, Appendix Figure D.2(b) shows that these groups see a relative decrease in demand following a loosening of LTV limits, while the youngest households, similar to first-time home buyers, exhibit a large relative increase in demand.

Next, Appendix Figure D.3 aggregates over age groups to plot total responses in the Life Cycle economy, and compares them to the Benchmark economy studied in Section 4. These responses show a clear and uniform pattern: the macroeconomic dynamics for the

life-cycle model are virtually identical to those in the Benchmark case without a life cycle.

This result can be understood in light of the intuition in Appendix A.6. Because the aggregate share of jointly-constrained households in the Life Cycle economy is 36.4%, which is very close to the corresponding share of 37.4% in the Benchmark economy, the approximate response to changes in PTI limits is nearly identical, which is also borne out in this exact nonlinear response. While, as discussed above, the responses to the change in LTV limits could differ between models due to comovement between $s^{LTV,j}$ and μ_j or ψ_j, in practice the discrepancy is minimal in this setting.

Last, I repeat this exercise for the endogenous shocks to LTV and PTI limits (via changes in inflation and housing preference ω). The results are displayed in Appendix Figures D.4 and D.5. As with the shocks to credit limits, these changes cause large variation in the response by age group, but exhibit minimal differences relative to the Benchmark model once aggregated. As a result, I conclude that the main model results appear highly robust to the inclusion of a life cycle, and following the argument in Appendix A.6, to many other potential extensions as well.

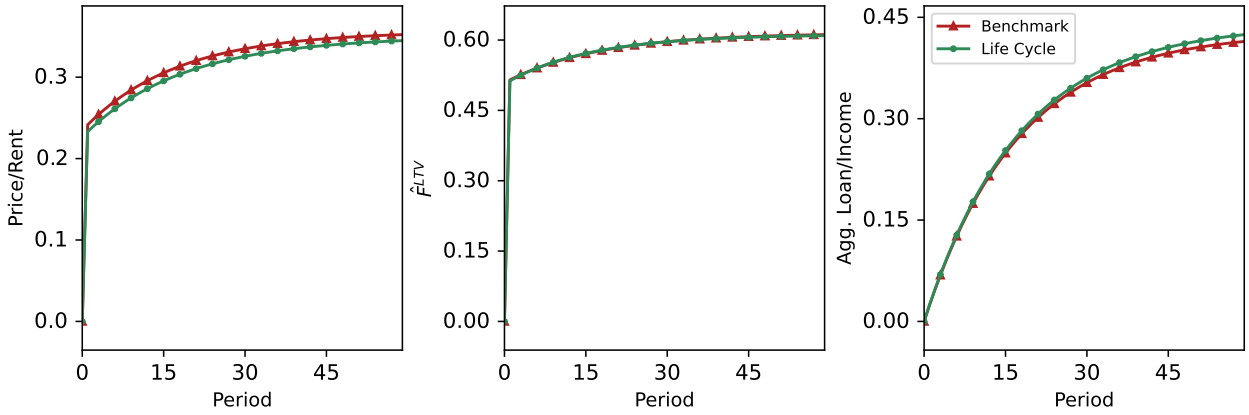

(a) Shock to θ^{PTI}

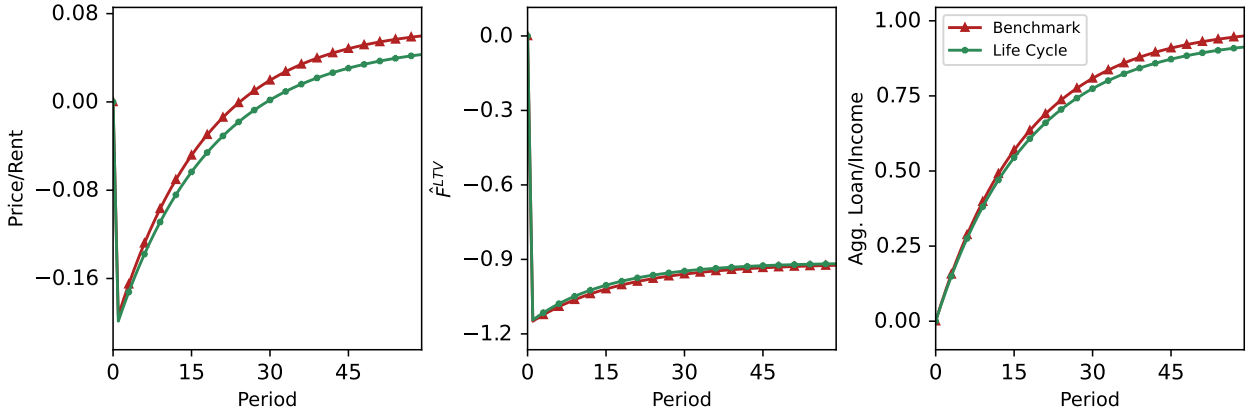

(b) Shock to θ^{LTV}

Figure D.3: Responses, Shocks to Credit Standards, Benchmark vs. Life Cycle Economies

Note: The figure displays model responses to a permanent increase of 0.01 to $\log \theta^{PTI}$ (Panel (a)) or $\log \theta^{LTV}$ (Panel (b)) in the Life Cycle Economy. A value of 1 represents a 1% increase relative to steady state, except for $\hat{F}^{LTV}$, which is measured in percentage point deviations from steady state. The price-rent ratio is defined as $p_t/(u^h_{B,t}/u^c_{B,t})$, where the denominator is the implied price of rental services. Aggregate LTI is defined as $M_t/\bar{Y}$.

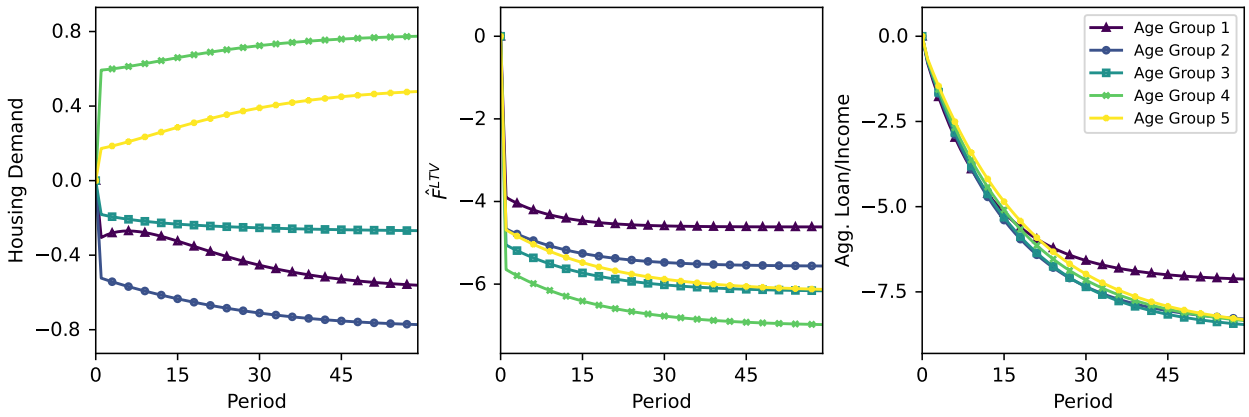

(a) High Inflation Experiment

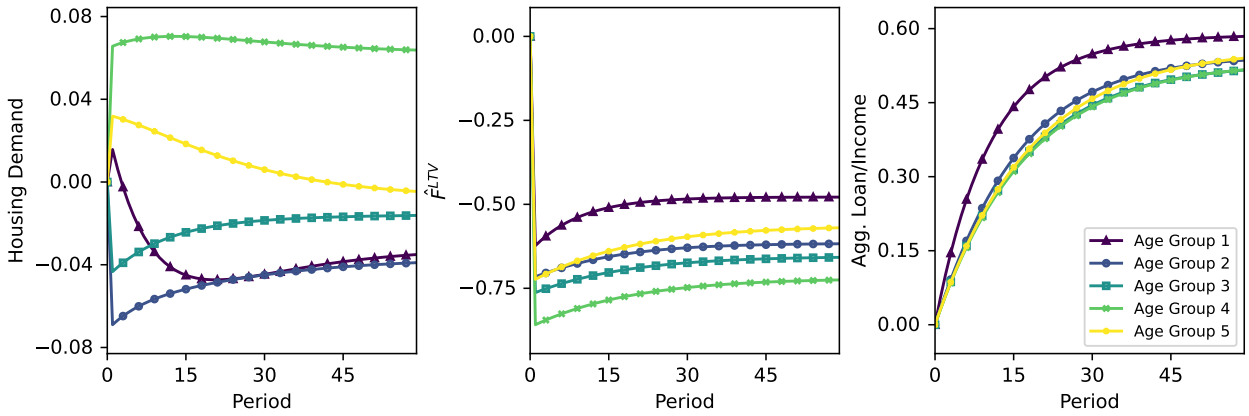

(b) High ω Experiment

Figure D.4: Responses by Age, Endogenous Changes to Credit Limits

Note: The figure displays model responses to a permanent increase of 0.0025 (1% annual) to $\log \pi$ (Panel (a)) or a permanent increase of 0.01 in ω (Panel (b)) in the Life Cycle economy. Each line indicates the response of one age group, as defined above. A value of 1 represents a 1% increase relative to steady state, except for $\hat{F}^{LTV}$, which is measured in percentage point deviations from steady state. The price-rent ratio is $p_t/(u^h_{B,t}/u^c_{B,t})$, where the denominator is the implied price of rental services. The aggregate loan-income ratio is $M_t/\bar{Y}$.

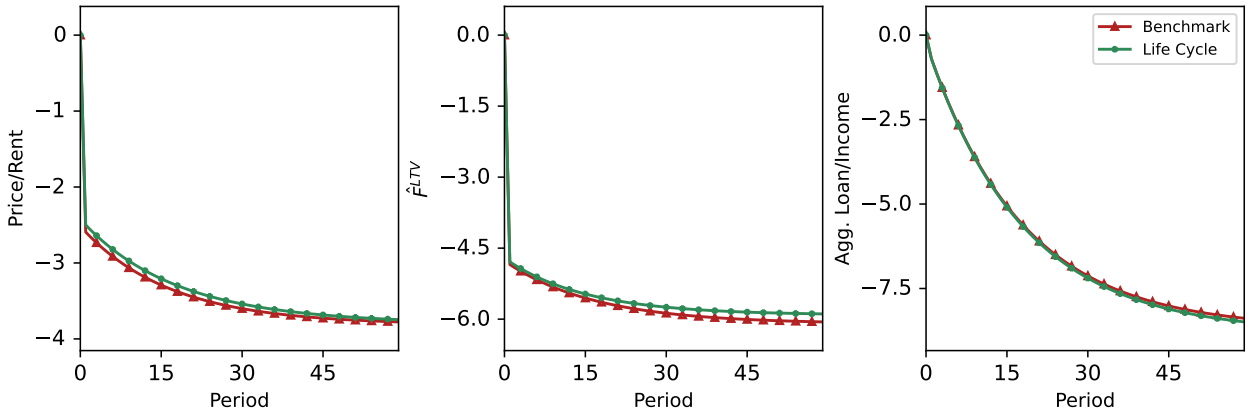

(a) High Inflation Experiment

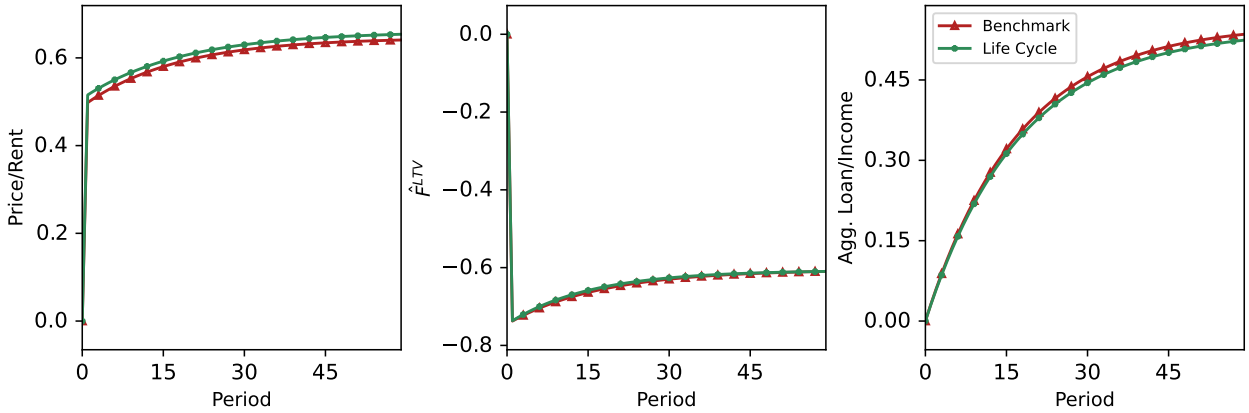

(b) High ω Experiment

Figure D.5: Responses, Endogenous Changes to Credit Limits, Benchmark vs. Life Cycle Economies

Note: The figure displays model responses to a permanent increase of 0.0025 (1% annual) to $\log \pi$ (Panel (a)) or a permanent increase of 0.01 in ω (Panel (b)). The responses are shown for the Benchmark Economy and the Life Cycle Economy. A value of 1 represents a 1% increase relative to steady state, except for $\hat{F}^{LTV}$, which is measured in percentage point deviations from steady state. The price-rent ratio is $p_t/(u^h_{B,t}/u^c_{B,t})$, where the denominator is the implied price of rental services. The aggregate loan-income ratio is $M_t/\bar{Y}$.