Managing a Housing Boom

Jason Allen, Daniel L. Greenwald*

January 2022

Abstract

We investigate how macroprudential policies intended to dampen rises in debt and house prices are influenced by segmentation in the housing and mortgage market. We develop a modeling framework with two mortgage submarkets: a government-insured sector with loose LTV limits and tight PTI limits, and an uninsured sector displaying the reverse pattern. This form of heterogeneity is modeled after the Canadian mortgage system, but is common in countries around the world. This multi-market structure implies that house prices are much more responsive to increases in latent demand, allowing for larger booms. While tightening payment-to-income (PTI) limits is highly effective at dampening a housing boom in a one-sector system, tightening these limits in the insured sector *only* is much weaker, due to substitutions into the uninsured sector. In contrast, the effect of tightening loan-to-value (LTV) limits in the uninsured sector is *strengthened* by market segmentation, causing price-rent ratios to fall, while the same tightening in the insured sector would counterproductively cause price-rent ratios to rise.

*The views in this paper are those of the authors and do not represent those of the Bank of Canada. Allen: Bank of Canada, email: jallen@bank-banque-canada.ca. Greenwald: NYU Stern School of Business; email: dlg340@stern.nyu.edu.

1 Introduction

The prevalence in recent years of large booms in house prices, and the disruptive consequences of large crashes, have drawn increased attention to the role of macroprudential policy in stabilizing housing and mortgage markets. But despite this uptick in salience, much remains to be learned theoretically about the effective conduct of these policies. For example, Jácome and Mitra (2015) note that while limits on loan-to-value (LTV) and payment-to-income (PTI) ratios are widespread around the world, there are few theoretical frameworks available to guide regulators in their application.

In this paper, we develop a realistic yet tractable general equilibrium framework, extending Greenwald (2018), that can be used to evaluate the effects of the essential macroprudential tools available to policymakers. Our key innovation is to capture an important but largely overlooked feature of many mortgage markets around the world: a segmented market, with different submarkets offering different credit standards. We design this framework to reflect the Canadian mortgage system, which is divided between a government-insured sector that offers high LTV ratios (maximum of 95%) but imposes tight PTI requirements (maximum of 44%), and an uninsured sector that uses much tighter LTV limits (maximum of 80%) but leaves PTI ratios essentially unlimited. While we focus on Canada, it is worth noting that such segmentation is common, and can be seen for example in the US, where the FHA, GSEs, and private label securitizers all employed different underwriting policies during the US housing boom.

Using this framework, we find that market segmentation has a first-order influence on the dynamics of the housing market and the consequences of macroprudential policy, whose effects depend crucially on which submarket-constraint combination is targeted. First, we demonstrate that, holding credit limits fixed, the multi-market structure allows for larger increases in house prices for a given increase in latent housing demand, allowing for larger housing booms. In a single-market setting, as house prices rise, more and more borrowers become constrained by binding PTI limits, slowing further increases in housing demand and credit growth. However, in this dual-market setting, borrowers facing binding PTI limits in the insured sector can escape them by switching to the uninsured sector. This dual market system not only allows for larger credit growth as a result, but also leads these switching borrowers to greatly increase their housing demand due to marginal incentives to acquire additional collateral in the uninsured space. As a result, a fixed PTI limit is a weaker defense against housing booms in the multi-market setting.

Next, we consider the implications of macroprudential policies that directly tighten the various credit limits. While a single sector model would show a tightening of PTI limits to be highly effective at dampening a rise in debt and house prices, we find that tightening these limits in the government-insured sector *only* — a relatively easy policy for regulators to implement — is substantially weakened. Essentially, borrowers react to this tightening by switching into the uninsured sector where PTI limits are loose but housing collateral is scarce, putting countervailing upward pressure on housing demand and house prices. Tightening PTI limits in the uninsured sector, in contrast, does not trigger such a countervailing force, since borrowers with this submarket-constraint pairing are far from the insured-uninsured boundary and therefore do not exhibit sector switching behavior.

A reverse intuition holds for shifts in LTV limits. Tightening LTV limits in the insured sector has a similar effect as in a single-sector model. As found in Greenwald (2018), this policy counterproductively causes price-rent ratios to *rise*, due to a "constraint switching effect" where an increase in the fraction of LTV-constrained borrowers — who can use additional housing to relax their borrowing limits — increases the demand for housing collateral. Perhaps surprisingly, this result does *not* hold when LTV limits are tightened in the uninsured sector. Importantly, this policy switches marginal borrowers from the uninsured sector, where they are LTV-constrained, to the insured sector, where they become PTI-constrained. This switch reduces those borrowers' collateral needs, dampening housing demand, and ensuring that price-rent ratios fall instead of rise.

Related Literature. Our paper connects to several literatures. The first relates to measuring or estimating the strength of different macroprudential policies. Papers measuring these quantities using structural models include Garriga and Hedlund (2018), Grodecka (2019), Rubio and Carrasco-Gallego (2014), while papers estimating these effects empirically include Benetton, Bracke, Cocco, and Garbarino (2019), Cerutti, Claessens, and Laeven (2017a), Everett, de Haan, Jansen, McQuade, and Samarina (2021), Igan and Kang (2011), Kelly, McCann, and O'Toole (2018), Kinghan, McCarthy, and O'Toole (2019), Kuttner and Shim (2016), Morgan, Regis, and Salike (2019), Richter, Schularick, and Shim (2019), Wong, Fong, Li, and Choi (2011). Our paper provides a new theoretical foundation extending these insights to a multi-market setting, showing that the effectiveness of policies depends crucially on the submarket in which they are applied.

Our paper is also closely related to work on the coexistence of multiple credit constraints, and the combination of LTV and PTI, in particular, including Grodecka (2019), Ingholt (2019), Johnson (2018), and Justiniano, Primiceri, and Tambalotti (2019). This paper's framework is most closely related to Greenwald (2018), who performs a similar analysis for a single-market setting, and finds that households switching between being constrained by the two limits can have large effects on house prices (the constraint switching effect). As a result, policies or events that shift the relative tightness of the two constraints is key to housing dynamics. This paper complements this existing work with a novel mechanism, due to households' ability to switch submarkets. Because these submarket switches most commonly entail a change in a household's binding constraint, this mechanism generates novel implications for house price responses with major divergences from the single-market setting.

Last, our paper links to the literature on the causes and consequences of real estate booms, including Cerutti, Dagher, and Dell'Ariccia (2017b), Crowe, Dell'Ariccia, Igan, and Rabanal (2013), Favilukis, Ludvigson, and Van Nieuwerburgh (2017), Glaeser and Gyourko (2006), Kaplan, Mitman, and Violante (2019), and Keys, Piskorski, Seru, and Vig (2013). We complement these works by demonstrating that a multi-market structure can amplify the ability of other latent factors influencing demand to drive a housing boom.

Overview. The rest of the draft is organized as follows. Section 2 provides background information on the Canadian mortgage market, as well as evidence supporting our modeling assumptions. Section 3 provides a numerical example that explains the intuition behind our results. Section 4 develops the theoretical model. Section 5 displays our findings, while Section 6 concludes.

2 Background: The Canadian Mortgage Market

The Canadian mortgage market is divided into two sectors. First, there is the "insured" sector, in which default risk on mortgages is guaranteed by the government. This sector borrowers to obtain loans with up to 95% LTV ratios, but restricts them to a maximum PTI ratio of 44% (also known as total debt service, or "TDS").[1] Second, there is an "uninsured" sector, in which default risk on mortgages is not guaranteed. Mortgages in this sector are required to have an LTV ratio below 80%, but do not have any formal cap on PTI limits.

The resulting system is therefore captured by the diagram in Figure 1. This figure uses "I" to denote the insured submarket, "U" to denote the uninsured market, and θ

[1] A total debt service ratio is the sum of a gross debt service ratio (mortgage payment-to-income) plus other debt payments relative to income.

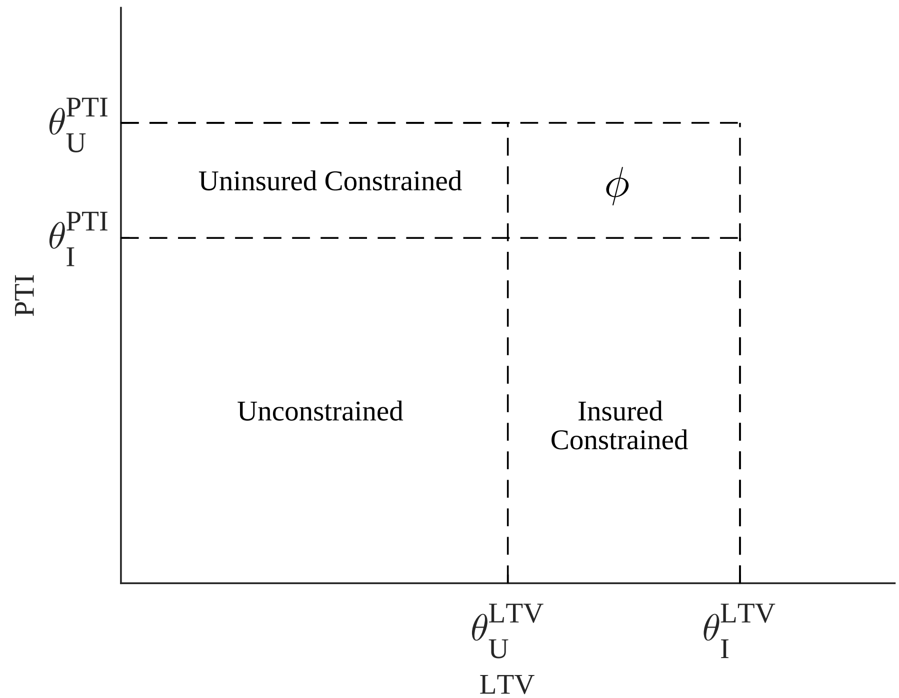

Figure 1: Borrowing Constraints by Sector

to denote a maximum ratio, so that e.g., θ_I^{LTV} indicates the maximum LTV ratio in the insured sector. According to our definitions of these submarkets, borrowers wishing to take out loans with a PTI ratio below θ_I^{PTI} and with an LTV ratio below θ_U^{LTV} are "unconstrained" meaning they do not hit either credit limit in either submarket, and are free to borrow from either. However, any borrower seeking to go beyond this rectangle will be restricted to a single market. In particular, borrowers seeking a PTI ratio higher than θ_I^{PTI} can only obtain such a loan in the uninsured sector, while borrowers seeking an LTV ratio higher than θ_U^{LTV} can only obtain such a loan in the insured sector. No loans can be given simultaneously violating both of these limits.

To make sure that this definition fits with real-world behavior, Figure 2 displays the LTV and PTI ratios for new purchase loans by sector. Corresponding almost perfectly with our definitions in Figure 1, these data confirm that nearly all high-LTV borrowers go to the insured sector, and among these, extremely few surpass the typical 44% PTI limit. Similarly, nearly all high-PTI borrowers go to the uninsured sector, and among these, virtually none take on more than 80% LTV. This pattern therefore validates our stylized definition of the two sectors, allowing us to proceed with construction of the model.

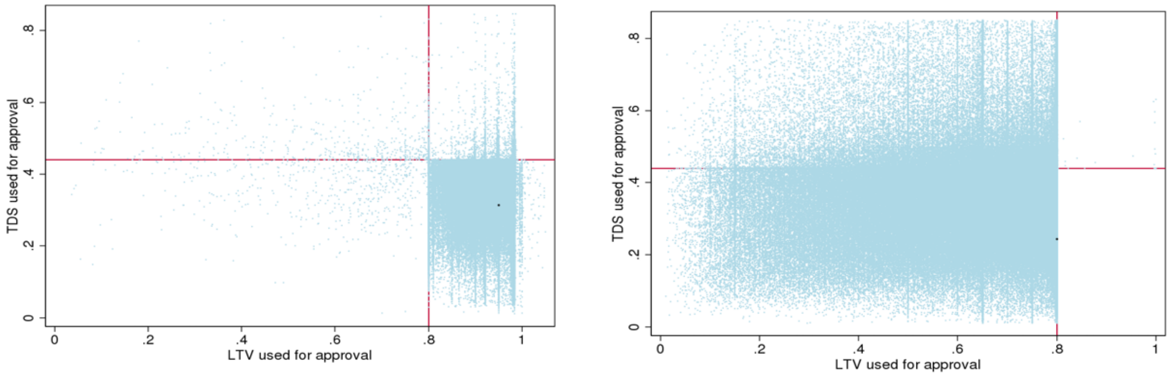

(a) Insured Sector

(b) Uninsured Sector

Figure 2: Borrower Ratios: Insured vs. Uninsured Sectors

3 Numerical Example

This section provides a numerical example demonstrating how this multiple-submarket structure influences debt limits and housing demand.

3.1 Debt Limits

We begin by illustrating how the multi-market structure influences borrowers' effective debt limits. For this example, we consider a borrower with an income of \$100,000, who faces an interest rate of 6% on a 30-year amortizing mortgage. Matching the data in Section 2, we assume a maximum LTV ratio of 95% in the insured sector, and a maximum LTV ratio of 80% in the insured sector. PTI ratios, in contrast, are limited to 44% in the insured sector but are unlimited in the uninsured sector. Since PTI ratios include outside debt obligations, we assume that the borrower has outside debt obligations equal to 5% of her income, leaving a maximum mortgage payment of 39% of her income in the insured sector.

Figure 3 display the maximum loan size this borrower can obtain for a given house value. Panel (a) displays the limits in each submarket. To begin, the red dotted line shows the maximum loan available in the insured submarket. For low enough house values, an insured loan is limited by the LTV constraint, which states that the loan cannot exceed 95% of the house's value. As a result, in this region, the maximum loan is increasing in house value with slope 0.95. However, once the house becomes valuable enough, a loan equal to 95% of the house's value would exceed the largest loan the borrower qualifies

Figure 3: Debt Limits

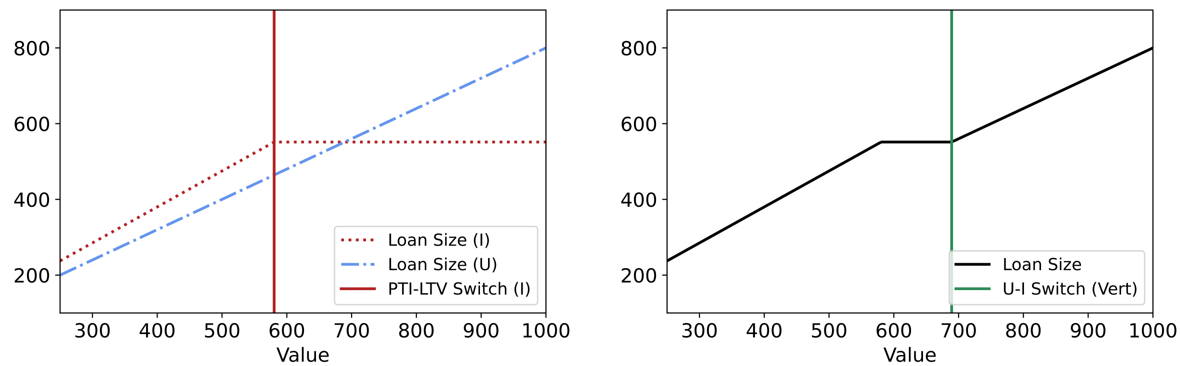

(a) Debt Limit by Submarket

(b) Overall Debt Limit

for under her PTI constraint, with the transition occurring at the point labeled "PTI-LTV Switch (I)." Since PTI limits are determined by interest rates and income, and do not depend on the house's value, the borrower cannot attain any additional credit in the insured market beyond this amount, so that the maximum loan size becomes flat.

Next, we can turn to the maximum loan size available in the uninsured submarket, displayed as the blue line in Panel (a). Since we have assumed that no PTI constraint exists in this submarket, the only relevant limit is the LTV constraint. As a result, the maximum loan size is equal to 80% of the house's value, generating an increasing line with slope 0.8. Because the LTV limit is tighter in this market, the uninsured market offers strictly lower maximum loan sizes for low house values. However, because the debt limit is strictly capped at the PTI limit in the insured market, but increases indefinitely in the uninsured market, the maximum loan size is strictly larger in the uninsured market for sufficiently large house values.

Since borrowers can typically choose between submarkets, the overall credit limit is the upper envelope of the two submarket limits. This overall debt limit is displayed in Figure 3 Panel (b). For low house values, a borrower seeking the largest possible loan will choose the insured market, so the overall limit is increasing in the region where this borrower is LTV constrained, then flat where the borrower is PTI constrained. However, once the borrower switches to the uninsured market, marked by the green vertical line in Panel 3b, the limit begins increasing again indefinitely.

3.2 Housing Demand Experiment

To understand the implications of this multi-market structure for the dynamics of a housing boom, we need to know how a borrower's choice set translates into actual housing purchases, which requires some assumptions on preferences. To this end, we assume that our example borrower is infinitely lived, with quasi-linear preferences and exogenous labor income, and intends to make a single housing purchase, then keep it forever. Her problem then becomes

$$V_0 = \max_{c,h,m} \; c + \sum_{t=0}^{\infty} \beta^t \xi \log(h)$$

subject to the budget constraint

$$c \leq \underbrace{\sum_{t=0}^{\infty} R^{-t} y_t}_{\text{labor income}} - \underbrace{(h - m)}_{\text{down payment}} - \underbrace{\sum_{t=1}^{T} R^{-t} q m}_{\text{mortgage payments}}$$

and the debt constraint $m \leq \bar{m}(h)$, where R is the gross interest rate, is an indicator for the mortgage still being paid, q is the coupon rate per unit of debt (interest plus principal) for a fully amortizing mortgage with interest rate R, m is the mortgage size, and $\bar{m}(h)$ is the function plotted in Figure 3 Panel (b).[2] Substituting in for c and m, and solving a geometric sum, we obtain

$$\begin{aligned} V_0 &= \sum_{t=0}^{\infty} R^{-t} y_t - (h - \bar{m}(h)) - \underbrace{\left(\sum_{t=1}^{T} R^{-t} q \right)}_{\kappa} \bar{m}(h) + \underbrace{\left(\frac{\xi}{1-\beta} \right)}_{\alpha} \log(h) \\ &= PV(y) + \alpha \log(h) - h + (1-\kappa)\bar{m}(h) \end{aligned}$$

for appropriately defined α and κ. Since labor income is constant regardless of the housing choice, we define the net present value of a particular house size choice as

$$NPV(h) = \underbrace{\alpha \log(h)}_{\text{PV benefit}} - \underbrace{(h - (1-\kappa)\bar{m}(h)}_{\text{PV cost}}.$$

[2] We have assumed without loss of generality that all consumption occurs at time 0 due to quasilinearity and that the price per unit of housing is equal to unity.

The "PV benefit" term represents the present value of housing services, while the "PV cost" represents the present value of all outlays paid for the house. Assuming the borrower has $\beta < R^{-1}$, the PV cost term differs from the time 0 price of the house h, because the present value of mortgage payments is lower than the number of dollars provided by the mortgage today. Differentiating the "PV benefit" and "PV cost" terms yields the marginal benefits and costs

$$MB(h) = \alpha h^{-1} \tag{1}$$

$$MC(h) = 1 - (1-\kappa)\bar{m}'(h) \tag{2}$$

where the marginal cost expression shows that additional housing is perceived to be cheaper when it can be used to collateralize additional debt ($\bar{m}'(h) > 0$), thus reducing the marginal cash outlay. We solve this model at quarterly frequency using an annual discount factor $\beta = 0.88$, implying $\kappa = 0.535$.

To understand how these features influence demand during a housing boom, Figure 4 displays the marginal cost, benefit, and NPV of each potential housing choice h. To make clear the contribution of the multi-market structure, the panel includes both the "Dual Market" system described above, as well a "Single Market" system in which the uninsured market does not exist, and all loans are obtained in the insured market, so that $\bar{m}(h)$ is equal to the red line in Figure 3 Panel (a). The top row displays these values for a baseline housing demand parameter $\alpha = 55$, while the bottom row displays these values for a higher "boom" housing demand parameter $\alpha = 70$.

To fix ideas, we will first consider the Single Market economy under baseline demand. Panel (a) displays the marginal benefit (1) in red, while the marginal cost under the single market is displayed as the solid blue line. Because the slope of $\bar{m}(h)$ drops from 0.95 to 0 when the borrower becomes PTI-constrained, this causes a discontinuous drop in $\bar{m}'(h)$, and therefore a discontinuous increase in $MC(h)$. As discussed at length in Greenwald (2018), this discontinuity implies that many borrowers will choose a house value that exactly equates their LTV and PTI limits, corresponding to the intersection of $MB(h)$ and $MC(h)$ in Panel (a), which can be seen to maximize $NPV(h)$ in Panel (b).

Incorporating the Dual Market structure does not change the marginal benefit curve $MB(h)$, but does change $\bar{m}(h)$ and therefore $MC(h)$. Specifically high enough house values h, the borrower will now switch to the uninsured market, increasing $\bar{m}'(h)$ from 0 to 0.8, and causing a discontinuous drop in the marginal cost curve, displayed as the dashed line in Panel (a). As a result, the marginal benefit curve can now intersect the marginal

Figure 4: Demand, Baseline vs. Boom Scenario

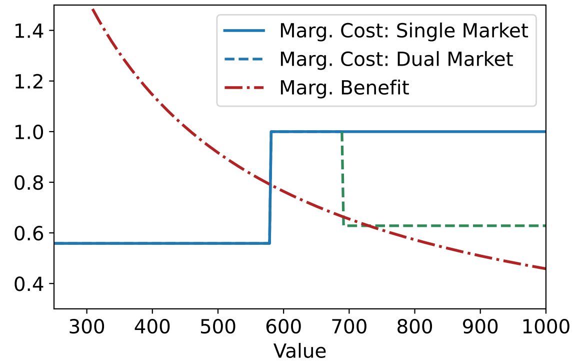

(a) Cost/Benefit (Baseline Demand)

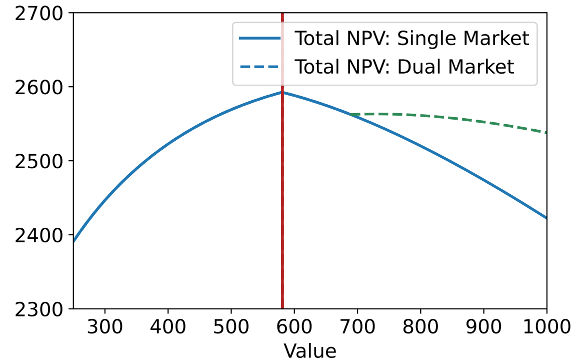

(b) Total NPV (Baseline Demand)

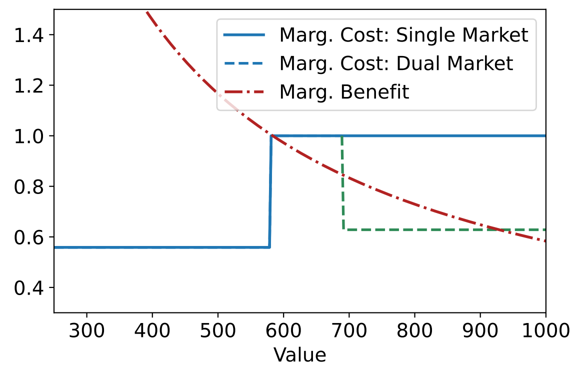

(c) Cost/Benefit (Boom Demand)

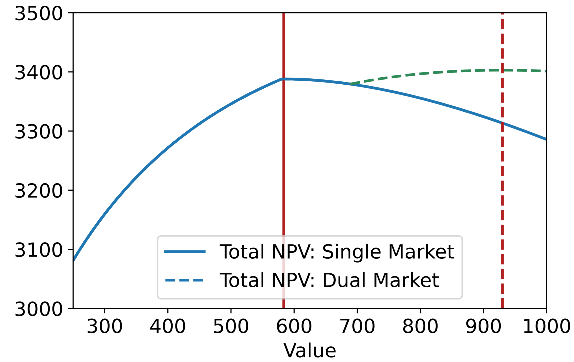

(d) Total NPV (Boom Demand)

Note: Blue-shaded lines above denote that the loan is obtained in the insured space, while green-shaded lines denote the uninsured space.

cost curve twice, as in Panel (a). Effectively, these jumps in marginal cost create a non-concave shape for $NPV(h)$, which includes two local maxima, one in the insured market, and one in the uninsured market. For the baseline parameters, however, the insured market option is still dominant, leading to the same housing allocation in the Dual Market problem as in the Single Market problem.

This changes, however, when housing demand increases. The bottom row of Figure 4 reproduces the figures displayed in the top row for $\alpha = 70$, which causes the $MB(h)$ and $NPV(h)$ curves to shift upward. In the Single Market economy, which has only a single potential solution, the increase in marginal benefit is not large enough to escape the discontinuous jump in marginal cost, leaving the borrower's optimal housing allocation unchanged. As in Greenwald (2018), these results implying that PTI limits are very effective at dampening housing booms in a single market environment.

In the Dual Market economy, however, this shift in demand changes the relative benefits of the two potential solutions. In particular, higher housing demand implies that the borrower now prefers to purchase a larger house financed by a mortgage from the uninsured sector, even though this requires making a substantially larger down payment. As a result, despite leaving the credit standards in each submarket unchanged, the borrower in the Dual Market sharply increases her housing demand. Aggregating over many borrowers in each economy, these results imply that the existence of a dual market, particularly the option to switch to a market without a PTI limit, can greatly expand the influence of housing preferences on housing demand and house prices, allowing for larger booms.

3.3 Tightening PTI Limits

Beyond investigating the impact of changes in preferences, we can consider the impacts of various macroprudential policy measures that exogenously tighten the various credit constraints. First, we can consider the impact of tightening the PTI constraint, which was found by Greenwald (2018) to be very effective at depressing house prices and credit in a single market setting. Figure 5 displays the effect of tightening PTI limits on the borrower's maximum loan size. The top row corresponds to an experiment tightening the maximum PTI ratio in the insured market from 44% to 34%. This tightening reduces the maximum loan size under a borrower's PTI limit, causing the insured market debt limit to flatten out earlier and at a lower level in Panel (a). Panel (b) shows that this change does lead to some tightening of the overall borrowing limit. However, because the uninsured market is not affected by this change, we find that not only do credit limits not tighten in the uninsured space, but that borrowers can switch from the insured to the uninsured market sooner, allowing even borrowers previously in the insured space to partially avoid the decline in credit limits.

The bottom row of Figure 5 displays the effect of tightening the uninsured PTI limit, decreasing it from an effective value of ∞ (no constraint) to a value of 49%. This reduces borrower credit limits, but in contrast to the experiment tightening the insured PTI limit, does not lead to any change in the U-I Switch point. This is because the uninsured space still has looser PTI limits relative to the insured space. As a result, households facing binding PTI limits in the uninsured sub-market are unable to access any additional credit by switching to the insured space, where PTI limits would be even tighter.

To understand the effects of this policy on housing demand, we can turn to Figure 6,

Figure 5: Debt Limit Effects, PTI Tightening

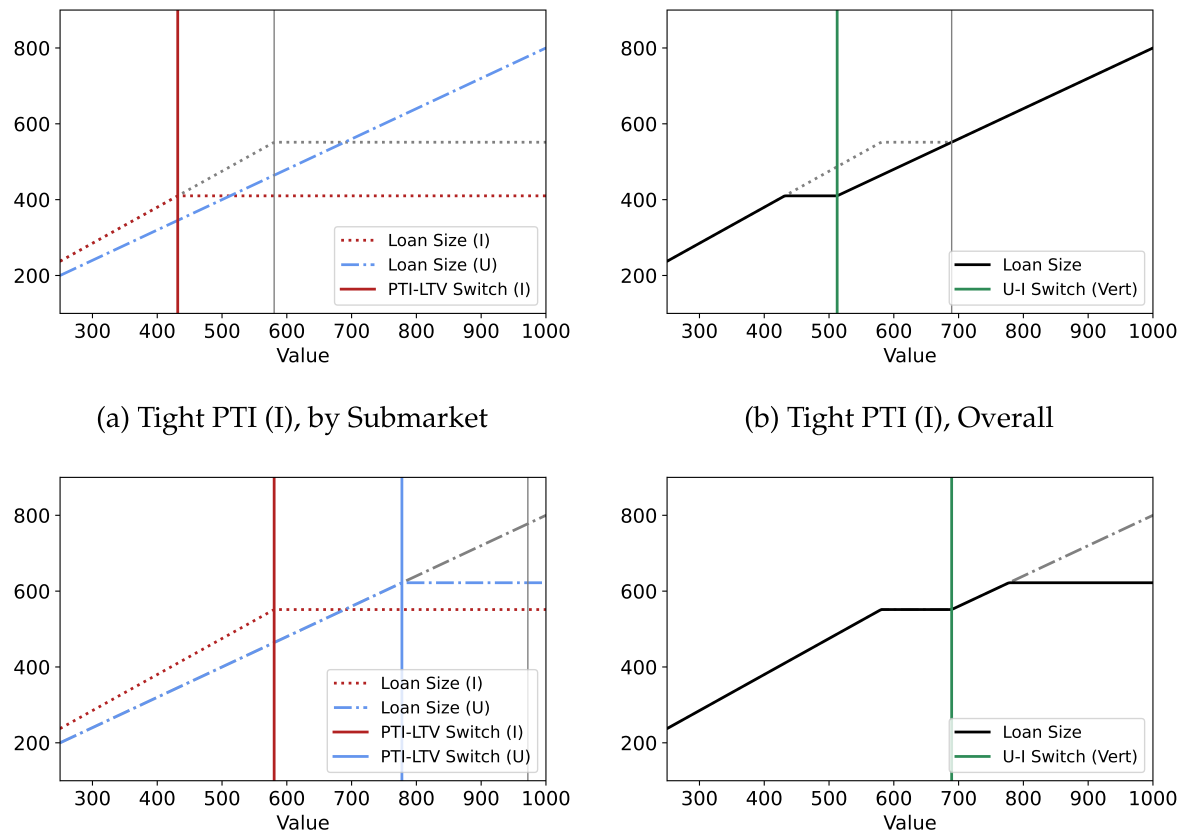

(a) Tight PTI (I), by Submarket

(b) Tight PTI (I), Overall

(c) Tight PTI (U), by Submarket

(d) Tight PTI (U), Overall

Note: Grey lines indicate the baseline case of Figure 3, while the remaining lines indicate the responses following a reduction in $\theta^{PTI}_{j,t}$.

which displays marginal benefits, costs, and NPVs for the baseline parameters, as well as under a tightening of PTI limits. To fix ideas, the top row shows the impact of tightening PTI in a single market structure, while the second and third rows show the impact of tightening the insured and uninsured submarket PTI ratios, respectively, in the dual market structure.

Under a single market structure (top row), this tightening of the PTI limit implies that the borrower becomes PTI-constrained at a lower house value. As a result, the marginal cost function jumps at an earlier value, leading the borrower to reduce her housing demand. These results imply, as in Greenwald (2018), that tightening PTI limits can very effectively reduce housing demand and house prices in a single market setting.

Under the dual market structure, however, the response is more complex. The second

Figure 6: Demand Effects, PTI Tightening

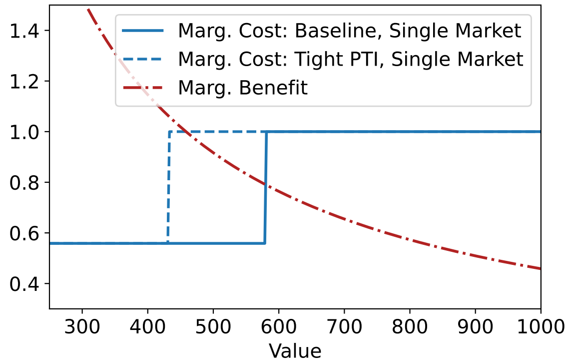

(a) Cost/Benefit (Tight PTI, Single)

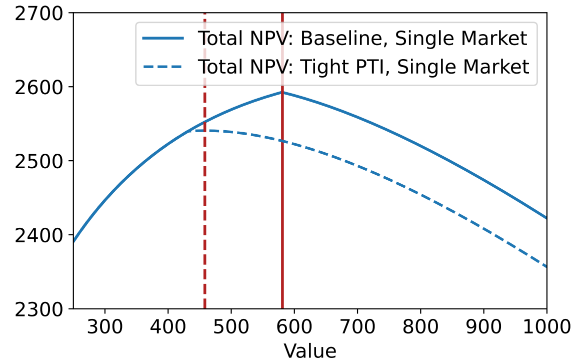

(b) Total NPV (Tight PTI, Single)

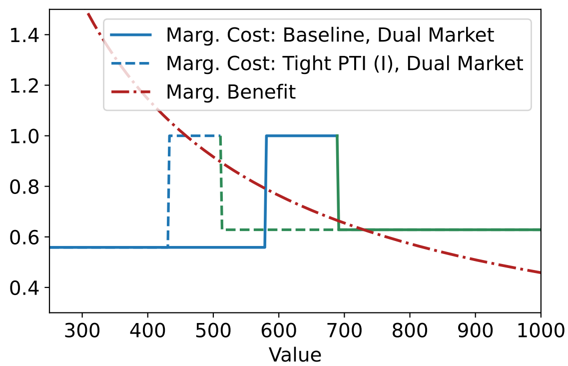

(c) Cost/Benefit (Tight PTI (I), Dual)

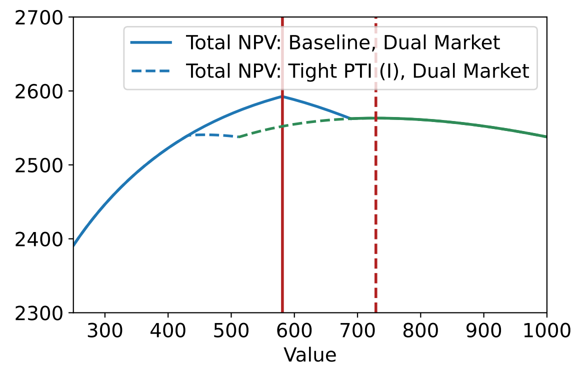

(d) Total NPV (Tight PTI (I), Dual)

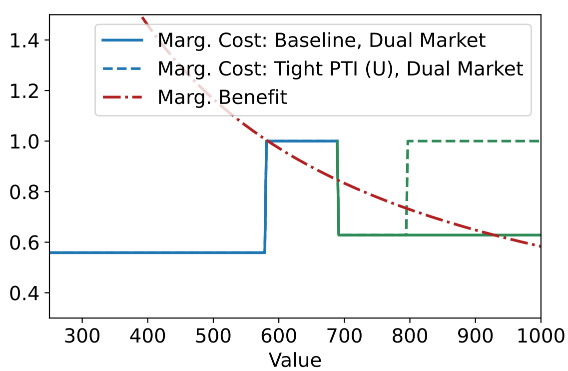

(e) Cost/Benefit (Tight PTI (U), Dual)

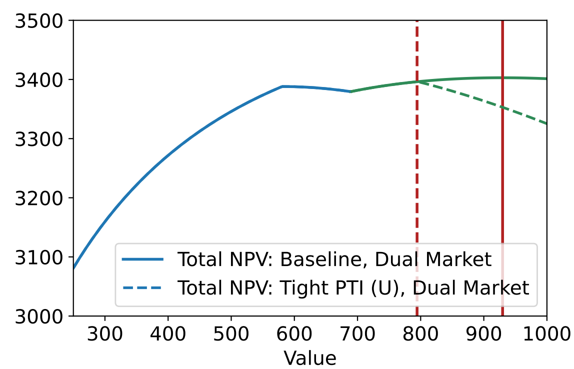

(f) Total NPV (Tight PTI (U), Dual)

Note: Blue-shaded lines above denote that the loan is obtained in the insured space, while green-shaded lines denote the uninsured space.

row of Figure 6 shows the impact of tightening the insured PTI ratio only. As shown in Figure 5, this shifts to the left both the discontinuous jump in marginal cost (where insured space borrowers become PTI constrained) and the discontinuous drop in marginal cost (where borrowers switch from being PTI-constrained in the insured space to being LTV-constrained in the uninsured space). As a result, each of the two potential solutions (intersections of MB and MC) shift left. However, because this tightening reduces the amount of available credit in the insured space, and therefore the NPV of purchasing a house to the left of the U-I Switch point in Figure 5, this policy can reorder the relative benefits of the insured and uninsured local maxima, leading the borrower to switch to the uninsured space, as shown in Panel (d). As a result, unlike in the single market case, tightening PTI limits can in fact cause a large *increase* in housing demand when it encourages borrowers to switch to a market with looser PTI and tighter LTV.

The third row of Figure 6 displays the effects of imposing a PTI limit in the uninsured market. To create a case where this is relevant to the borrower, we assume that the borrower has the high housing demand $\alpha = 70$, so that she previously preferred the uninsured market as in Figure 4 Panel (d). In this case, adding a PTI limit generates a second jump in the marginal cost curve in Figure 6 Panel (e) as the borrower becomes PTI-constrained in the uninsured space, flattening the $\bar{m}(h)$ curve. In our example, this causes the local optimum in the uninsured space to shift left, reducing borrower housing demand. For a large enough tightening, it could even incentivize the borrower to switch back to the local optimum in the insured space, causing an even large drop in housing demand. Either way, these results show that tightening the PTI limit in the uninsured space should have an unambiguously negative effect on housing demand, as in the single market case.

3.4 Tightening LTV Limits

We now repeat this exercise, tightening LTV limits (increasing minimum down payment ratios) by 10pp for both the single and dual market structures. As before, Figure 7 displays the updated debt limits following a tightening in the insured sector (top row) and in the uninsured sector (bottom row). Beginning with the insured sector, Panel (a) shows that a tightening reduces debt limits for LTV-constrained borrowers in the insured space, but also shifts to the right the point at which borrowers switch from being LTV-constrained to being PTI-constrained, since the maximum loan size increases more slowly in the LTV-constrained region. Panel (a) shows that this reduces debt limits for low house values, but has no influence on the relative orderings of debt limits in the two markets. The intuition, similar to that provided for a tightening of uninsured PTI, is that LTV constraints are still looser in the insured market than in the uninsured market, implying that LTV-constrained borrowers in the insured sector cannot obtain additional credit by switching to the uninsured sector, where constraints are even tighter.

Figure 7: Debt Limit Effects, LTV Tightening

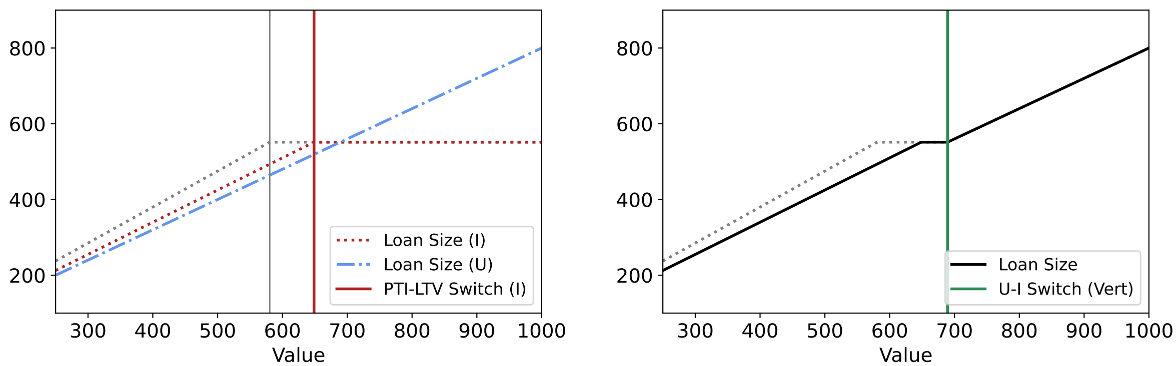

(a) Debt Limit by Submarket ($\theta_I^{LTV} \downarrow$) (b) Overall Debt Limit ($\theta_I^{LTV} \downarrow$)

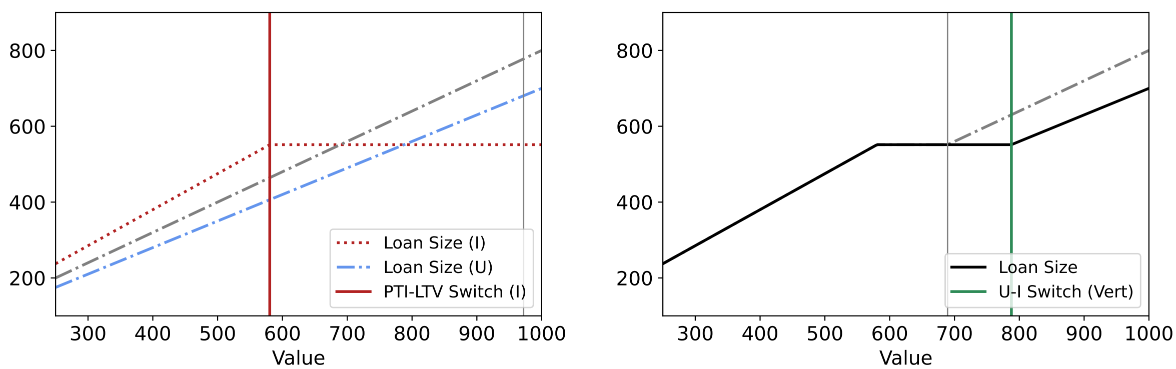

(c) Debt Limit by Submarket, ($\theta_U^{LTV} \downarrow$) (d) Overall Debt Limit, ($\theta_U^{LTV} \downarrow$)

Note: Grey lines indicate the baseline case of Figure 3, while the remaining lines indicate the responses following a reduction in $\theta_{j,t}^{LTV}$.

The bottom row of Figure 7 shows the impact of tightening LTV limits in the uninsured space. This policy tightens credit limits for uninsured borrowers, but also expands the region in which the insured market offers larger loan values, since borrowers who find themselves newly constrained by LTV limits in the uninsured sector may be able to obtain additional credit under the loose LTV limits of the insured sector.

For the impact of this policy on housing demand, we now turn to Figure 8. As before, we present the effect on single market in the top row, and the effect of tightening the insured and uninsured market ratios in a dual market setting in the second and third

Figure 8: Demand Effects, LTV Tightening

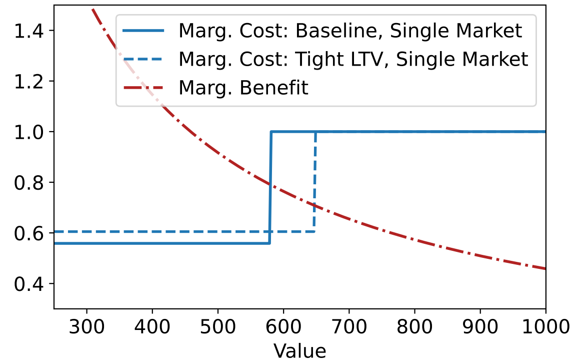

(a) Cost/Benefit (Tight LTV, Single)

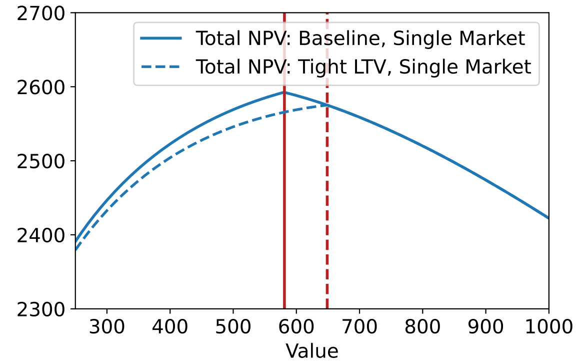

(b) Total NPV (Tight LTV, Single)

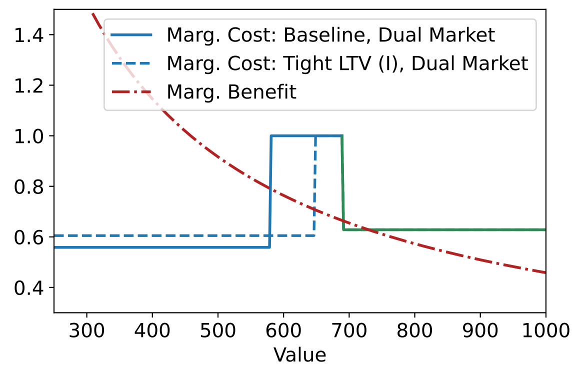

(c) Cost/Benefit (Tight LTV (I), Dual)

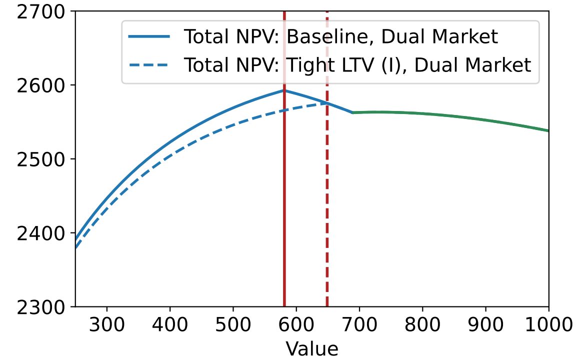

(d) Total NPV (Tight LTV (I), Dual)

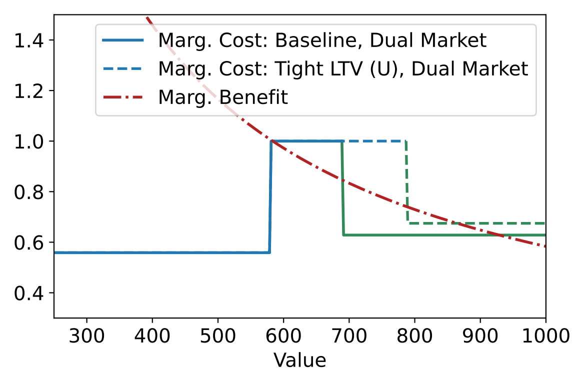

(e) Cost/Benefit (Tight LTV (U), Dual)

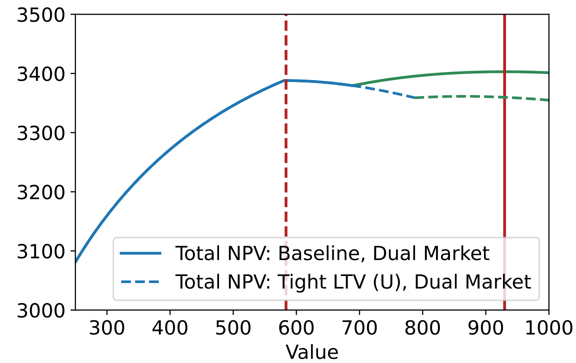

(f) Total NPV (Tight LTV (U), Dual)

Note: Blue-shaded lines above denote that the loan is obtained in the insured space, while green-shaded lines denote the uninsured space.

rows.

Beginning with the top row (Single Market), we observe that the LTV tightening has two effects on marginal cost. First, it raises marginal cost slightly in the LTV-constrained (left) region, since lower LTV maps directly into a lower slope of $\bar{m}(h)$ and thus lower collateralizability through $\bar{m}'(h)$. Second, a lower LTV ratio shifts the value at which borrowers become PTI-constrained to the right. Because the threshold value at which the LTV-PTI switch occurs is equal to the maximum PTI loan size plus the minimum down payment, a larger minimum down payment (tighter LTV limit) increases this value, even though the borrower's PTI limits have not changed. Since marginal cost now jumps discontinuously at a higher value, this can cause borrowers constrained by both limits (as in this example) to *increase* their housing demand, as in Greenwald (2018). Panel (b) shows that, while tightening LTV has made borrowers worse off, it has done so disproportionately at lower house values, pushing the optimal housing allocation up. As a result, Greenwald (2018) finds that tightening LTV limits is often an ineffective method for dampening housing booms in a single market setting.

Turning to the second row, we see that the effect of tightening the insured LTV limit in the Dual Market economy is very similar to the effect in the single market. Because this policy tightens LTV limits where they are currently loosest, there is no incentive to switch between markets. As a result, we observe the same shift in marginal cost, with the discontinuous jump, and optimal policy, once again pushed to the right.

Last, the third row displays the effect of tightening the uninsured limit in the Dual Market economy. Again, to make this shift relevant to the borrower, we begin in the environment with elevated demand ($\alpha = 70$) so that the borrower is currently choosing an uninsured mortgage. As before, this has two effects on the marginal cost curve in Panel (e). First, in the rightmost region, it causes a slight upward increase in marginal cost, due to the same mechanism as in the other two rows. Unlike in the other examples, however, this change in LTV limits pushes the discontinuous fall in marginal cost — caused by the switch between the insured and uninsured markets — to the right. Because this switch occurs when the uninsured market can offer larger loans than the insured market, a reduced LTV ratio in the uninsured market causes this switch to occur at a higher house value, explaining the rightward shift.

Importantly, however, this change in LTV limits also influences the relative benefits of the insured and uninsured local optima. In particular, tighter credit in the uninsured market pushes down the NPV curve in the uninsured market (Panel (e)). For borrowers close to marginal between the two solutions, this can push the borrower to switch to the insured space, leading to a large drop in housing demand. Thus, unlike LTV policy in a single market or in the insured space, reducing LTV limits in the uninsured space (where they are initially tighter) can potentially be a more effective candidate for dampening house prices during a boom.

Summary. We conclude our numerical example by summarizing the main findings. In a single market setting, binding PTI limits can prevent housing demand from rising, even when housing utility is increased during a boom episode. However, in a dual market setting, the ability to switch to an uninsured market without a PTI limit can lead borrowers to demand much more housing in such an episode, leading to larger house price and credit booms.

While tightening PTI ratios is an effective way to dampen housing demand in a single market setting, as does tightening the uninsured PTI ratio in the dual market (where PTI limits are loosest), tightening the insured PTI ratio in the dual market (where PTI limits are tightest) can actually increase housing demand. This occurs because the tightening incentivizes some borrowers to switch to the uninsured market, where the ability to use housing to collateralize additional debt at the margin generates a large increase in housing demand.

Last, tightening LTV ratios does not effectively dampen housing demand in a single market setting, or when the LTV limit is tightened in the insured market (where LTV limits are tightest). In these settings, tighter LTV limits imply that borrowers only become PTI-constrained at higher valuations, so that borrowers can continue to use housing to collateralize more debt, pushing up housing demand. However, tightening LTV ratios in the uninsured space (where LTV limits are tightest) can substantially reduce housing demand to the extent that it leads borrowers to switch to the insured market, where the local optimum for housing demand is much lower.

The common thread between these results is that tightening constraints in submarkets where that constraint is the loosest (the LTV limit in the insured market, the PTI limit in the uninsured market) has effects similar to those of a single market setting. The reason is that, because the constraints are even tighter in the alternative market, this change does not provide any incentives to switch markets. However, tightening constraints in the submarkets where that constraint is the tightest (the PTI limit in the insured market, the LTV limit in the uninsured market) incentivizes borrowers to avoid the newly binding constraint by switching to the alternative market, where it is looser. Since the local optimum for housing demand is much higher in the uninsured space, these switches cause very different implications for housing demand, so that tightening insured PTI can increase housing demand for these borrowers, while tightening uninsured LTV can decrease it.

4 Model

This section translates the intuition from Section 3 into a more rigorous structural setting. The overall approach follows Greenwald (2018), with the main innovations appearing in the treatment of the mortgage submarkets.

Demographics. The economy is populated by two families of agents — borrowers and savers — who are denoted by the subscripts b and s, respectively. Agents are infinitely lived and types are permanent, with a fixed measures χ_b of borrowers and $\chi_s = 1 - \chi_b$ of savers. A complete set of contracts over consumption and housing services can be traded within each family, but not across families.

Preferences. Each agent of type $j \in \{b, s\}$ maximize expected utility of the form

$$V_j = \mathbb{E}_t \sum_{k=0}^{\infty} \beta_j^k u(c_{j,t+k}, h_{j,t+k}, n_{j,t+k}) \tag{3}$$

where the inputs to the utility function are: nondurable consumption c, housing services h, and labor supply n. Period utility takes the separable form

$$u_j(c, n, h) = \log(c) + \xi \log(h) - \eta_j \frac{n^{1+\varphi}}{1 + +\varphi}.$$

For notation, let e.g., $u^c_{j,t}$ denote the derivative of the utility function of agent j with respect to c, and let

$$\Lambda_{j,t+1} = \beta_j \frac{u^c_{j,t+1}}{u^c_{j,t}}$$

denote the stochastic discount factor of an agent of type j.

Mortgages. The mortgage sector consists of two submarkets: one for government insured mortgages, denoted I, and one for uninsured mortgages, denoted U. Borrowers obtaining new loans choose freely which submarket they prefer to enter. A mortgage in sector j is a nominal perpetuity with a fixed interest rate and geometrically decaying coupon. This means that the borrower in submarket j pays back a constant fraction ν_j of the remaining principal balance in each period, so that the payment at time $t+k$ on \$1 of debt issued at time t is $\$(1-\nu_j)^k(r_t^* + \nu_j)$ for all k until the loan is renewed, where r_t^* is the fixed interest rate. Each mortgage in submarket j is renewed each period with probability ρ_j, at which time the borrower prepays her existing balance and can take out a new loan.

To induce changes in real mortgage rates similar to shifts in mortgage spreads or term premia, we introduce a proportional tax $\Delta_{q,t}$ on all future mortgage payments received by the saver on loans in submarket j originated at time t, subject to the process

$$\Delta_{q,t} = (1-\phi_q)\mu_q + \phi_q \Delta q, t-1 + \varepsilon_{q,t}.$$

This tax does not correspond to any real-world policy, but is instead a parsimonious way to create a wedge between long rates and average future short rates, which is needed to match the large and volatile discrepancy between these two items in the data. As a result, we rebate the proceeds from this tax back to the savers in lump sum fashion each period.

The size of a new loan is limited by both an LTV constraint and a PTI constraint, defined by

$$\frac{m_{i,t}^*}{p_t^h h_{i,t}^*} \leq \theta_{,t}^{LTV} j,t, \qquad \frac{(r_{j,t}^* + \nu_j + \alpha)m_{i,t}^*}{w_t n_{i,t} e_{i,t}} + \omega \leq \theta_{,t}^{PTI} j,t$$

respectively. The LTV constraint caps the ratio of the balance on the new loan $m_{i,t}^*$, against that borrower's housing collateral $p_t^h h_{i,t}^*$, where $h_{i,t}^*$ is the quantity of newly purchased housing. These constraints are applied at origination only.

The PTI constraint caps the ratio of the borrower's debt and related payments to her income. The numerator on the left hand side is the initial mortgage payment, where the offset term α is used in calibration to adjust for property taxes and insurance payments, as well as to adjust for the difference in amortization between true fixed-rate mortgages and the model's geometrically decaying coupon loans. The denominator is equal to labor income, which is shifted for each borrower by the idiosyncratic productivity shock $e_{i,t}$, drawn i.i.d. across borrowers and time from a distribution with c.d.f. Γ_e. Due to complete markets within the borrower family, this income shock has no impact on borrower consumption allocations, but is instead used to create variation among borrowers, allowing for endogenous sorting into the two submarkets, and ensuring that endogenous fractions are limited by each constraint. Finally, ω is used in calibration to adjust for debt obligations other than mortgages.

These constraints imply the maximum loan balances

$$\bar{m}^{LTV}_{i,j,t} = \theta^{LTV}_{j,t} p^h_t h^*_{i,t} \qquad\qquad \bar{m}^{PTI}_{i,j,t} = \frac{(\theta^{PTI}_{j,t} - \omega) w_t n_{i,t} e_{i,t}}{r^*_{j,t} + \nu_j + \alpha}$$

which define the maximum that borrower i can obtain in submarket j under each limit. Since both limits must be satisfied simultaneously, the maximum loan balance in submarket j is defined by $\bar{m}_{i,j,t} = \min(\bar{m}^{LTV}_{i,j,t}, \bar{m}^{PTI}_{i,j,t})$.

At equilibrium, the optimal policy will be to choose the insured space — which has looser income-based constraints and tighter collateral-based constraints — if and only if $e_{i,t}$ exceeds an endogenous threshold e^*_t. We define $F_{j,t}$ to be the fraction of borrowers who choose sector j, so that $F_{U,t} = \Gamma_e(e^*_t)$ and $F_{I,t} = 1 - \Gamma_e(e^*_t)$.

Housing. The final asset in the economy is housing, which is divisible in fixed total supply $\bar{H}$ and produces a service flow equal to its stock. To focus on the use of housing as a collateral asset, we fix the saver's demand for housing at $h_{s,t} = \bar{H}_{s,t}$, which ensures that the borrower is the marginal pricer of housing at equilibrium.

Taxation. Each household's labor income is subject to proportional taxation at rate τ. Tax revenues are returned to borrowers and savers in lump-sum transfers equal to the average amount paid by that type.

Bonds. We introduce a risk-free one-period bond that can be used as a policy instrument by the central bank. An investment of \$1 at time t yields a guaranteed nominal payoff of $\$R_t$ at time $t+1$. This bond is in zero net supply and cannot be shorted, which means that it will be held by savers only in equilibrium.

Representative Borrower's Problem. The individual borrower's problem aggregates to that of a representative borrower. This representative borrower's endogenous state variables are the principal balances $m_{j,t-1}$ and promised interest payments $x_{j,t-1}$ and promised payments $x_{j,t-1}$ for each submarket $j \in \{I, U\}$, as well as total borrower housing $h_{b,t-1}$, which follow the laws of motion

$$m_{j,t} = \rho F_{j,t} m^*_{j,t} + (1-\rho)(1-\nu_j)\pi_t^{-1} m_{j,t-1} \tag{4}$$

$$x_{j,t} = \rho F_{j,t} r^*_{j,t} m^*_{j,t} + (1-\rho)(1-\nu_j)\pi_t^{-1} x_{j,t-1}, \tag{5}$$

where $F_{j,t}$ is the fraction of borrowers taking out new loans who choose submarket j. The representative borrower chooses nondurable consumption $c_{b,t}$, labor supply $n_{b,t}$, the size of newly purchased houses $h^*_{b,t}$, and the face value of newly issued mortgages in each sector m^*_t, and the income threshold at which to send borrowers to the insured sector e^*_t to maximize (3) using the aggregate utility function

$$u_b(c_{b,t}, h_{b,t}, n_{b,t}) = \log(c_{b,t}/\chi_b) + \xi \log(h_{b,t}/\chi_b) + \eta_j \frac{n_{b,t}^{1+\varphi}}{1+\varphi}$$

subject to the budget constraint

$$c_{b,t} \leq \underbrace{(1-\tau_y) w_t n_{b,t}}_{\text{labor income}} - \underbrace{\rho p^h_t (h^*_{b,t} - h_{b,t-1})}_{\text{housing purchases}} - \underbrace{\delta p^h_t h_{b,t-1}}_{\text{maintenance}} + \underbrace{T_{b,t}}_{\text{transfers}}$$
$$+ \sum_{j\in\{I,U\}} \left\{ \underbrace{\rho \left(m^*_{j,t} - (1-\nu_j)\rho \pi_t^{-1} m_{j,t-1} \right)}_{\text{net new debt issuance}} - \underbrace{\pi_t^{-1}(1-\tau_y) x_{j,t-1}}_{\text{interest payment}} - \underbrace{\pi_t^{-1} \nu_j m_{j,t-1}}_{\text{principal payment}} \right\}$$

and the aggregate borrowing constraints

$$m^*_{U,t} \leq \int^{\bar{e}_{U,t}} \bar{m}^{PTI}_{U,t} e_i \, d\Gamma_e(e_i) + \bar{m}^{LTV}_{U,t} \left(\Gamma_e(e^*_t) - \Gamma_e(\bar{e}_{U,t}) \right)$$
$$m^*_{I,t} \leq \int_{e^*_t}^{\bar{e}_{I,t}} \bar{m}^{PTI}_{I,t} e_i \, d\Gamma_e(e_i) + \bar{m}^{LTV}_{I,t} \left(1 - \Gamma_e(\bar{e}_{I,t}) \right).$$

where

$$\bar{e}_{j,t} = \frac{\bar{m}^{LTV}_{j,t}}{\bar{m}^{PTI}_{j,t}}, \quad j \in \{U, I\}$$

represent the threshold income levels in each submarket at which a borrower crosses from being PTI-constrained to being LTV-constrained.

Representative Saver's Problem. The individual saver's problem similarly aggregates to that of a representative saver. This representative saver's problem largely parallels the representative borrower's problem. The state variables are the balances and promised payments on each type of loan $m_{j,t}$ and $x_{j,t}$, again following (4) and (5). The representative saver chooses nondurable consumption $c_{s,t}$, labor supply $n_{s,t}$, and the average balance on newly issued mortgages m_t^*. [TBC]

Productive Technology. The productive technology follows the standard New Keynesian assumptions of a competitive final goods producer and a continuum of monopolistically competitive intermediate goods producers. Both types of firm are owned by the saver at equilibrium. The final goods producer solves the static problem

$$\max_{y_t(i)} P_t \left(\int y_t(i)^{\frac{\lambda-1}{\lambda}} \, di \right)^{\frac{\lambda}{\lambda-1}} - \int P_t(i) y_t(i) \, di,$$

where $y_t(i)$ is the intermediate good produced by firm i, $P_t(i)$ is the price of that good, and P_t is the price of the final good.

Intermediate firms choose their price $P_t(i)$, and operate the linear production function

$$y_t(i) = a_t n_t(i)$$

to meet the demand of the final good producer, where $n_t(i)$ denotes labor hours employed by firm i, and a_t is total factor productivity (TFP), which follows an AR(1) in logs:

$$\log a_{t+1} = (1 - \phi_a)\mu_a + \phi_a a_t + \varepsilon_{a,t+1}.$$

Intermediate goods producers are subject to pricing frictions of the Calvo-Yun form, meaning that fraction ζ of firms cannot adjust their price in a given period, while the other $1 - \zeta$ fraction are free to do so.

Monetary Authority. The monetary authority follows a Taylor rule similar to the specification in Smets and Wouters (2007), given by

$$\begin{aligned} \log R_t = \log \bar{\pi}_t + \phi_r \left(\log R_{t-1} - \log \bar{\pi}_{t-1} \right) \\ + (1 - \phi_r) \left[\left(\log R_{ss} - \log \pi_{ss} \right) + \psi_\pi \left(\log \pi_t - \bar{\pi}_t \right) \right], \end{aligned}$$

where the subscript "ss" indicates the steady state value. The inflation target $\bar{\pi}_t$ follows an AR(1) in logs

$$\log \bar{\pi}_t = (1-\phi_{\bar{\pi}}) \log \pi_{ss} + \phi_{\bar{\pi}} \log \bar{\pi}_{t-1} + \varepsilon_{\bar{\pi},t}.$$

These shocks correspond to near-permanent changes in inflation that can cause nominal rates to move relative to real rates.

4.1 Model Solution

This section provides the key equilibrium conditions of the model. The full set of conditions will be available in the complete draft of the paper.

For the dynamics of debt and the allocation of borrowers across submarkets, the key equilibrium relation is the borrower's optimality condition for e_t^*

$$\mu_{U,t} \bar{m}_{U,t}^{LTV} = \mu_{I,t} \bar{m}_{I,t}^{PTI} e_t^*$$

which sets the benefit to the threshold borrower of taking out a loan to be equal in both submarkets, under the assumption (verified at equilibrium) that the threshold borrower would be LTV-constrained in the uninsured space, but PTI-constrained in the insured space. This net benefit is equal to the product of the size of the loan that could be taken out and the benefit of an additional dollar of debt in that submarket, which is equal to the multiplier on the debt limit for that submarket. Solving yields the threshold value

$$e_t^* = \frac{\mu_{U,t} \bar{m}_{U,t}^{LTV}}{\mu_{I,t} \bar{m}_{I,t}^{PTI}}.$$

We consider the case in which the amortization rates ν_j are identical across submarkets, which implies $\mu_{I,t} = \mu_{U,t}$. In this case, optimality implies that borrowers should simply take whichever submarket offers them a bigger maximum loan size, yielding $e_t^* = \bar{m}_{U,t}^{LTV} / \bar{m}_{I,t}^{PTI}$.

For the dynamics of house prices, we can turn to the borrower's optimality condition for new housing:

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

The term $\mathcal{C}_t$ represents the marginal collateral value of housing, defined by

$$\mathcal{C}_t = \sum_{j \in \{I,U\}} \mu_{j,t} F_{j,t}^{LTV} \theta_{j,t}^{LTV} \tag{6}$$

where $F_{j,t}^{LTV}$ is the fraction of borrowers who are in submarket j and constrained by LTV, computed as

$$\begin{aligned} F_{U,t}^{LTV} &= \Gamma_e(e_t^*) - \Gamma_e(\bar{e}_{U,t}) \\ F_{I,t}^{LTV} &= 1 - \Gamma_e(\bar{e}_{I,t}). \end{aligned}$$

This term $\mathcal{C}_t$ equals the value to the borrower from an extra dollar of average housing wealth due to the extra debt it allows her to obtain. An extra dollar of housing wealth relaxes LTV constraints in submarket j by $\theta_{j,t}^{LTV}$. But since not all borrowers are LTV-constrained, this relaxation only loosens overall credit limits by $F_{j,t}^{LTV}\theta_{j,t}^{LTV}$ in each submarket. Finally, each additional dollar of credit is valued at rate $\mu_{j,t}$, the multiplier on the credit constraint in submarket j. In the special case with only a single mortgage sector, (6) would collapse to the condition

$$\mathcal{C}_t = \mu_t F_t^{LTV} \theta_t^{LTV} \tag{7}$$

from Greenwald (2018), where μ_t is the multiplier on the single budget constraint, and F_t^{LTV} is the total LTV-constrained fraction of the population.

An increase in the collateral value $\mathcal{C}_t$ raises house prices and price-rent ratios. This gives rise to a "constraint switching" effect, as in Greenwald (2018). Specifically, since LTV-constrained borrowers are willing to pay a premium for additional housing collateral, which they can use to relax their borrowing limits, an increase in the fraction of LTV-constrained borrowers — here $F_{U,t}^{LTV}$ or $F_{I,t}^{LTV}$ — increases housing demand. This maps directly into the examples from Section 3, where LTV-constrained borrowers faced a strictly positive $\bar{m}'(h)$, reducing their marginal cost and increasing housing demand, while PTI-constrained borrowers have $\bar{m}'(h) = 0$, increasing their marginal cost and reducing housing demand.

Incorporating submarkets, however, allows for additional dynamics through the market-switching cutoff e_t^*. Recall that borrowers with income shocks just below e_t^* will be LTV-constrained in the uninsured market, while borrowers with income shocks just above e_t^* will be PTI-constrained in the insured market. As a result, when e_t^* increases, meaning that borrowers are switching from the insured to the uninsured space, those borrowers will also be switching from being PTI-constrained to being LTV-constrained. This increases $\mathcal{C}_t$, pushing up housing demand. Through this force, flows into the uninsured space to increase house prices and price-rent ratios, while flows out of the uninsured space reduce them. We denote the share of borrowers choosing the uninsured space as $F_{U,t} = F_{U,t}^{LTV} + F_{U,t}^{PTI}$, and note that equilibrium house prices are generally increasing in this value.

4.2 Calibration

The calibration is a work in progress, but largely follows the methodology of Greenwald (2018) adapted to Canadian data. Calibrated values can be found in Table 1.

5 Results

We now present the quantitative implications of our structural model. These can be considered as the general equilibrium equivalents of the simpler numerical examples considered in Section 3.

5.1 Housing Boom Experiment

To begin, we simulate a housing boom environment to observe the influence of the dual market structure on the economy's response. To replicate a boom in prices relative to rents, we increase expectations of future housing utility, which could also stand in for overoptimism or bubbly beliefs on the part of homebuyers. Specifically, we run an experiment in which agents learn in 2001:Q1 that the housing utility parameter ξ will increase by 120% after an 80Q delay, chosen to match the growth in the value-to-income ratio in Canada from 2001 to 2016. The solutions are then computed as nonlinear perfect foresight paths.

To demonstrate the influence of the dual market structure, we compare our Benchmark economy, described above, with two alternative single market economies. The first, labeled "All Insured," represents an economy in which the uninsured market does not exist, and where the only PTI and LTV limits are those of the insured market. Symmetrically, the "All Uninsured" economy represents an economy in which the insured market does not exist, and all borrowers obtain uninsured mortgages.

Figure 9: Housing Boom Experiment

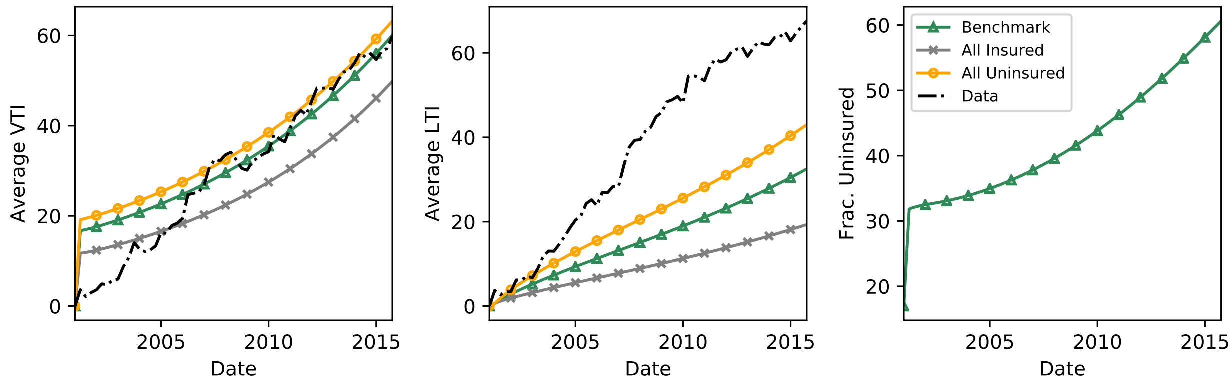

Note: Average VTI is the ratio of total housing value to total income in the economy: $p^h_t \bar{H}/y_t$. Average LTI is the ratio of total mortgage debt to total income in the economy: $(m_{I,t} + m_{U,t})/y_t$. Frac. Uninsured is the share of new mortgage borrowers choosing mortgages from the uninsured market $F_{U,t}$. Black dash-dotted lines indicate the data, while colored lines indicate model experiments. Average VTI and Average LTI are presented as percent deviations from the initial steady state, while Frac. Uninsured is presented as a percentage in levels.

The results of this experiment for the three economies just described are displayed in Figure 9. As can be seen, this very large increase in the expected future value of housing translates into sustained house price growth in all three economies. To understand the variation between the paths, recall that rising prices increase the loan sizes allowed under all LTV constraints. In the All Insured, which imposes a strict PTI limit on all borrowers, this leads more borrowers to switch from being LTV-constrained to being PTI-constrained, at the point where marginal cost jumps in Figure 4 Panel (c). Since this jump in marginal cost partially constrains increases in housing demand, binding PTI limits dampen house price growth in this economy. For an alternative explanation following the intuition in Section 4, we note that the share of borrowers constrained by LTV drops, reducing F^{LTV}_t and depressing collateral values in equation (6). In contrast, the All Uninsured model does not impose any PTI limit. As a result, this constraint switching effect is never active, and we observe a larger increase in prices and VTI ratios.

Figure 9 shows that the Benchmark model, containing both sectors, is unsurprisingly between the two. In this world, LTV-constrained borrowers in the insured sector initially become PTI-constrained, partially dampening their increase in demand. However, as their demand increases further, they eventually switch to the uninsured market, increasing their housing demand, and undoing this dampening effect. Interestingly, the path of price and VTI ratios in the Benchmark economy is much closer to that of the All

Uninsured economy to that of the All Insured economy, even though more than 80% of borrowers choose mortgages from the insured sector in the initial steady state. However, the large rise in the fraction of unconstrained mortgages $F_{U,t}$ (right panel of Figure 9), representing a shift in mass from $F_{I,t}^{PTI}$ to $F_{U,t}^{LTV}$ increases housing demand through the market switching effect, leading to an outsized effect on house prices.

These results not only create variation in house price growth, but also in credit growth. This occurs through two channels. First, there is a direct channel through the debt limits themselves. For example, some PTI-constrained borrowers in the All Insured economy would be able to obtain larger uninsured loans in the Benchmark economy, increasing debt levels. Second, however, is an indirect effect through house prices. In all three economies, most borrowers are LTV-constrained. Since rising house prices increase the maximum loan sizes allowed under LTV limits, additional house price growth further increases credit growth in each economy. Combined, these forces imply that growth in the loan-to-income (LTI) ratio is much higher in the Benchmark economy than in the single market All Insured economy, and is once again closer to the All Uninsured economy.

All told, the structural model reproduces the intuition of Section 3, finding that the dual-market structure allows the same increase in expected housing utility to generate a larger boom in house prices and credit in the Benchmark economy compared to its single-market All Insured counterpart.

5.2 Tightening PTI Limits

We next turn to the experiments tightening PTI limits, as in Section 3.3. Our experiment is based on an October 2016 policy in Canada that PTI ratios in the insured submarket must be evaluated at the "posted" interest rate rather than the "contract" interest rate. Since the posted interest rate is roughly 200bp higher, on average, this effectively translates into a 16.5% tightening of the PTI limit. We impose this tightening as a one-time unanticipated change, beginning from the model's steady state, and trace out the nonlinear perfect foresight path as the economy transitions to the new steady state. To clarify the contribution of the multi-market structure, we again produce responses both for the Benchmark economy, as well as for a No Uninsured economy in which all borrowers use the insured submarket.

The resulting responses are displayed in Figure 10. The main result is that this policy is much less effective in the Benchmark economy, leading to less than half the decline in VTI and LTI ratios. The intuition largely follows that of Section 3.3. In the No Uninsured econ-

Figure 10: Tight PTI Experiment

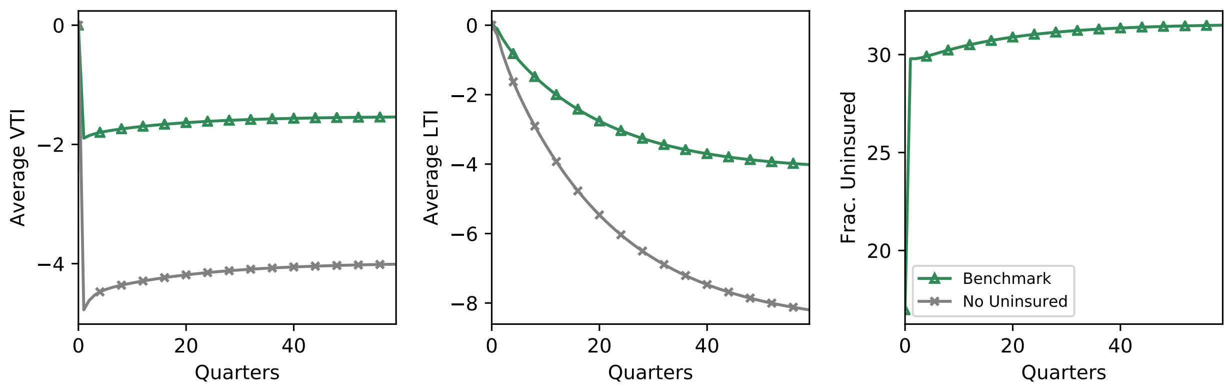

Note: Average VTI is the ratio of total housing value to total income in the economy: $p_t^h \bar{H}/y_t$. Average LTI is the ratio of total mortgage debt to total income in the economy: $(m_{I,t} + m_{U,t})/y_t$. Frac. Uninsured is the share of new mortgage borrowers choosing mortgages from the uninsured market $F_{U,t}$. Black dash-dotted lines indicate the data, while colored lines indicate model experiments. Average VTI and Average LTI are presented as percent deviations from the initial steady state, while Frac. Uninsured is presented as a percentage in levels.

omy, tightening PTI limits cause many borrowers to switch from being LTV-constrained to being PTI-constrained, reducing their housing demand. In the Benchmark economy, however, some of these borrowers switch to the uninsured sector, where their marginal cost is much lower, and incentives to acquire housing are much higher, as seen in the rightmost panel of Figure 10. Translating the intuition of Section 3 into the F^{LTV} terms used in the structural model, a tighter PTI limit reduces F^{LTV} substantially in the single market No Uninsured model, but this is partially undone by a shift in mass from F_I^{PTI} to F_U^{LTV} in the Benchmark model.

As above, due both due to the increase in housing demand, and to credit limits being weakly looser in the Benchmark economy, we also observe a much smaller decline in credit growth in the Benchmark economy relative to the No Uninsured economy.

5.3 Tightening LTV Limits

For the final experiment, we consider the impact of tightening LTV limits, both in the Benchmark economy. For this experiment, we separately tighten the LTV limit in each submarket by 10pp, so that the insured LTV limit falls from 95% to 85%, and the uninsured LTV limit falls from 80% to 70%. Each change is assumed by households to be permanent, so that each response plots the nonlinear transition back to steady state.

Figure 11: Tight LTV Experiment

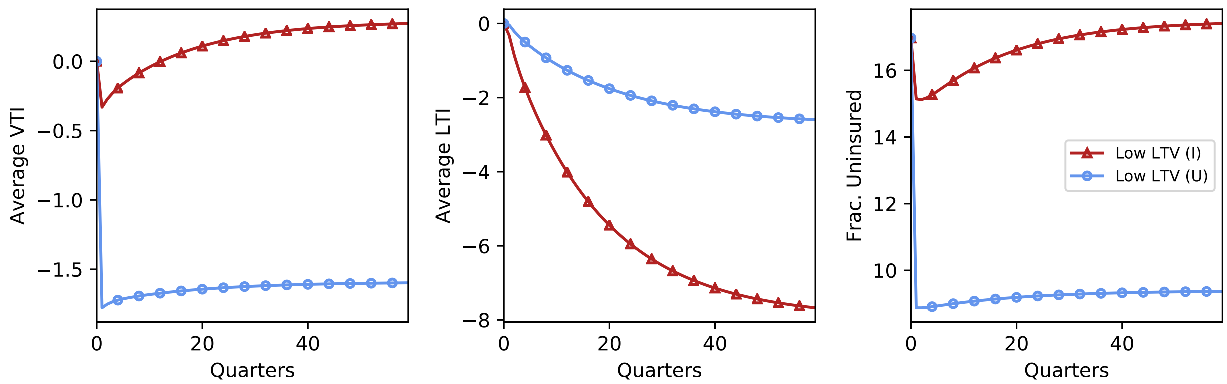

Note: Average VTI is the ratio of total housing value to total income in the economy: $p_t^h \bar{H}/y_t$. Average LTI is the ratio of total mortgage debt to total income in the economy: $(m_{I,t} + m_{U,t})/y_t$. Frac. Uninsured is the share of new mortgage borrowers choosing mortgages from the uninsured market $F_{U,t}$. Black dash-dotted lines indicate the data, while colored lines indicate model experiments. Average VTI and Average LTI are presented as percent deviations from the initial steady state, while Frac. Uninsured is presented as a percentage in levels.

The resulting responses are displayed in Figure 11, and exhibit sharp differences from each other. Beginning with the left panel, tightening the LTV limit in the uninsured sector has a much larger impact on house prices and VTI ratios than the same tightening in the insured sector. This result follows the intuition from Section 3. Since tightening the insured LTV limit shifts the jump in marginal cost to the right, many borrowers increase their housing demand following the shock. Mapping this into the structural model, we note that tightening the insured LTV limit increases F_I^{LTV}, while leaving the overall insured/uninsured split roughly constant, since LTV-constrained borrowers in the insured space cannot gain from switching submarkets. This increase in housing demand and collateral value through F^{LTV} lead to a small *increase* in house prices following the policy change.

In contrast, tightening LTV ratios in the uninsured sector leads to a substantial decline in house prices and VTI ratios. This tightening leads nearly half of the borrowers in the uninsured sector to switch to the insured sector, as shown in the rightmost panel of Figure 11. As shown in Section 3.4, these borrowers choose a much smaller housing allocation following the switch, depressing demand. Mapped into the language of the structural model, since all borrowers in the uninsured space are LTV-constrained, tightening the LTV limit in that space leads more borrowers to switch to the insured space and become PTI-constrained, counterintuitively *decreasing* the share of borrowers who are LTV-constrained. These forces lead to a large decline in housing and collateral demand, and hence house prices.

Turning to the middle panel of Figure 11, we observe the opposite pattern for credit, where tightening insured LTV limits leads to a much larger decrease in credit volumes. This is largely due to the fact that the vast majority of borrowers in the Benchmark economy obtain loans from the insured sector at steady state, and most of these borrowers are LTV-constrained. As a result, the larger direct effect of tightening credit limits for this large portion of the population from tightening the insured LTV limit outweighs the larger indirect effect of reducing house prices from tightening the uninsured LTV limit.

These different responses imply that the appropriate macroprudential tool to use may depend on the policymaker's ultimate goal. To dampen house price growth while leaving aggregate credit balances stable, a tightening of LTV limits in the uninsured sector (sector with tighter LTV, looser PTI) is the better choice. To reduce credit volumes while leaving aggregate house prices largely unchanged, tightening LTV limits in the insured sector (sector with looser LTV, tighter PTI) is the better choice.

6 Conclusion

In this paper, we studied the role of mortgage market segmentation in determining the effectiveness of macroprudential policies. The main finding is that the existence of a multi-market structure, where one submarket has relatively loose LTV limits, and the other has relatively loose PTI limits, has important implications for housing market dynamics and macroprudential policy. The ability to switch between submarkets seriously weakens the ability of stable PTI limits to slow or prevent large increases in house prices in the presence of increased latent demand.

Moreover, the impacts of macroprudential tightenings of the various credit constraints depend crucially on which submarket they are applied to. Importantly, while tightening PTI limits in a single market model is a highly effective policy for dampening growth in house prices and credit during a boom, tightening PTI limits in the sector with loose LTV limits and tight PTI limits only is much less effective in a multi-market setting, as borrowers switching to the alternative sector both evade these tighter credit limits and demand more housing due to their financing incentives. The impact of tightening LTV limits similarly has strong sector dependence. Tightening LTV limits in the insured space where LTV limits are looser fails to depress house prices, as in a single-market setting, but is effective at reducing credit growth. Conversely, tightening LTV limits in the uninsured space is effective at reducing house prices, but with a much smaller impact on credit growth.

Table 1: Parameter Values: Baseline Calibration

Parameter	Name	Value	Internal	Target/Source
Demographics and Preferences				
Fraction of borrowers	χ_b	0.319*	N	1998 Survey of Consumer Finances'
Income dispersion	σ_e	0.411*	N	Fannie Mae Loan Performance Data
Borr. discount factor	β_b	0.981	Y	Value-to-income ratio (Survey)
Saver discount factor	β_s	0.990	N	Avg. 10Y rate
Housing preference	ξ	0.25	N	Davis and Ortalo-Magné (2011)
Borr. labor disutility	η_b	8.796	Y	$n_{b,ss}/\chi_b = 1/3$
Saver labor disutility	η_s	5.838	Y	$n_{s,ss}/\chi_s = 1/3$
Inv. Frisch elasticity	φ	1.0	N	Standard
Housing and Mortgages				
Mortgage amortization (U)	ν_U	0.522%	N	See text
Mortgage amortization (I)	ν_I	0.522%	N	See text
Income tax rate	τ_y	0.204	N	Elenev, Landvoigt, and Van Nieuwerburgh (2016)
Max LTV ratio (I)	$\bar{\theta}_I^{LTV}$	0.95	N	See text
Max LTV ratio (U)	$\bar{\theta}_U^{LTV}$	0.8	N	See text
Max PTI ratio (I)	$\bar{\theta}_I^{PTI}$	0.44	N	See text
Max PTI ratio (U)	$\bar{\theta}_U^{PTI}$	0.7	N	See text
Renewal Rate	ρ	0.05	N	5-Year Contract
PTI offset (taxes, etc.)	α	0.313%	Y	$q^*_{ss} + \alpha$ = 9.97% (annualized)
PTI offset (other debt)	ω	0.05	N	See text
Term premium (mean)	μ_q	0.124%	Y	Avg. mortgage rate
Term premium (pers.)	ϕ_q	0.852*	N	Autocorr. of (mort. rate - 1Y rate), US data
Log housing stock	$\log \bar{H}$	2.178	Y	$p^h_{ss} = 1$
Log saver housing stock	$\log \bar{H}_s$	1.867	Y	Canadian Survey Data
Housing depreciation	δ	0.005	N	Standard
Productive Technology				
Productivity (mean)	μ_a	1.099	Y	$y_{ss} = 1$
Productivity (pers.)	ϕ_a	0.964	N	Garriga, Kydland, and Sustek (2015)
Variety elasticity	λ	6.0	N	Standard
Price stickiness	ζ	0.75	N	Standard
Monetary Policy				
Steady state inflation	π_{ss}	1.005	N	Avg. infl. expectations
Taylor rule (inflation)	ψ_π	1.5	N	Standard
Taylor rule (smoothing)	ϕ_r	0.89	N	Campbell, Pflueger, and Viceira (2014)
Infl. target (pers.)	$\phi_{\bar{\pi}}$	0.994	N	Garriga et al. (2015)

Note: The model is calibrated at quarterly frequency. Parameters denoted "Y" in the "Internal" column are internally calibrated, meaning that they are not set explicitly in closed form, but are instead chosen implicitly to match a particular moment at steady state.

References

BENETTON, M., P. BRACKE, J. F. COCCO, AND N. GARBARINO (2019): "Housing Consumption and Investment: Evidence From Shared Equity Mortgages," .

CAMPBELL, J. Y., C. PFLUEGER, AND L. M. VICEIRA (2014): "Monetary Policy Drivers of Bond and Equity Risks," Tech. rep., National Bureau of Economic Research.

CERUTTI, E., S. CLAESSENS, AND L. LAEVEN (2017a): "The Use and Effectiveness of Macroprudential Policies: New Evidence," *Journal of Financial Stability*, 28, 203–224.

CERUTTI, E., J. DAGHER, AND G. DELL'ARICCIA (2017b): "Housing Finance and Real-Estate Booms: A Cross-Country Perspective," *Journal of Housing Economics*, 38, 1–13.

CROWE, C., G. DELL'ARICCIA, D. IGAN, AND P. RABANAL (2013): "How to Deal With Real Estate Booms: Lessons From Country Experiences," *Journal of Financial Stability*, 9, 300–319.

DAVIS, M. A. AND F. ORTALO-MAGNÉ (2011): "Household Expenditures, Wages, Rents," *Review of Economic Dynamics*, 14, 248–261.

ELENEV, V., T. LANDVOIGT, AND S. VAN NIEUWERBURGH (2016): "Phasing Out the GSEs," *Journal of Monetary Economics, Forthcoming*.

EVERETT, M., J. DE HAAN, D.-J. JANSEN, P. MCQUADE, AND A. SAMARINA (2021): "Mortgage Lending, Monetary Policy, and Prudential Measures in Small Euro-Area Economies: Evidence From Ireland and the Netherlands," *Review of International Economics*, 29, 117–143.

FAVILUKIS, J., S. C. LUDVIGSON, AND S. VAN NIEUWERBURGH (2017): "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium," *Journal of Political Economy*, 125, 140–223.

GARRIGA, C. AND A. HEDLUND (2018): "Housing Finance, Boom-Bust Episodes, and Macroeconomic Fragility," Working Paper, Federal Reserve Bank of St. Louis.

GARRIGA, C., F. E. KYDLAND, AND R. SUSTEK (2015): "Mortgages and Monetary Policy," Working Paper 19744, National Bureau of Economic Research.

GLAESER, E. L. AND J. GYOURKO (2006): "Housing Dynamics," Tech. rep., National Bureau of Economic Research.

GREENWALD, D. L. (2018): "The Mortgage Credit Channel of Macroeconomic Transmission," MIT Sloan Research Paper No. 5184-16.

GRODECKA, A. (2019): "On the Effectiveness of Loan-To-Value Regulation in a Multiconstraint Framework," *Journal of Money, Credit and Banking*.

IGAN, D. AND H. KANG (2011): "Do Loan-To-Value and Debt-To-Income Limits Work? Evidence From Korea," *IMF Working Papers*, 1–34.

INGHOLT, M. M. (2019): "Multiple Credit Constraints and Time-Varying Macroeconomic Dynamics," Tech. rep., Danmarks Nationalbank Working Papers.

JÁCOME, L. I. AND S. MITRA (2015): "LTV and DTI Limits–Going Granular," Tech. rep., IMF Working Paper.

JOHNSON, S. (2018): "Mortgage Leverage and House Prices," Working Paper, Northwestern University.

JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2019): "Credit Supply and the Housing Boom," *Journal of Political Economy*, 127, 1317–1350.

KAPLAN, G., K. MITMAN, AND G. L. VIOLANTE (2019): "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, Forthcoming.

KELLY, R., F. MCCANN, AND C. O'TOOLE (2018): "Credit Conditions, Macroprudential Policy and House Prices," *Journal of Housing Economics*, 41, 153–167.

KEYS, B. J., T. PISKORSKI, A. SERU, AND V. VIG (2013): "4. Mortgage Financing in the Housing Boom and Bust," in *Housing and the financial crisis*, University of Chicago Press, 143–204.

KINGHAN, C., Y. MCCARTHY, AND C. O'TOOLE (2019): "How Do Macroprudential Loan-To-Value Restrictions Impact First Time Home Buyers? A Quasi-Experimental Approach," *Journal of Banking & Finance*, 105678.

KUTTNER, K. N. AND I. SHIM (2016): "Can Non-Interest Rate Policies Stabilize Housing Markets? Evidence From a Panel of 57 Economies," *Journal of Financial Stability*, 26, 31–44.

MORGAN, P. J., P. J. REGIS, AND N. SALIKE (2019): "LTV Policy as a Macroprudential Tool and Its Effects on Residential Mortgage Loans," *Journal of Financial Intermediation*, 37, 89–103.

RICHTER, B., M. SCHULARICK, AND I. SHIM (2019): "The Costs of Macroprudential Policy," *Journal of International Economics*, 118, 263–282.

RUBIO, M. AND J. A. CARRASCO-GALLEGO (2014): "Macroprudential and Monetary Policies: Implications for Financial Stability and Welfare," *Journal of Banking & Finance*, 49, 326–336.

SMETS, F. AND R. WOUTERS (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 586–606.

WONG, T.-C., T. FONG, K.-F. LI, AND H. CHOI (2011): "Loan-To-Value Ratio as a Macroprudential Tool-Hong Kong's Experience and Cross-Country Evidence," *Systemic Risk, Basel III, Financial Stability and Regulation*.

A Appendix

A.1 Optimality Conditions

The borrower's optimality conditions with respect to $m^*_{j,t}$ are:

$$1 = \Omega^m_{b,j,t} + \Omega^x_{b,j,t} r^*_{j,t} + \mu_{j,t}$$

where $\Omega^m_{b,j,t}$ and $\Omega^x_{b,j,t}$ are the marginal continuation costs of an extra dollar of principal balance and promised interest payment on a sector j mortgage, respectively, and $\mu_{j,t}$ is the multiplier on the sector j borrowing constraint. [To be completed.]