



Munich Personal RePEc Archive

Mortgage Lending Limits and Housing Demand: Evidence from Bunching in FHA Borrowing

Heilbron, John

18 November 2021

Online at <https://mpra.ub.uni-muenchen.de/125714/>
MPRA Paper No. 125714, posted 27 Aug 2025 08:28 UTC

Mortgage Lending Limits and Housing Demand: Evidence from Bunching in FHA Borrowing *

John Heilbron [†]

November 18, 2021

Abstract

I adapt the bunching framework to measure the loan-to-value elasticity of housing demand. Unlike existing literature, my estimator can identify the effect of credit supply while remaining agnostic about how households form beliefs over future housing returns. I measure a statistically significant elasticity of demand, suggesting that households are credit constrained at the time of home purchase. The elasticity is economically small, $\sim 14\text{-}25\text{bp}$, suggesting that shocks to credit supply drove housing demand largely through the channel of household beliefs.

Keywords: Household Finance, Governmental Loan Guarantees, Housing Demand

JEL Codes: D12, H81, R21

*I would like to thank Agustin Hurtado, Elisabeth Kempf, Joshua Morris-Levenson, Stefan Nagel, Scott Nelson, Pascal Noel, Kelly Posenau, Amir Sufi, Constantine Yannellis and my peers in Amir Sufi's working group.

[†]University of Chicago, Booth School of Business

1 Introduction

Mortgage credit supply and expectations of future house price growth are prominent among explanations of the housing boom and bust in the 2000s (Mian and Sufi, 2009; Kaplan et al., 2020). Under the credit supply story, households with newly relaxed credit-constraints demanded more housing services and, consequently, more housing. Under the expectations story, household demand for housing services did not change, per se, but demand for housing increased because of a pervasive and counter-factual belief in high returns on housing as an asset class.

In principle, both channels may have contributed to the housing cycle. Empirically determining the relative contribution of each is fraught, however, because the channels are potentially interrelated: changing aggregate credit conditions could affect rational or irrational beliefs and vice versa. In particular, the credit supply story has been most rigorously established by instrumental variable approaches that identify the effect of increased credit supply on local housing price indices in the cross-section (Favara and Imbs, 2015; Johnson, 2019). But this approach captures both a direct effect on demand for housing services as well as a knock-on effect of demand for housing due to changing beliefs.

In this paper, I develop and employ a different tack to better establish the relationship between credit supply and demand for housing by homeowners, the intuition of which I explain here. Households face two primary constraints when borrowing through major institutional channels to finance a home purchase: the loan-to-value limit, \bar{L} , and initial loan balance limit, \bar{B}^o , implied, under fairly weak assumptions, by an institutional debt-to-income limit. For the first \$1 of housing, and for every \$1 a household demands before reaching the initial loan balance limit, \bar{B}^o , the household must pay $$(1 - \bar{L})$ more in down-payment. For the first $1 of housing after reaching \bar{B}^o , and for every $1 thereafter, the household must pay $1 more in down-payment.$

If, outside these institutional lending programs, households are not constrained, and can substitute without cost to another line of credit, e.g. a second mortgage, credit card, or personal loan, then this discontinuity in marginal down-payment requirements should not affect their demand for housing. In fact, by borrowing from other lines of credit, households can eliminate this programmatic discontinuity. For example, having reached the initial loan balance limit, \bar{B}^o , to finance an additional \$1 of housing, households could simply make the \$1 down-payment required by the lender, and borrow \bar{L} on another line of credit, for a net marginal down-payment at the time of purchase of $$(1 - \bar{L})$.$

However, if households are elsewhere constrained, e.g. when home purchase may be financed only with savings and a mortgage loan obtained from a major institutional channel, the lending limits ensure discontinuously higher marginal down-payment and therefore a discontinuously higher utility cost associated with an additional dollar of housing. As a result, households for whom the

marginal benefit of additional housing falls between the utility costs associated with a marginal downpayment of $\$(1 - \bar{L})$ and $\$1$ would demand additional housing until they reached this discontinuity in marginal down-payment requirements. The effect of changing down-payment requirements can therefore be observed in the data as the extent of bunching at the down-payment amount, D^* , associated with this marginal down-payment discontinuity. In a later section, I formalize and clarify the interpretation of this bunching behavior.

This approach to analyzing the link between demand for housing and credit availability has several advantages over those in the literature. Chief among these, it is implausible that identified behavioral responses in this setting are the result of changing expectations. For one, this strategy does not rely on aggregate (e.g. zip code or county-level) treatment, which might affect home-buyer beliefs through its effect on other households or investors. Moreover, even at the individual level, the information treatment and credit constraint quasi-treatment are disentangled. Borrowers may update their beliefs about housing markets as they come to learn the rules that govern their own and others' borrowing limits, but this change in beliefs occurs separately from the selection of a loan from a kinked opportunity set. These features of the empirical design allow for isolation of the effects of credit constraints on housing demand, holding constant households' beliefs about the market.

There are ancillary benefits to this approach as well. For one, the fact that bunching is measured from individual-level outcomes, it is possible to look at different segments of the borrower population to understand the extent to which changing mortgage borrowing constraints influences their demand for housing. This is productive in light of research suggesting that capital gains on housing during the early 2000s varied in the cross-section of housing quality ([Landvoigt et al., 2015](#)). Second, this strategy isolates effects of borrowing requirements on housing demand, a real outcome, rather than first mortgage balances, a financial outcome which may be offset by borrowing along additional margins ([DeFusco and Paciorek, 2017](#)).

This bunching design is not without limitations. Primarily, there is a challenge in comparing the results of this paper, which quote the effect of credit supply on individual quantities of housing demanded, to the existing literature, which quotes the effect of credit supply on equilibrium prices. To make this comparison requires two steps, each of which requires additional assumptions. First, it is necessary to aggregate changes in housing demand resulting from changing credit conditions in the early 2000s. Second it is necessary to map from aggregate change in quantities demanded to equilibrium price changes. Additionally, the analysis requires merging, at the individual level, household incomes, back-end debt-to-income ratios, loan balances, initial interest rates, transaction prices, as well as information on their county of residence and up-front mortgage insurance premia. My ability to do this in the data, at present, extends only to the years 2018-2019 and only for FHA loans, for which subsample both HMDA and McDash provide enough fields to merge the required

variables.

In Section (2), I consider in closer detail how my analysis builds on two robust literatures, housing demand in the 2000s and bunching estimators. In Section (3), I provide context for my analysis by describing the history and features of the FHA lending program; I describe how I use the features of this program to construct the loan-to-value limit and the initial loan balance limit for each borrower. In Section (4), I describe my empirical approach by sketching the household problem being examined (4.1), how observable quantities can be used to estimate the relevant elasticity (4.2), and how to interpret this elasticity (4.3). In Section (5), I describe the way I construct a sample from available data that suits the needs of my estimator (5.1) and technical details of implementing the bunching estimator (5.2). Section (6) presents the results of the analysis and Section (7) concludes.

2 Related Literature

2.1 Beliefs and Credit Constraints

Since the housing crisis of the 2000s, a central line of inquiry has sought to understand the determinants of housing demand and the housing cycle. Early research substantiated the role of credit supply in generating both the boom and bust in housing prices. By disaggregating to the level of the zip code, [Mian and Sufi \(2009\)](#) note that areas with high mortgage credit growth in the boom years tended to have relatively low income growth, suggesting that credit supply rather than demand drove the rise in household leverage and house prices. The credit supply story has theoretical as well as empirical appeal, as it plays an important role in building a cohesive narrative of the major facts surrounding the housing cycle in the early 2000s. [Justiniano et al. \(2019\)](#) observes that for increased leverage to be consistent with falling interest rates, as was true in the early 2000s, credit supply must have been expanding rather than credit demand.

Other scholarship has pursued the role of household expectations; under this hypothesis, home prices increased because optimistic beliefs about the path of future home prices in the future generated demand for housing access in the present. [Kaplan et al. \(2020\)](#) argues that, though credit supply helps to explain changes in homeownership (see also, [Acolin et al. \(2016\)](#)) and default rates during the crisis, it provides little explanatory power for house prices. Constrained households may obtain equivalent housing services from the rental market, they argue, such that relaxed credit constraints do not generate increased demand for housing units. Rather, they attribute the boom to a counter-factual belief that housing would appreciate in value. [Mian and Sufi \(2018\)](#) studies home purchases by investors and notes that investors' optimistic beliefs drive their purchasing behavior and therefore home prices during the boom.

A central challenge in empirical scholarship on this topic is disentangling beliefs from credit constraints, which are easily confounded: changing aggregate credit conditions could affect rational or irrational beliefs and vice versa. The credit supply story has been most rigorously established by instrumental variable approaches that identify the effect of increased credit supply on local housing price indices in the cross-section (Favara and Imbs, 2015; Johnson, 2019). But this approach captures both a direct effect on demand for housing services as well as a knock-on effect of demand for housing due to changing beliefs.

Some empirical work has attempted to disentangle credit access from beliefs in generating housing demand. Empirical tests to rule out the beliefs story have appealed to regions with high housing supply elasticity in which a belief in long-run price fluctuations might be unreasonable (Mian and Sufi, 2009). But because these exercises have not been able to directly measure beliefs, it is more difficult to control for unreasonable beliefs. Bailey et al. (2019) use social networks to identify shocks to beliefs about housing values. They argue that changes in beliefs do not generate enough change in housing demand to account for the housing cycle of the early 2000s. They suggest that credit access may play a role instead but their research design does not enable them to measure this in the data.

Finally, some work has examined bunching in data at loan-limits, an approach which, though not the studies' explicit intent, is more persuasive in controlling for beliefs. Anenberg et al. (2017) looks for a mass point at the "loan frontier", an empirically imputed maximum loan balance available to a borrower of a given credit-score, down-payment, and income. This mass point is a poor test of credit constraints, though, as down-payments are chosen contemporaneously with home purchase. DeFusco and Paciorek (2017) identifies bunching at the jumbo loan threshold where interest rates increase, suggesting that first mortgage demand is responsive to interest rates. This measures a different elasticity altogether, though, and the link from mortgage demand to housing demand is unclear because of the possibility of putting other lines of credit to use.

2.2 Bunching Estimators

There is widespread use of bunching estimators in the economics literature to identify elasticities in the data. An older literature considers information content available in kinked opportunity sets, but the availability of administrative data set off the present use of bunching estimators, first formalized by Saez (2010). Kleven and Waseem (2013) provides a theoretical analysis of notches and relates the empirically estimated elasticities to structural elasticities of interest. These estimators have been used in a wide variety of "real" contexts (Yelowitz, 1995; Sallee and Slemrod, 2012; Ramnath, 2013; Manoli and Weber, 2016; Blundell and Shephard, 2012; Blundell and Hoynes, 2004; Kleven, 2016).

The introduction of bunching estimators to the finance literature is somewhat more recent. [DeFusco and Paciorek \(2017\)](#) estimates first mortgage borrowing elasticity relative to the interest rate by exploiting the interest rate variation at the GSE conforming loan limit. [Best et al. \(2018\)](#) estimates the elasticity of inter-temporal substitution using interest rate variation at particular LTV ratios in the British mortgage market. [Dagostino \(2019\)](#) estimates the elasticity of municipal bond issuance size using changes in available yields due to bank qualification status. [Bachas et al. \(2020\)](#) estimate the elasticity of loan size using a discrete change in the loan guarantee rate. [Ebrahimian \(2020\)](#) uses bunching as a moment to identify a structural IO model of student debt.

Relative to this literature, I exploit policy parameters that generate a kink from borrowing rather than budget constraints. A households' additional costs enter through increasing down-payment requirements at the kink point rather than increasing interest expense. My analysis explains how it is possible to use a bunching estimator to estimate and interpret a financial elasticity of interest, the loan-to-value elasticity of housing demand.

2.3 FHA Lending

There is a small literature on FHA lending, particularly in the aftermath of the financial crisis, during which it served as an important source of mortgage lending as the private market contracted. [Hwang et al. \(2016\)](#) study the effects of HERA, which raised FHA loan limits and induced more borrowing, but primarily through cash-out refinancing and higher LTV ratios at purchase rather than the purchase of larger home. [Park \(2016\)](#) studies the same limit increases and finds an increase in FHA-qualified loan originations in 2008 but no corresponding decrease in 2014 when the limit increases expire; he concludes that borrowers substitute toward the private market as it recovers in the interim years. [Passmore and Sherlund \(2016\)](#) find real effects of FHA lending during this period; in particular, they document that counties with ex-ante higher FHA lending experience smaller declines in mortgage purchase originations, house prices, and new automobile purchases as well as smaller increases in unemployment rates during the crisis. [DeFusco and Mondragon \(2020\)](#) document a crisis-era FHA policy excluding unemployed borrowers from refinancing severely reduced the refinancing activity of those with greatest demand to do so, curtailing the effects of monetary policy. I document several new facts relative to this literature, namely, the bunching behavior of borrowers at the county loan limit, and the responsiveness of borrowers in the upper quantiles of the borrowing distribution to the cross-sectional increase in the limits.

3 The Federal Housing Administration (FHA)

3.1 The History and Function of the FHA

The Federal Housing Administration, or FHA, was founded in 1934 to help stabilize the housing market during the Great Depression. By providing insurance for mortgage principle in the event of borrower default, the government provided assurances to help keep lenders operating in the market (DeFusco and Mondragon, 2020). The program was designed to be self-funding, with insurance premiums sufficient to cover program costs due to mortgage default.

Since its creation, the prominence of FHA lending in the mortgage market has waxed and waned. In the aftermath of World War II, the Veteran's Administration entered mortgage lending and the FHA lost market share. Amid civil rights legislation during the 1960s, the FHA was brought under the regulation of the Department of Housing and Urban Development, or HUD, and authorized to expand in scope by congress. In the late 20th century, the FHA lost market share as the growth of securitization and private mortgage insurance made mortgage credit more easily available to sub-prime borrowers in the private market. During the housing crisis of the late 2000s, FHA market share grew substantially from 4.5% to about 25% of mortgage originations. Since then, it has dwindled again, but remains around 15-20%. (Immergluck, 2011)

The FHA program has remained self-sufficient, with a single exception in the aftermath of the housing crisis (Puzzanghera, 2013). Insurance premiums through the program are two-fold, and consist of an up-front mortgage insurance premium (UFMIP) and an annual premium. The size of these premiums may be updated by regulators and is disclosed in the FHA's Mortgagee Letters. This historical values of these premium rates are depicted in Figure (11); at present, the UFMIP rate is 1.75% and the annual MIP rate is 0.85%.

The Government National Mortgage Association, or GNMA, was founded in 1968. Since its founding, it has served to guarantee timely payment of principal and interest on MBS backed by pools of government-insured mortgage loans, most frequently loans with FHA insurance. Although it does not buy mortgages or issue MBS itself, this guarantee provides FHA loans with access to the secondary market. (FDIC, 2018)

At present, loans that qualify for FHA insurance tend to serve low-income borrowers and first-time home buyers who cannot afford down-payments for conventional mortgages and do not qualify for private mortgage insurance. Table (1) summarizes a 1pp sample from HMDA including all first mortgages on owner-occupied homes in 1-4 family units originated during 2018 and 2019. FHA-insured loans tend to be smaller and more standardized in repayment structure, and to have higher LTV and DTI ratios. Borrowers obtaining these loans tend to be younger, have lower income, and be hispanic or Black. The purchased homes are in poorer census tracts with higher minority populations and lower home-ownership rates.

3.2 FHA Loan Limits and the Bunching Framework

The FHA restricts access to mortgage credit in a variety of different ways including characteristics of the borrower, such as their FICO score, characteristics of the purchased housing, which must meet basic sanitary and safety requirements, and the characteristics of the loan itself. The main identification strategy pursued in this paper relies on the presence of two limits, a loan-to-value limit, \bar{L} , and an initial balance limit, \bar{B}^o . In this section, I describe the borrowing limits imposed on FHA loans and how these limits correspond to the limits central to the identification strategy.

To qualify for insurance through the FHA, loans must conform to three limits: a county-level cap on the initial loan balance, \bar{B}^o_{jt} , a cap on the loan-to-value ratio of the loan, \bar{L}_t , and a cap on the resulting debt-to-income ratio of the borrower following the loan, \bar{D}_{ijt} . Consider borrower i in county j and year t with annual (monthly) income y_{ijt} (\hat{y}_{ijt}) and non-mortgage debt service expenses κ_{ijt} ($\hat{\kappa}_{ijt}$) who purchases a home at price P_{ijt} with a 30-year fixed-rate mortgage characterized by initial balance B^o_{ijt} and annual interest rate r_{ijt} (and implied monthly amortizing factor $\hat{r}_{ijt} \equiv \frac{r_{ijt}/12}{1-(1+r_{ijt}/12)^{-30*12}}$). The FHA restrictions may be represented as follows:

$$B^o_{ijt} \leq \bar{B}^o_{jt} \quad (1)$$

$$\frac{B^o_{ijt}}{P_{ijt}} \leq \bar{L}_t \quad (2)$$

$$\frac{\hat{r}_{ijt}B^o_{ijt} + \hat{\kappa}_{ijt}}{\hat{y}_{ijt}} \leq \bar{D}_{ijt} \quad (3)$$

Assume that the rate of the loan is pinned down by borrower rather than transaction characteristics (i.e. that it does not depend on the loan size). We can then solve Equation (3) for B^o_{ijt} and combine this with Equation (1) to obtain an individual-level loan balance limit:

$$B^o_{ijt} \leq \bar{B}^o_{ijt} \equiv \min \left\{ \bar{B}^o_{jt}, \frac{1}{\hat{r}_{ijt}} \left[\bar{D}_{ijt} \hat{y}_{ijt} - \hat{\kappa}_{ijt} \right] \right\} \quad (4)$$

We can rewrite Equations (2) and (4) in terms of prices and down-payments so that they correspond more closely with the bunching framework outlined in Section (4.1). For this change of variables, we use the fact that the down-payment and initial balance together make the transaction price, $P_{ijt} = D_{ijt} + B^o_{ijt}$.

$$P_{ijt} \stackrel{(2)}{\leq} P_{ijt} \frac{1 - B^o_{ijt}/P_{ijt}}{1 - \bar{L}_{ijt}} = \frac{D_{ijt}}{1 - \bar{L}_{ijt}} \quad (5)$$

$$P_{ijt} = D_{ijt} + B^o_{ijt} \stackrel{(4)}{\leq} D_{ijt} + \bar{B}^o_{ijt} \quad (6)$$

Finally, we can combine Equations (5) and (6), differentiate with respect to the price of the home purchased, and write the piece-wise function for the marginal down-payment relative to a dollar increase in housing demand. Note that we define a reference house price, $P_{ijt}^* = \frac{\bar{B}_{ijt}^o}{\bar{L}_{ijt}}$, and reference down-payment, $D_{ijt}^* = (1 - \bar{L}_{ijt})P_{ijt}^*$, as the price and down-payment at which (5) and (6) are both binding.

$$\frac{dD_{ijt}}{dP_{ijt}} \geq \begin{cases} 1 - \bar{L}_{ijt} & P_{ijt} < P_{ijt}^* \\ 1 & P_{ijt} \geq P_{ijt}^* \end{cases} \quad (7)$$

3.3 FHA Loan Limit Assignment Rules

Identifying the choice set in the space of prices and down-payments available to an individual borrower in the data requires imputing the loan-to-value and borrower-level initial balance limits assigned to them in the FHA program. The loan-to-value limit was 0.965 for all borrowers over the entire sample period. Borrower-level initial balance limits are more sophisticated. Per Equation (4), I assign these as $\bar{B}_{ijt}^o \equiv \min \left\{ \bar{B}_{jt}^o, \frac{1}{\hat{r}_{ijt}} [\bar{D}_{ijt} y_{ijt} - \kappa_{ijt}] \right\}$. Below, I review the method for assigning county-level loan limits, \bar{B}_{jt}^o , and debt-to-income limits, \bar{D}_{ijt} , in the FHA and in my data. I review how I identify household debt-service, κ_{ijt} , which is not explicitly available as a field in the merged data.

The county-level loan limit, \bar{B}_{jt}^o , is computed as the product of HUD's county-median home price index the previous year, $\tilde{P}_{j,t-1}$, and a median home price multiplier, θ_t , unless this value is below the FHA national loan limit floor, $\underline{\bar{B}}_t^o$, or above the FHA national loan limit ceiling, $\overline{\bar{B}}_t^o$, in which case it takes on those values, respectively. The FHA national loan limit floor and ceiling are set according to the conventional conforming national loan limit floor, $\bar{B}_{G,t}^o$, as well as multipliers for the floor, $\underline{\varphi}_t$, and ceiling, $\overline{\varphi}_t$. (In particular, $\underline{\bar{B}}_t^o = \underline{\varphi}_t \bar{B}_{G,t}^o$ and $\overline{\bar{B}}_t^o = \overline{\varphi}_t \bar{B}_{G,t}^o$.) Formally, this may be represented as follows:

$$\bar{B}_{jt}^o = \begin{cases} \underline{\bar{B}}_t^o & \theta_t \tilde{P}_{j,t-1} \leq \underline{\bar{B}}_t^o \\ \theta_t \tilde{P}_{j,t-1} & \theta_t \tilde{P}_{j,t-1} \in [\underline{\bar{B}}_t^o, \overline{\bar{B}}_t^o] \\ \overline{\bar{B}}_t^o & \theta_t \tilde{P}_{j,t-1} \geq \overline{\bar{B}}_t^o \end{cases} \quad (8)$$

Figure (??) depicts the historical values of the median price multiplier, θ_t , floor and ceiling multipliers, $\underline{\varphi}_t$ and $\overline{\varphi}_t$, as well as the conventional conforming national loan limit floor, $\bar{B}_t^{o,GSE}$, and FHA national loan limit floor and ceiling, $\underline{\bar{B}}_t^o$ and $\overline{\bar{B}}_t^o$. Figure (12) depicts loan limits assigned to each county by the FHA against the median home price used to assign the limit. A separate panel is used for every period in which HUD enforced different limits. It is clear that the assigned limits

largely follow the truncated linear assignment rule. In the immediate aftermath of the financial crisis, 2008-2013, additional provisions preventing a fall in limits are responsible for the observable noise in the limits from the stated rules. I can merge these limits into the data using the county and origination month or year of the transaction.

The debt-to-income limit imposed by the FHA varies according to whether the borrower is approved through the Fannie Mae Automated Underwriting System (AUS) or requires approval through the Manual Underwriting System (MUS). Typically, lenders will attempt to obtain approval through the AUS, which has a maximum debt-to-income limit of 0.569. If borrowers do not qualify under the AUS, then the MUS determines a borrower's debt-to-income limit according to borrower FICO score and a number of "compensating factors", which are indications that the borrower may have sufficient assets or income to afford the loan. In particular, for borrowers with a FICO score between 500 and 580, the MUS sets a debt-to-income limit of 0.43. For borrowers with a FICO score between 580 and 850, the MUS sets a debt-to-income limit of 0.43 for no compensating factors, 0.47 for one, and 0.50 for two compensating factors.

In practice, it is difficult to identify whether a loan was underwritten by the MUS or AUS. It is also difficult to identify the number of compensating factors enjoyed by a borrower in the absence of the administrative data. Although I have the FICO score of borrowers in the data, there are no borrowers in the merged sample with a FICO between 500 and 580, so this is of no value in discriminating debt-to-income limits. Instead, I assign borrowers the lowest debt-to-income limit that is still above that which they received in the data.

Although non-mortgage debt service, κ_{ijt} , is not a field in the merged data, it can be obtained by subtracting the monthly mortgage expense on the observed loan from the total debt service implied by the the debt-to-income ratio and income of the borrower, $\kappa_{ijt} = y_{ijt}D_{ijt} - \hat{r}_{ijt}B_{ijt}^o$.

4 Empirical Strategy

4.1 The Household Problem

A household has preferences over housing, H_t , and non-housing consumption, X_t , in each of two periods, t_0 and t_1 . To simplify the derivations shown in Appendix (C), I assume standard concavity, that housing and non-housing consumption are separable within period, and that the borrower discounts with factor β . For exogenous reasons, the borrower opts to purchase a home at t_0 , and sell the home in favor of renting in t_1 . The borrower is endowed with an income made available in each period, y_0 and y_1 , and chooses how to allocate this wealth as consumption. Of particular interest is the intensive margin of household demand for housing, H_0 .

I normalize the initial price of non-housing consumption to 1; it rises according to the risk-free

rate to $(1 + r_f)$ at t_1 . The price of housing is initially p_H and evolves as $p_H(1 + \tilde{r}_H)$. It is the borrower's beliefs that matter when choosing how much housing to purchase and, for simplicity in the derivations, I assume these are dogmatic and degenerate, $\mathbb{E}_i^0[\tilde{r}_H] = v_i$.

Households have two financial vehicles to facilitate their consumption decision. They may save or borrow at the risk-free rate along what I call a “liquid” margin, analogous to a checking account or credit card. Along this margin, they face a standard borrowing constraint, which I normalize to 0 for simplicity. They may also take out a mortgage loan to finance the purchase of their housing, also at the risk-free rate, r_f . As is standard in mortgage lending, this loan is subject to two constraints, a loan-to-value limit, \bar{L} , and a debt-to-income limit, the latter of which I rewrite as an initial mortgage balance limit, \bar{B}^0 .

To choose the optimal quantity of housing at t_0 , the household considers the marginal costs and benefits, in utility terms, of an additional unit of housing. I depict this choice in Figure (??). The marginal benefits to an additional unit of housing are standard. They are the sum total of utility from t_0 housing services, U_H^0 , and the discounted utility from additional t_1 non-housing consumption the borrower anticipates purchasing out of excess returns on housing relative to the risk-free asset, $\beta U_X^1 p_H \mathbb{E}_i^0[\tilde{r}_H - r_f]$. In general these are falling due to concavity in both X and H .

The household's marginal costs feature a potential non-standard discontinuity due to the two constraints on the mortgage loan. A household exhausting their loan-to-value limit can rely on additional mortgage financing until they also exhaust their initial balance limit. This occurs at the reference quantity of housing, $H^* = \frac{1}{p_H} \frac{\bar{B}^0}{\bar{L}}$. Below this quantity, the utility cost of an additional unit of housing derive from lost t_0 consumption due to additional down-payment and lost t_1 consumption due to additional costs of eventual mortgage repayment, $p_H[(1 - \bar{L})U_X^0 + \bar{L}U_X^1]$. Beyond this quantity, the price of an additional unit of housing is financed entirely from a down-payment out of t_0 consumption, and the utility costs are $p_H U_X^0$.

The household's additional margin for borrowing and saving outside of the mortgage contract plays a key role in generating the discontinuity in marginal costs. If households are not at their “liquid” borrowing constraint, then they may smooth consumption by equating marginal utility of non-housing consumption in both periods. In this case, there will be no discontinuity in marginal costs, $p_H[(1 - \bar{L})U_X^0 + \bar{L}U_X^1] \stackrel{U_X^0 = U_X^1}{=} p_H U_X^0$. This is akin to the Modigliani-Miller benchmark. However, if borrowers are constrained along other margins, then the household would benefit from more consumption smoothing and so $U_X^0 > U_X^1$, generating the upward discontinuity in marginal costs.

In the absence of constrained households, therefore, there will be a smooth distribution of households across the reference quantity of housing, H^* . In Figure (9), in the counterfactual case, borrowers b and m choose different quantities of housing. Denote this counterfactual distribution as $f(H)$. By contrast, when households are constrained, they tend to bunch at the reference quantity.

In Figure (9), buncher b and marginal household m both choose the reference quantity of housing, H^* .

4.2 Identification

To identify the loan-to-value elasticity of housing demand, I rely on as-good-as-random assignment of counterfactual household demand for housing across the reference housing quantity. At and above the reference quantity, the decreased effective loan-to-value limit induce an increase in utility costs of additional housing for borrowing constrained households. For this reason, households with counterfactual demand above the reference quantity may be considered the “treated” households and they respond by adjusting their demand downward. Households with counterfactual demand below the reference quantity, the “control” group, do not respond. The size of the resulting bunching relative to the unperturbed distribution measures the behavioral response of the “treated” households.

I measure the loan-to-value limit elasticity of housing demand, $\widehat{\epsilon}_L^H = \frac{\% \Delta H}{\% \Delta L}$. I consider the behavioral response of the marginal buncher to the reduction in the loan-to-value limit on housing units above the reference quantity from \bar{L} to 0. I also observe that house prices are simply house quantities scaled by a constant price level. The estimator can therefore be rewritten as:

$$\frac{\% \Delta H}{\% \Delta L} = \frac{\frac{H^* - H_m^o}{H_m^o}}{\frac{0 - \bar{L}}{\bar{L}}} = \frac{[H_m^o - H^*]}{H^* + [H_m^o - H^*]} = \frac{[P_m^o - P^*]}{P^* + [P_m^o - P^*]}$$

The reference house price, P^* , can be obtained from the household’s borrowing limits. The adjustment of the marginal borrower, $[P_m^o - P^*]$, can be obtained from the bunching equation:

$$B = \int_{P^*}^{P_m^o} f(P) dP \approx f(P^*) [P_m^o - P^*]$$

The bunching mass, B , can be measured in the realized distribution. And the counterfactual density at the reference house price $f(P^*)$, can be approximated from the realized density as well.

Below, I discuss the appropriate interpretation of the estimate. I distinguish between loan-to-value limits on average and marginal units of housing, establish the nature of the measured elasticity relative to these, and describe the relationship of the measured elasticity relative to loan-to-value and debt-to-income policy parameters. I then distinguish between partial and general equilibrium loan-to-value limit elasticities and characterize one advantage of estimating partial equilibrium elasticities.

4.3 Interpretation

4.3.1 Marginal vs. Average LTV Limit Elasticity

Because loan-to-value limits determine the trade-off between down-payment supply and housing demand, they function as prices in the household problem.¹ In this capacity, their impact resembles compensated and uncompensated demand elasticities from the standard consumer theory. Households may reduce housing demand in response to a lower loan-to-value limit on the entire loan (i.e. the average unit of housing purchased), akin to an uncompensated price elasticity. Households may also reduce housing demand in response to a lower loan-to-value limit on merely the last unit of housing, akin to a compensated price elasticity. These behavioral responses are depicted in Figure (2).

Compensated and uncompensated price elasticities are related in the famous Slutsky equation by income effects. In this context, the marginal, $\epsilon_L^{H,mg.}$, and average, $\epsilon_L^{H,avg.}$, loan-to-value limit elasticities are related analogously by a quasi-income effect, where the quasi-income is the loan balance obtainable before making any down-payment. Following the sketch in Kleven (2016), I use this Slutsky-type relation to show in Appendix (C.3) that the estimated elasticity is a weighted average of these elasticities:

$$\widehat{\epsilon}_L^H = \left[1 - \frac{\Delta a}{\Delta \mu}\right] \epsilon_L^{H,mg.} + \left[\frac{\Delta a}{\Delta \mu}\right] \epsilon_L^{H,avg.} \quad (9)$$

The estimator contains some information about average loan-to-value limit elasticity, but there is clearly a margin for bias. Concerns about bias may be mitigated for two reasons. First, the value of $\frac{\Delta a}{\Delta \mu}$, which ranges from 0 to 1, increases in the sharpness of the kink and in the present setting the kink is quite sharp. Second, to the extent that quasi-income effects are null, we have that the marginal and average loan-to-value limit elasticities are identical and the estimated elasticity captures both. In the bunching literature on tax elasticities of earnings, structural models tend to impose this.

4.3.2 Policy LTV and DTI Limit Elasticities

Loan-to-value and debt-to-income limits are policy parameters that can be used to regulate the mortgage and housing markets. Below, I review the nature of these policy elasticities and the ways in which they relate to and differ from the conceptual elasticities introduced above. Though the estimator in this setting captures some information about behavioral responses to changes in

¹Between the price level of housing, p_H , and the loan-to-value limit, \bar{L} , this suggests that households face two prices associated with housing. This comports with the intuition that households make credit constrained down-payments at the time of purchase and unconstrained mortgage payments at later dates. In other words, there are two separate margins along which households sacrifice consumption.

policy loan-to-value limits, it does a better job of estimating responses to changes in policy debt-to-income limits. A related research design exploiting features of the mortgage market may better capture responses to policy loan-to-value limits.

The effect of adjusting policy loan-to-value limits bears closest resemblance to the average loan-to-value limit elasticity because it affects down-payment requirements on all units of housing up to the reference quantity. To the extent that debt-to-income limits restrict household behavioral response to a change in the policy limit, the effect of policy may understate the average loan-to-value limit elasticity. Setting aside this caveat, the decomposition of the estimated elasticity above indicates that it captures some information about effects of changing loan-to-value policy limits, with associated caveats.

In spite of its shortcomings, this research design suggests another that might better capture the policy effects of changing loan-to-value limits. Conventional mortgages and FHA mortgages have different loan to value limits and, as described, FHA mortgages have limits on initial balances. This variation is depicted in Figure (3) and serves as a close analogy to the bunching notch designs in the literature, which do a better job measuring average rather than marginal effects.

The effect of adjusting policy debt-to-income limits is more similar to the marginal loan-to-value limit because they affect effective loan-to-value limits only on units of housing beyond the reference quantity and not on initial units. Again, if the adjusted debt-to-income limit remains binding, the policy effect may understate the elasticity. Additionally, because non-infinitesimal changes in debt-to-income limits imply quasi-income effects, the policy effect may overstate the marginal loan-to-value limit elasticity. These quasi-income effects are precisely the reason for the non-zero weight placed on the average loan-to-value limit elasticity in Equation (9). For this reason, the policy effects of changing debt-to-income limits are those best captured by this framework. This notion is depicted in Figure (4).

4.3.3 Partial vs. General Equilibrium LTV Limit Elasticity

In partial equilibrium, an increase in the availability of mortgage credit may increase housing demand by relaxing household credit constraints, all else equal. In general equilibrium, aggregate increases in the availability of mortgage credit may affect housing demand through effects in other markets. For example, as household wealth increases due to wages or non-housing asset values, wealth effects may further drive demand. The strategy in this paper identifies the partial equilibrium effect of loan-to-value limits on housing demand.

An advantage of measuring the partial equilibrium elasticity is that it can better isolate the effects of easy credit on household demand for housing without capturing potentially confounding effects, such as beliefs about housing returns. Indeed, the relative roles of changing beliefs and easy credit remains an enduring question about the housing cycle of the early 2000s. A predom-

inant empirical approach in the literature has been to exploit heterogeneity in the cross-section of geographies and to find instruments at the geographic level for the availability of credit. A shortcoming of this approach is that it can only identify the pure effects of relaxed credit constraints under fairly strong assumptions about household beliefs.

To see this, consider a simple model of housing demand in which demand is a function of county-level loan-to-value limits and individual beliefs about housing:

$$H_{ic} \sim \beta_0 + \beta_L \bar{L}_c + \beta_r \mathbb{E}_i[\tilde{r}_{H,c}] + \varepsilon_{ic} \quad (10)$$

Further specify a fairly general model of household beliefs as a linear combination of everything in the information set of the household:

$$\mathbb{E}_i[\tilde{r}_{H,c}] \sim \rho_0 + \rho_L \bar{L}_c + \rho'_F \mathcal{F}_i \setminus \bar{L}_c + \eta_{ic} \quad (11)$$

The county-level loan-to-value limit is part of the household's information set and is represented because of its possible dual role in mitigating credit constraints and shaping beliefs.

The parameter of interest is β_L , which, as described above, is the parameter for which this paper provides an unbiased estimate. Ordinarily, using a geographic instrument to resolve the problem of an omitted variable, beliefs, would also provide an unbiased estimate. However, because any county-level treatment that eases a household's credit constraints doubles as an information treatment that may shape their beliefs about housing returns in the area, this strategy may produce a biased estimate of the pure credit availability effects.

To see this, construct a candidate instrument, z_c , that satisfies relevance:

$$\bar{L}_c \sim \lambda_0 + \lambda_z z_c + u_c \quad \lambda_z \neq 0$$

Further stipulate that the candidate is nearly exogenous in the sense that it is orthogonal to all other shocks and information:

$$z_c \perp \mathcal{F}_i \setminus \bar{L}_c, \eta_{ic}, \varepsilon_{ic}$$

The candidate instrument cannot be orthogonal to beliefs because it is non-orthogonal to loan-to-value limits by construction and, as information, these enter beliefs. For this reason, the instrumental variable estimator may include bias:

$$\hat{\beta}_L^{IV} = \frac{\frac{Cov(H_{ic}, z_c)}{Var(z_c)}}{\frac{Cov(\bar{L}_c, z_c)}{Var(z_c)}} = \dots = \beta_L + \beta_r \rho_L$$

In practice, there may be no bias if either return expectations have no effect on housing demand,

$\beta_r = 0$, or local credit availability has no effect on return expectations, $\rho_L = 0$. Because a household with more optimistic return expectations can increase their perceived wealth by weighting their portfolio toward housing, it is unlikely that the former holds. If the local housing market has perfectly elastic housing supply, it is long-run expectations that enter the demand equation, and households have rational expectations, then the latter may hold.

Rational expectations is a useful benchmark but a fairly strong assumption. In the present setting, it is possible to weaken this assumption because the effective treatment of loan-to-value limits takes place at the individual rather than aggregate level. The availability of credit experienced by the borrower in no way differentially informs them about the availability of credit experienced by other borrowers and therefore it is less plausible that they might draw conclusions about the future price of housing in the area.

5 Estimation

5.1 Sample Construction

A key obstacle for implementing the estimator outlined in the previous section is the ability to reconstruct the borrower’s choice set from their characteristics so that their transaction may be located in that choice set relative to the kink. Loan-to-value limits are fairly standard across borrowers, but to impute effective initial balance limits requires knowledge of both a borrower’s debt-to-income limit and their income at the time of origination. I overcome data limitations by merging HMDA and CRISM data. In doing so, I follow the observation from [Bartlett et al. \(2019\)](#) that these data are mergeable; I have a somewhat different purpose in doing so, and am somewhat more limited in my ability to do so extensively.

HMDA data cover 90% of mortgage originations in the United States ([Bartlett et al., 2019](#)), provide borrower-level data, and include minimally-redacted information on borrower income. However, HMDA does not cover debt-to-income limits or loan realizations until 2018 and these are heavily redacted. CRISM is a merge of Equifax credit bureau fields to BlackKnight mortgage loan origination and performance data; since 2005, it covers roughly 60% of mortgage originations ([Adelino et al., 2013](#)).

By merging the datasets, I obtain both income and realized debt-to-income values on a sample of FHA loans. My sample is limited to FHA loans in the years 2018 and 2019 for reasons of data availability. CRISM tends to report 5-digit, as opposed to 3-digit, zip codes only for FHA loans; HMDA contains enough origination variables, like home price and interest rate, only beginning in 2018. In principle, with better data, it would be possible to compute the analysis on any loan with both a loan-to-value and debt-to-income limit. I describe the merge procedure in more detail

below.

I begin by constructing a large sample of FHA loans, which I refer to as the CRISM 2009-2020 Sample. This consists of loans in CRISM that are vanilla (non-IO and non-balloon), fixed-rate, 15 or 30 year, first mortgages for purchase of single-family, owner-occupied housing and are originated through the FHA between 2009 and 2020. I discard a handful of observations which are not onboarded to CRISM promptly, or have unrealistically high LTV ratios at origination. I use HUD crosswalks to identify the counties associated with the 5-digit zip codes in these samples, and use county and month of origination to merge in up-front mortgage insurance premium rates, the HUD county median home price index, county-level FHA loan limits and county-level GSE conforming loan limits, all hand-collected from mortgagee letters on HUD's website.²

Next, I construct another sample of FHA loans from HMDA, which I refer to as the HMDA 2018-2019 Sample. I begin by identifying loans in HMDA that are vanilla (non-IO, non-balloon, non-HELOC, non-reverse, and non-negatively amortizing), fixed-rate, 15 or 30 year, first mortgages for purchase of single-family, owner-occupied housing and are originated through the FHA between 2018 and 2019.

Finally, I assemble a sample for the main analysis in the paper, which I refer to as MSAMP3. For each loan in the CRISM 2009-2020 Sample, I search for a loan in the HMDA 2018-2019 Sample according to the following procedure. First, I consider any loan in this sample that with the same state, county, zip code, origination year, and loan term (15 or 30 years). Next, score the proximity of the loans along four characteristics, the initial balance, the transaction price, the loan-to-value ratio, and the interest rate. Because of censoring in both HMDA and CRISM, I do not penalize differences in the initial balance or transaction price up to \$5k. Outside of this bandwidth, and in the case of the loan-to-value and interest rate, I apply a linear penalty for deviations, until at a large-enough bandwidth, I say there is no match between the characteristics. I require the candidate HMDA loan to match the CRISM loan along all four characteristics and then keep a single match from any that remain that performs the best on a combination of all four dimensions.

I depict the quality of the merge in Figure (15). The vast majority of these merges are exact matches in interest rates and loan-to-value ratios. Further evidence of the match quality can be obtained by observing the sharp decrease in matches for which the CRISM loan amount or house price is more than \$5k away from its HMDA counterpart. This is consistent with an exact match given the redaction methods in the HMDA data.

I describe the characteristics of the merged sample in Table (2). I am able to identify a match for 75% of the loans in CRISM. Because the CRISM sample is so small, however, this is only about

²Note that it is possible, in principle, to extend this sample to 2005; the ease of collecting FHA and GSE conforming loan limits beginning in 2009 was the motivation for the sample start date. Furthermore, obtaining 5-digit zip codes in CRISM, which is necessary both for identifying county-level limits and for a successful merge to HMDA data, requires focusing on FHA loans rather than conventional mortgages in the version of CRISM to which I have access.

15% of the near-universal HMDA sample. Still, the contents of the CRISM sample look broadly similar to the more representative HMDA sample on the attributes available in CRISM (columns 1 and 2). And the merged sample also appears broadly similar, though it is slightly weighted towards loans that remain on balance sheet than the HDMA sample.

5.2 Estimator Implementation

Taking the identification strategy to the data requires that three assumptions hold. First, that the counterfactual distribution would be smooth if not for the presence of the notch. Second, that bunchers SSecond, that bunchers come from a continuous set such that there exists a well-defined marginal buncher. Third,

To calculate identify $P_m^o - P^*$, I estimate the counterfactual distribution that would have occurred in the absence of a kink. I identify each transaction's "relative price", the distance between the transaction price and the household's reference price, where the reference price is imputed for each household according to the procedures explicated above. I estimate the counterfactual distribution of relative prices by fitting a polynomial to the counts, but excluding data near the kink, with the following specification:

$$N_j = \sum_{i=0}^q \beta_k^0 (p_i)^k + \sum_{i=R_\ell}^{R_r} \delta_j^0 \mathbb{1}\{i = j\} + \epsilon_j^0 \quad (12)$$

Above, transactions are divided into 250 bins indexed by $j \in \{-200, \dots, 0, \dots, 49\}$. Each bin has a width of \$400 so that the relative prices used in the estimation range from -\$80k to \$20k. For each bin, j , the number of loans is denoted N_j and the relative price at the midpoint of the bin is denoted p_j . R_ℓ and R_r denote the index of the left and right-most bins deemed to contain bunching mass, which I set at -12 and 0 , respectively. And q is the order of the polynomial, which I set to 15.

I use the predicted values from Equation (12) to estimate the counterfactual distribution:

$$\hat{N}_j = \sum_{k=0}^{15} \hat{\beta}_k (d_j)^k \quad (13)$$

I use the difference between observed and counterfactual bin counts in the bunching region, $i \in \{-12, 0\}$, to estimate the excess mass:

$$\hat{B} = \sum_{j=-12}^0 (N_j - \hat{N}_j) = \sum_{j=-12}^0 \hat{\delta}_j \quad (14)$$

I do not observe salient missing mass above the kink point, a common feature of the bunching

estimator literature. Because the bunching mass in the realized distribution must derive from mass above the kink point in the counterfactual distribution, a more accurate estimate of the counterfactual distribution would increase its mass by that of the estimated excess bunching mass. This is a common feature of bunching strategies and I employ an iterative procedure to identify the counterfactual distribution used in the literature (Bachas et al., 2020; Chetty et al., 2011; Kleven and Waseem, 2013). Having estimated a counterfactual distribution, I inflate the mass of transactions above the bunching region by the amount of the bunching mass. In particular, having estimated \hat{B} from 12 and 14, we estimate the following:

$$N_j * \left[1 + \mathbb{1}\{j > R_r\} \frac{\hat{B}}{\sum_{j=R_r+1}^{\infty} N_j} \right] = \sum_{i=0}^q \beta_k(p_i)^k + \sum_{i=R_\ell}^{R_r} \delta_j \mathbb{1}\{i = j\} + \varepsilon_j \quad (15)$$

This yields a new value of \hat{B} which can be used again in Equation (15) iteratively until the value of \hat{B} reaches a fixed point.

I then define the adjustment of the marginal buncher, $P_m^o - P^*$, as the ratio of the bunching mass to the average density of the counterfactual distribution in the bunching region:

$$P_m^o - P^* = \frac{\hat{B}}{\sum_{j=R_\ell}^{R_r} \hat{N}_j / (R_r - R_\ell + 1)} \quad (16)$$

I compute standard errors by following a non-parametric bootstrap procedure. In particular, following Bachas et al. (2020); Chetty et al. (2011), I create new bins of loans by sampling from the estimated vector ε_j and adding these to the estimated distribution. I then use the outlined procedure to compute a new estimate of the marginal buncher's adjustment. To compute the elasticity, I use the median reference price as the reference price. I define the standard error of the elasticity estimate as the standard deviation of estimates due to repeated sampling.

Because down-payment and house prices are related by a constant, I can repeat the estimation procedure with down-payments. In this specification, I consider 200 bins indexed by $j \in \{-100, \dots, 0, \dots, 99\}$, each with a width of \$60 so that the relative down-payments are in the range of -\$6k to \$6k. I set R_ℓ to -1 and R_r to 0. I retain a polynomial of order 15.

6 Results

6.1 Suggestive Evidence

Suggestive evidence of bunching in the cross-section of relative housing quantity comes from simple histograms of realized loan-to-value and debt-to-income ratios on loans. I plot these in

Figure (16a) for the full HMDA FHA sample and in Figure (16b) for the merged sample. There is considerable bunching in both samples at the loan-to-value limit of 96.5% and some modest bunching at the lowest of debt-to-income limits of 43%.³ Although neither form of bunching is exactly the bunching of interest, the fact that both forms show up in the full sample mitigates concerns about using a selected sample for the primary analysis.

The intuition conveyed in these histograms is reiterated in Figure (17). In this figure, I plot the relative down-payments and house prices for transactions in MSAMP3 that take place in Cook County, Illinois. The density of transactions near an apparent left and upper boundary comports with the nature of the household problem described in the section on the empirical approach. The transactions near the left-most boundary are a depiction of bunching in the LTV constraint; transactions near the upper boundary are a depiction of bunching in the DTI constraint. The bunching of interest in our estimator is excess density where the two constraints intersect.

6.2 Loan-to-value Elasticities of Housing Demand

More direct primary evidence comes from Figure (19), which plots down-payments and house prices relative to the reference amount.⁴ Note that the upper and lower panel are really the result of collapsing a full-sample version of Figure (17) onto the x- and y-axes, respectively. The bunching at the reference amount suggests that households are, in fact, credit constrained and therefore rely on financing terms available through the FHA for marginal units of housing purchase.

Figure (18) plots the distribution of reference down-payments and housing prices for the individuals whose transactions are plotted in Figure (19). It is no coincidence that the upper and lower panels look nearly identical; recall that the reference down-payment is the reference price scaled by a function of the loan-to-value limit, and this loan-to-value limit is identical for nearly all FHA transactions. The presence of bunching in this plot is due to the presence of absolute county initial balance limits, which bind instead of DTI limits for all individuals of sufficiently high income in a county. This plot provides evidence that the reference points aligned to generate Figure (19) are heterogeneous due to both differences in income and geography. It allays concerns that the documented bunching is due to some coincident threshold.

Implementation of the bunching estimator is depicted in Figure (20). The dotted black lines designate the assigned area of bunching, the orange line depicts the counterfactual distribution, and the red line depicts the counterfactual location of the marginal buncher. I use the median reference price and down-payment in the formula for the bunching estimator. Using down-payment amounts,

³Where the bunching in debt-to-income realizations appears somewhat off the standard debt-to-income limits, this is because the FHA allows a borrower to roll the mortgage insurance premium into the balance of the loan without counting it toward the debt-to-income limit.

⁴These plots can only be constructed for the merged sample and so sample selection concerns remain.

I obtain an elasticity estimate of 0.025 ± 0.007 ; using house price amounts, I obtain 0.014 ± 0.004 .

These suggest fairly small adjustments to housing demand in response to changes in credit availability. For a hypothetical household formerly bound by a debt-to-income requirement who experiences a relaxation of that requirement, this estimate suggests they would only increase their housing demand by 2.5%. This effect, however, is restricted to households bound by the debt-to-income requirements. Hypothetical households bound by loan-to-value requirements experiencing a relaxation of this requirement of 18% (when mortgage financing becomes more widely available at 95% rather than 80% LTV), would only increase their housing demand by $0.18 * 2.5\% \approx 0.5\%$.⁵

As a final exercise, I test whether the estimate varies in the distribution of incomes. I divide households into groups by whether their annual income at the time of origination was in $[\$0k, \$40k)$, $[\$40k, \$60k)$, $[\$60k, \$80k)$, or $[\$80k, \$\text{inf})$. I plot the distribution of relative downpayments in Figure (21). The evidence suggests some consistency with the idea that higher income households are less credit constrained. Implementing the bunching estimator within groups confirms this in the estimated values, which are .031, .028, .022, and .022, in respective income bins. However, the standard errors on these values are too large to distinguish them statistically. Moreover, the variation in the magnitude does not appear to be sufficiently large to explain differing housing returns in different markets.

7 Conclusion

In this paper, I develop a new estimator of the loan-to-value limit elasticity of housing demand. This estimator has the benefit that it captures only partial equilibrium effects and is therefore more plausibly excludes the confounding effects of changing house price expectations. In addition to providing extensive analysis on how the estimator is constructed and interpreted, I apply it to a sample of FHA loans from recent years. While I find evidence that credit constraints play a role in limiting these households' access to housing on the intensive margin, the magnitude appears to be somewhat limited.

⁵This latter number could change substantially if considering the quasi-income effects of additional loan balance becoming available.

A Tables

Table 1: HMDA Sample[†] Summary Statistics

	Conventional N=93936	FHA-insured N=21299
Loan Characteristics:		
Loan Amount (\$k)	283 (264)	215 (107)
Property Value (\$k)	472 (8881)	243 (550)
Interest Rate (%)	4.43 (0.92)	4.50 (0.65)
Total Costs (\$k)	4.12 (3.93)	7.58 (6.74)
Rate Spread (%)	0.48 (0.86)	1.37 (8.15)
LTV (%)	73.3 (20.3)	92.4 (10.3)
DTI (%)	34.7 (9.90)	42.9 (10.1)
Purpose: Purchase	53.8%	70.8%
Term: 30y	77.0%	96.1%
Rate Type: Fixed Rate	87.6%	98.1%
Repayment Type: Vanilla	91.8%	96.6%
Borrower Characteristics:		
Annual Income (\$k)	178 (6486)	113 (2272)
Age (y)	47.0 (14.7)	42.4 (13.8)
Ethnicity: Hispanic or Latino	10.9%	21.7%
Race:		
Asian	7.31%	2.65%
Black or African American	4.37%	14.9%
White	85.4%	79.1%
Sex:		
Female	23.2%	29.1%
Joint	44.7%	34.3%
Male	32.1%	36.6%
Census Tract Characteristics:		
Population (k)	5.70 (3.18)	5.94 (3.34)
Minority Population (%)	28.7 (23.8)	36.0 (27.5)
Median Income (\$k)	93.8 (39.7)	75.4 (26.8)
1-4 Unit Homeownership Rate (%)	74.3 (14.8)	70.4 (14.6)
Median Age of Housing Units (y)	35.5 (17.8)	34.9 (17.6)

[†] 1pp; 2018-19; 1st Lien; Owner-Occupied; 1-4 Unit Dwelling

Table 2: Merge Process Summary Statistics

	CRISM Sample [†] N=2934	HMDA Sample [†] N=14282	MSAMP3 N=2184
Loan Characteristics:			
Loan Amount (\$k)	206 (100)	213 (101)	206 (102)
Property Value (\$k)	218 (109)	221 (107)	215 (107)
Interest Rate (%)	4.66 (0.64)	4.59 (0.64)	4.65 (0.63)
Total Costs (\$k)	. (.)	7.65 (9.54)	7.24 (3.45)
Rate Spread (%)	. (.)	1.40 (0.56)	1.44 (0.55)
LTV (%)	94.9 (6.48)	96.0 (5.31)	95.8 (4.62)
DTI (%)	41.2 (9.02)	43.3 (9.61)	40.4 (8.99)
Borrower Characteristics:			
Annual Income (\$k)	. (.)	71.9 (37.4)	70.0 (36.0)
FICO Score	673 (45.9)	. (.)	675 (46.3)
Age (y)	. (.)	38.9 (11.9)	38.3 (11.9)
Ethnicity: Hispanic or Latino	.%	23.8%	24.1%
Race:			
Asian	.%	2.50%	2.00%
Black or African American	.%	15.2%	15.1%
White	.%	78.9%	80.1%
Sex:			
Female	.%	28.2%	28.3%
Joint	.%	33.7%	31.9%
Male	.%	38.1%	39.7%
Funding Characteristics:			
Purchaser Type:			
Balance Sheet	.%	41.3%	47.7%
FNMA & FDMC	.%	0.21%	0.10%
GNMA	.%	44.6%	38.3%
Source Type:			
Correspondent	33.8%	.%	37.1%
Retail	50.5%	.%	46.6%
Wholesale	15.6%	.%	16.3%

[†] 1pp; FHA-insured; 2018-19; 1st Lien; Owner-Occupied; 1-4 Unit Dwelling;
Purchase; 30-year; Fixed Rate; Vanilla; Conforming

Table 3: Detailed Summary Statistics for MSAMP3

	N	\bar{x}	s_x	p0	p25	p50	p75	p100
Loan Characteristics:								
Loan Amount (\$k)	218k	206	98.6	20	137	187	257	967
Property Value (\$k)	218k	215	103	20	142	195	267	1150
Interest Rate (%)	218k	4.65	0.64	2	4.25	4.62	5.12	7.12
Total Costs (\$k)	208k	7.35	6.69	0	5.04	6.82	9.05	1710
Rate Spread (%)	215k	1.44	0.56	-4.96	1.05	1.38	1.78	21
LTV (%)	218k	95.9	4.41	13.9	95.5	97.5	98.2	117
DTI (%)	68.9k	41.2	8.97	1	36	42	47	95
Borrower Characteristics:								
Annual Income (\$k)	218k	69.9	36.3	-40	46	63	86	990
FICO Score	191k	675	46.6	342	643	668	702	850
Age (y)	218k	38.5	11.8	20	30	40	50	80
Census Tract Characteristics:								
Population (k)	218k	5.6	2.89	0.18	3.85	5.1	6.66	53.8
Minority Population (%)	218k	36.8	27.9	0	13.5	28.9	56.1	100
Median Income (\$k)	218k	71.5	24.6	0	54.1	68.4	85.3	282
1-4 Unit Homeownership Rate (%)	217k	69.3	14.8	0.31	60	71	80.3	100
Median Age of Housing Units (y)	216k	39.4	17.3	4	26	38	52	76

B Figures

B.1 Housing Choice Sets




Figure 1: The HH choice set in downpayment-housing space is depicted in yellow. It is bounded on the left by the the LTV constraint and above by the DTI constraint, which is really an initial balance constraint. The reference downpayment, D^* , and housing quantity, H^* , are depicted at the point of intersection of the two constraints. Solving the constraints for D and differentiating in H , we obtain the marginal down-payment requirement for an additional unit of housing:

$$\frac{dD}{dH} = \begin{cases} p_H * 1 & H < H^* \\ p_H * (1 - \bar{L}) & H \geq H^* \end{cases}$$




Figure 2: Changing LTV requirements on a loan may apply either to marginal units of housing or to the average unit of housing, i.e. the entire loan. The dashed black line depicts a baseline in which borrowers must comply with a LTV limit, \bar{L}^0 , over the entire loan. The solid black line depicts a regime in which borrowers must comply with the lower LTV limit, \bar{L}' , for marginal units above a reference quantity, H^* . The solid grey line depicts a regime in which borrowers must comply with the lower LTV limit, \bar{L}' , on the entire loan. The behavioral response of households is depicted and, in principle, may differ.

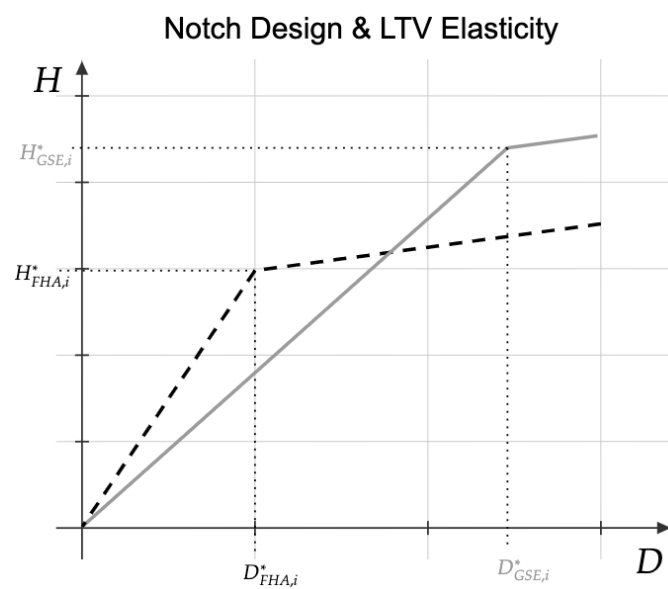


Figure 3: default




Figure 4: default

B.2 Housing Bunching Estimator




Figure 5: The marginal benefits from housing are the sum of utility due to housing services, U_H^0 , and utility due to future consumption afforded by expected excess returns on housing investment, $p_H \mathbb{E}_i [\tilde{r}_H - r_f] * \beta U_X^1$. For the bunching counterfactual, in which a borrower is not subject to a debt-to-income or initial balance limit, the marginal costs are due to a down-payment requirement out of present consumption, $p_H (1 - \bar{L}) * U_X^0$, and mortgage repayment out of future consumption, $p_H \bar{L} (1 + r_f) * \beta U_X^1$. Borrowers obtain the optimal quantity of housing by equating the marginal costs and benefits of an additional unit, H_b^0 .



Figure 6: For the bunching counterfactual, in which a borrower is not subject to a debt-to-income or initial balance limit, households exhibit a distribution of demand for quantities of housing. This heterogeneity may be driven by preferences for housing relative to non-housing consumption, or even wealth orthogonal to the extent of household credit constraints. The optimal choices of four households, $\{i, b, m, e\}$, are depicted. The distribution of optimal housing quantities in the counterfactual is summarized as $f(H)$.

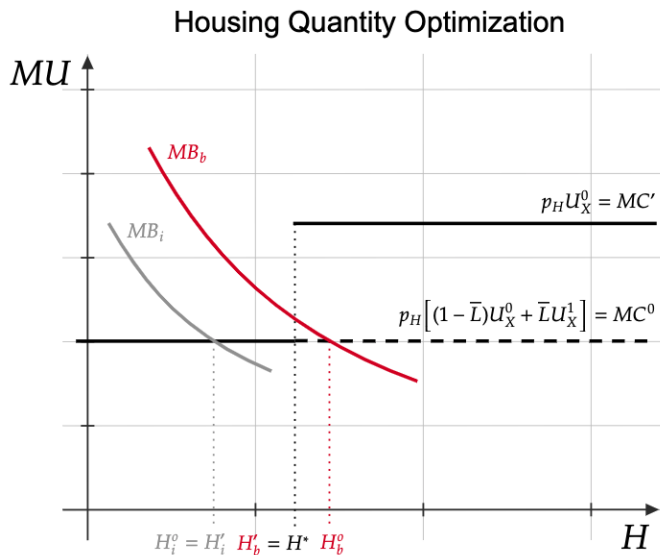


Figure 7: With the introduction of the initial balance constraint, the marginal costs of housing are unchanged under the reference quantity, H^* . Above the reference quantity, the marginal costs are due entirely to down-payment requirements out of present consumption, $p_H * U_X^0$. For unconstrained households, consumption smoothing delivers $U_X^0 = U_X^1$ and $MC' = MC^0$. For constrained households, the constraint implies $U_X^0 > U_X^1$ and $MC' > MC^0$. Borrowers adjust their housing consumption if their counterfactual optimum consumption lies above the reference quantity. Borrower i , in grey, is unaffected. Borrower b , in red, reduces demand from H_b^0 to $H_b' = H^*$.

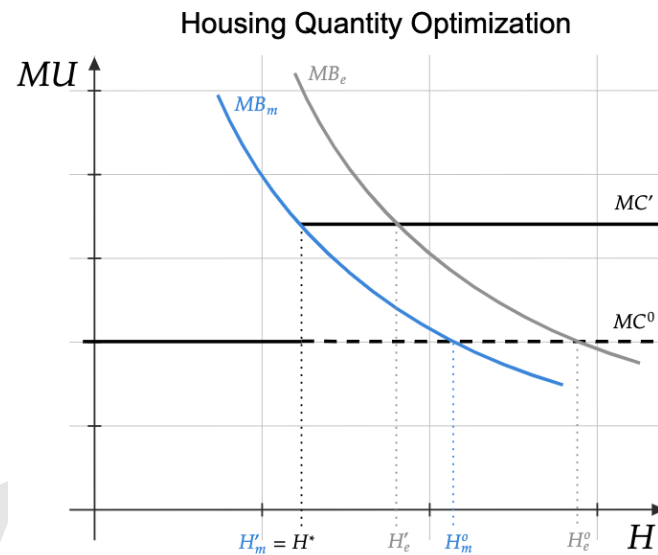


Figure 8: The marginal borrower, m , in blue, is the borrower who adjusts to the reference quantity of housing, H^* , in the presence of the initial balance constraint, and whose marginal costs and benefits of housing are equal after adjustment. Extra-marginal borrowers, e.g. e , may adjust in the presence of the DTI constraint, but does not adjust to the reference quantity.

Housing Quantity Optimization




Figure 9: Bunching is due to all borrowers with housing demand in the counterfactual regime between the reference quantity, H^* , and the marginal buncher's demand, H_m^o . The mass of bunching may be measured as the mass of such borrowers and approximated as $B = f(H^*)[H_m^o - H^*]$. This is the density of borrowers at the reference quantity (under the counter-factual distribution) times the marginal borrower's behavioral response.

B.3 FHA Limits and Lending




Figure 10: Parameters for FHA initial balance limit. Source: HUD website.




Figure 11: Up-front and annual mortgage insurance premium rates for FHA loans. Source: HUD Mortgage Letters.

FHA County-Level Initial Balance Limits, 2008-2020

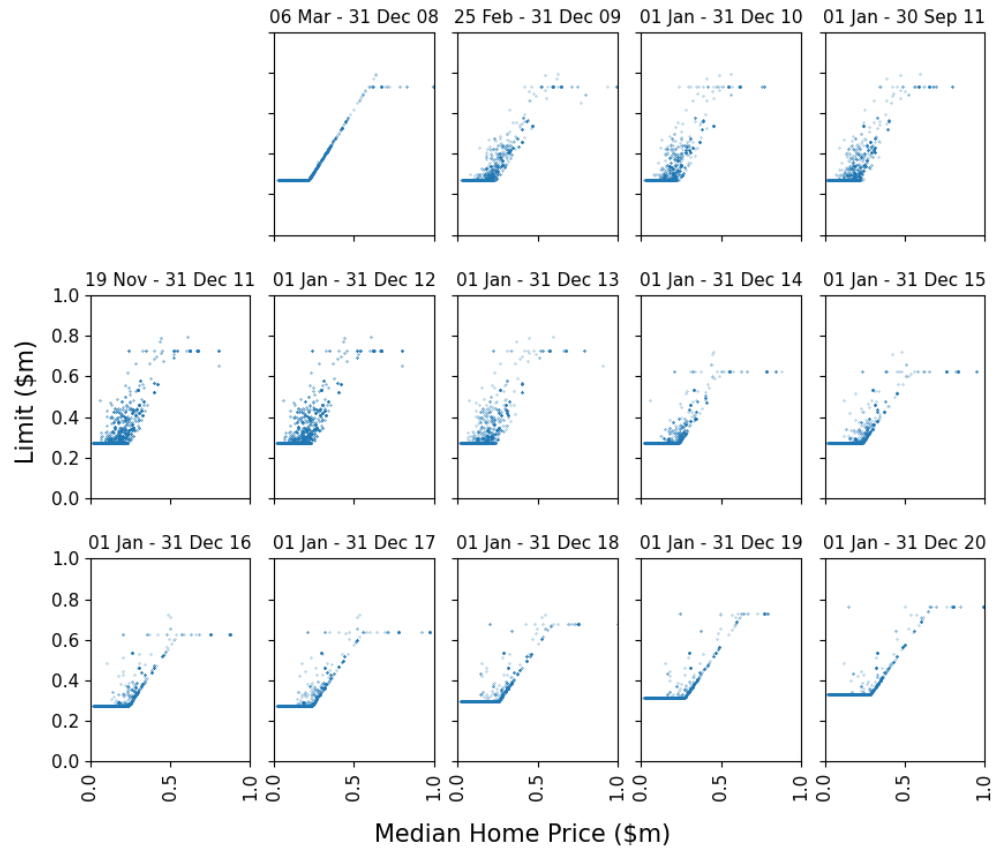


Figure 12: County-level FHA loan limit by median house price and time in force, as designated by the Mortgagee Letter. County limits largely follow assignment rules; noise follows exceptions introduced in the aftermath of the financial crisis. Source: HUD website and HUD Mortgage Letters.




Figure 13: FHA loans have historically constituted 10%-30% of purchase originations. During the Great Recession, they grew to about 1 in 3 purchase originations. They remain around 20% of the market at present.

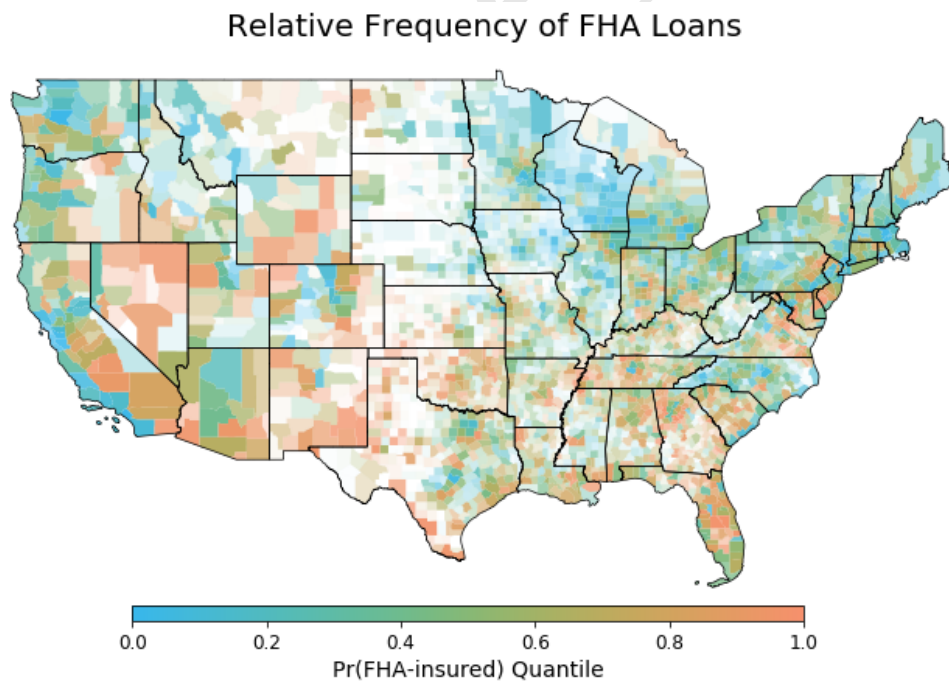


Figure 14: Darker counties have a higher count of FHA purchase originations. Redder counties have a higher relative rate of FHA originations. Urban centers have the most FHA lending though it tends to be relatively high at urban center peripheries.

B.4 Results




Figure 15: This figure depicts the merged sample in the state of California. The four panels depict differences in the HMDA and CRISM recorded values of the Loan Amount, Price, LTV, and Rate for each loan in the final sample. **[I]** The y-axis is a log-scale so that the modest central spikes, for which merged characteristics are almost identical, vastly outnumber those instances of mis-matched errors. **[II]** The steep decrease in match frequency of loans with either loan amounts or loan prices more than \$5k apart. This drop in match frequency is not the result of our match scoring procedure, which moves continuously across the threshold. Instead, it is likely due to the fact that HMDA censors loan amounts and transaction prices at \$5k (CRISM censors at \$1k). Therefore, the censored values of correctly matched loans would have a difference of anywhere from -\$5k to \$5k, depending on their pre-censored values. The merged characteristics are consistent with this.




Figure 16: Above are histograms of LTV and DTI for both the HMDA FHA sample and the merged sample. In both samples, there is bunching in both the LTV and the DTI at limit amounts, denoted in dashed black lines. The disagreement between the limits and the location of the bunching is due to adjustments made for mortgage-insurance premiums. Neither form of bunching is the bunching of interest for the purpose of the analysis but both are suggestive of the existence of such bunching. Finding comparable suggestive evidence in the HMDA sample mitigates concerns about the representativeness of the merged sample.




Figure 17: Above are transactions from the merged sample in Cook County, IL, plotted in downpayment-house price space and normalized to the reference downpayment and house price. The household choice set becomes apparent empirically from this plot and resembles the choice set established in the household problem. The bunching in LTV and DTI are apparent in the relative frequency of borrowers along the left and upper edges of the choice set, respectively. The bunching of interest in this paper concerns the relative frequency of borrowers at the origin, specifically.




Figure 18: Above are histograms of the reference downpayment and house price for every transaction in the merged sample. There is considerable variation in these values, suggesting that any bunching obtained is not the result of a coincident policy. The apparent bunching in these plots is the result of county-level loan limits which are binding for borrowers of sufficiently high income in the county. This is not the bunching of interest but does serve to improve the variability of the kink point relative to borrower characteristics.




Figure 19: Primary evidence of bunching in both downpayment and house prices. Note that both measures map to housing quantity according to $P_i = p_H H_i$ and $D_i = (1 - \bar{L}_i)P_i$.




Figure 20: The bunching estimator implemented on the data in both downpayments and house prices. The black dashed lines designate the area in which the bunching is determined to be present. The red line designates the counterfactual behavior of the bunching agent, as measured by the bunching estimator. Following the literature, an iterative procedure is used to back out the counterfactual density, depicted in orange.



Figure 21: default

C Derivations

C.1 Bunching Behavior under Credit Constraints

Consider a household choosing consumption today X_0 , consumption tomorrow X_1 , housing assets H , a housing down-payment D , and savings A_1 . The household faces a unit price of housing p_H and a liquid savings rate r_f . The household has degenerate beliefs, $\mathbb{E}_i[\tilde{r}_H]$, about the realization of net-of-depreciation housing capital gains rate \tilde{r}_H . The household makes its purchases with income at time-0 and time-1, Y_0 and Y_1 , and capital gains from the sale of the housing asset. It is subject to the standard budget constraints and liquid borrowing constraints.

The household is assumed to purchase housing assets outright so that the decision problem focuses on the intensive margin of housing demand. It may finance the housing purchase using a mortgage loan with an interest rate equal to the savings rate r_f . This mortgage loan must satisfy a standard loan-to-value limit \bar{L} and an initial balance limit \bar{B}^o which may be considered to be derived from an institutional debt-to-income constraint.

The household's problem may be solved in two steps for clarity. For a given choice of housing H and a down-payment D , it chooses the best possible allocation of time-0 and time-1 non-housing consumption, X_0 and X_1 . It then chooses the best possible combination of housing H and down-payment D subject to borrowing constraints, \bar{L} and \bar{B}^o .

Formally, we write:

$$\begin{aligned} \max_{\{H_0, D\}} \quad & \Phi(H_0, D) \\ \text{s.t.} \quad & p_H H_0 - D \leq \bar{B}^o \quad (\zeta) \\ & 1 - \frac{D}{p_H H_0} \leq \bar{L} \quad (\xi) \end{aligned}$$

Where:

$$\begin{aligned} \Phi(H_0, D) = \max_{\{X_0, X_1, A_1, H_1\}} \quad & U(X_0, H_0) + \beta \mathbb{E}[U(X_1, H_1) \mid \tilde{r}_H = \mathbb{E}_i[\tilde{r}_H]] \\ \text{s.t.} \quad & Y_0 = X_0 + \frac{A_1}{1 + r_f} + D \quad (\lambda_0) \\ & Y_1 + A_1 + p_H H_0(1 + \tilde{r}_H) = X_1 + (p_H H_0 - D)(1 + r_f) + r_H H_1 \quad (\lambda_1) \\ & A_1 \geq 0 \quad (\mu_1) \end{aligned}$$

Note that the constraints indexed by ζ and ξ may be written in (H, D) -space. ζ corresponds to Equation (4), which describes an initial balance limit, and ξ corresponds to Equation (5), which describes a loan-to-value limit. These form the boundaries of the kinked opportunity set depicted

in Figure (??).

To identify the household indifference curves in (H, D) -space, plug the budget constraints into the objective function and use the degenerate beliefs to simplify the expectations operator:

$$\begin{aligned} \Phi(H_0, D) = \max_{\{A_1, H_1\}} & U\left(Y_0 - \frac{A_1}{1+r_f} - D, H_0\right) + \beta U\left(Y_1 + A_1 + p_H H_0 \mathbb{E}_i[\tilde{r}_H - r_f] + D(1+r_f) - r_H H_1, H_1\right) \\ \text{s.t.} \quad & A_1 \geq 0 \quad (\mu_1) \end{aligned}$$

Use the first-order condition for H_1 to obtain:

$$\frac{U_H^1}{U_X^1} = r_H \quad (17)$$

Use the first-order condition for A_1 to solve for the unconstrained optimal savings \hat{A}_1 . The constrained optimal savings is given by $A_1^* = \max\{\hat{A}_1, 0\}$. For simplicity, I assume separability in U and that $\beta(1+r_f) = 1$. Then we have:

$$\hat{A}_1 = \frac{1+r}{2+r} \left[Y_0 - Y_1 - p_H H \mathbb{E}_i[\tilde{r}_H - r_f] + r_H H_1^* \right] - D(1+r) \quad (18)$$

Note that Equation (18) with $\hat{A}_1 = 0$ defines a line in (H, D) -space that divides cases in which the household is constrained from those in which it is not. This line is depicted in Figure (??), dividing red from blue.

Plug the closed-form representation of A_1^* back into the objective function to obtain a closed-form representation of Φ . For clarity, define the household's wealth given its housing investment, $W(H_0)$, as its income and excess capital gains from housing discounted to time 0. Formally, $W(H_0) = Y_0 + \frac{1}{1+r} Y_1 + \frac{1}{1+r} p_H H_0 (\mathbb{E}_i[\tilde{r}_H] - r_f) - \frac{1}{2+r} r_H H_1^*$. Then we have:

$$\Phi(H_0, D) = \begin{cases} U\left(\frac{1+r}{2+r} W(H_0), H_0\right) + \beta U\left(\frac{1+r}{2+r} W(H_0), H_1^*(H_0)\right) & \hat{A}_1 \geq 0 \\ U\left(Y_0 - D, H_0\right) + \beta U\left(Y_1 + p_H H_0 \mathbb{E}_i[\tilde{r}_H - r_f] + D(1+r) - r_H H_1^*, H_1^*\right) & \hat{A}_1 \leq 0 \end{cases}$$

Using this closed-form representation, it is possible to derive some intuitive characteristics of indifference curves in (H, D) -space. I show below that these indifference curves are convolutions of consumer preferences and (non-mortgage) borrowing constraints. I depict these indifference curves in Figure (??).

In the case when the household is not constrained, $\hat{A}_1 \geq 0$, the down-payment, D , does not enter its utility. Intuitively, in the unconstrained case, it is possible for the household to borrow or save around a down-payment. As a result, the borrower has horizontal indifference curves in

(H, D) -space. This is represented in blue in Figure (??). Horizontal indifference curves will not produce bunching and so any bunching in the data constitutes evidence of financial frictions.

In the case when the borrower is constrained, making additional down-payment effectively sacrifices consumption today for consumption tomorrow, which is costly in utility terms because a high marginal utility of consumption today is what drives to consumer against her borrowing constraint. We can see this by setting $\Phi = \bar{\Phi}$ in the constrained case. Consider H_0 an implicitly defined function of D , differentiate, and solve for $H'_0(D)$. For clarity, substitute the values of consumption, \hat{X}_0, \hat{X}_1 , implied by the choice of down-payment and housing. Use the results from the first-order condition in H_1 , Equation (17), to cancel terms in $H_1^*(D)$.

$$H'_0(D) = \frac{U_X^0 - U_X^1}{U_H^0 + \beta U_X^1 p_H \mathbb{E}_i[\tilde{r}_H - r_f]}$$

To help with interpretation, consider the case in which HH do not believe that housing will have any excess capital gains relative to liquid assets. And note that the difference in marginal utility of consumption at time 0 and 1 is the Lagrange multiplier on the liquid savings constraint.

$$H'_0(D) \stackrel{\hat{r}=r_f}{=} \frac{U_X^0 - U_X^1}{U_H^0} = \frac{\mu_1}{U_H^0}$$

These indifference curves are represented in red in Figure (??).

Intuitively, additional down-payment at time 0 incurs a utility cost that is the difference between marginal utility of non-housing consumption in time-0 and time-1. The more money put into the down-payment, the more recovered from the sale of the home. Additional time-0 housing has utility value due to the housing services enjoyed at time-0 as well as potential investment value that can be enjoyed as consumption at time-1.

The numerator of this expression is positive because households run up against their liquid borrowing constraints because they are having trouble moving consumption to the present. More technically, the Kuhn-Tucker conditions require that $\mu_1 > 0$ when $A_1^* = 0$. The denominator is positive provided that housing is not expected to perform poorly because the marginal utility of additional housing services is positive. During the run-up to the crisis, housing was generally expected to do well.

Finally, it is worth noting that extreme housing optimism flattens the indifference curve.

C.2 The Bunching Estimator and Curvature of the Indifference Curve

Here, I argue that the bunching estimator, which recovers an empirical measure of the down-payment adjustment of the marginal buncher, Δd , can be used to understand the curvature of the

marginal buncher's (indirect) indifference curve, \tilde{h} .

Consider a kinked opportunity set in (d, h) -space. Denote the slope of the lower and upper portion of the kinked opportunity set as μ_0 and μ_1 , respectively. Denote the downpayment and housing quantity at the kink point be denoted by \hat{d} and \hat{h} , respectively.

The marginal buncher, m , is then defined as optimizing relative to the upper portion of the kinked opportunity set at (\hat{d}, \hat{h}) and optimizing relative to the lower portion of the kinked opportunity set at $(\hat{d} + \Delta d, \hat{h} + \Delta h)$. Suppose the kink is small, so that the buncher remains on the same indifference curve; we denote this indifference curve as $\tilde{h} : D \rightarrow H$. Then we have:

$$\tilde{h}(\hat{d}) = \hat{h} \quad (19)$$

$$\tilde{h}'(\hat{d}) = \mu_1 \quad (20)$$

$$\tilde{h}(\hat{d} + \Delta d) = \hat{h} + \Delta h \quad (21)$$

$$\tilde{h}'(\hat{d} + \Delta d) = \mu_0 \quad (22)$$

Begin with an approximation of $\tilde{h}''(\hat{d})$:

$$\tilde{h}''(\hat{d}) = \frac{\tilde{h}'(\hat{d} + \Delta d) - \tilde{h}'(\hat{d})}{\Delta d}$$

Multiply both sides by $\frac{\hat{d}}{\tilde{h}'(\hat{d})}$ and rearrange:

$$\frac{\frac{\tilde{h}'(\hat{d} + \Delta d) - \tilde{h}'(\hat{d})}{\Delta d}}{\frac{\hat{d}}{\tilde{h}'(\hat{d})}} = \frac{\hat{d} \tilde{h}''(\hat{d})}{\tilde{h}'(\hat{d})}$$

We can use 20 and 22 in the above and rewrite terms for clarity to obtain:

$$\frac{\% \Delta \mu}{\% \Delta d} = \frac{d \ln \tilde{h}'(d)}{d \ln d} \Big|_{d=\hat{d}} = \epsilon_d^{\tilde{h}'}$$

Note that administrative parameters governing borrowing constraints through the FHA can be used to recover $\% \Delta \mu$ and \hat{d} and the bunching estimator can be used to recover Δd , so it is possible to estimate $\epsilon_d^{\tilde{h}'}$.

C.3 Compensated and Uncompensated Elasticity of Housing Demand

This proof follows the observation in Saez (2010), the formalization of this observation and sketched proof in Kleven (2016).

Consider the introduction of a kink into the loan-offer curve facing a borrower. Consider any borrower, i , who responds to the introduction of the kink and whose indifference curves in (h, d) -space are tangent to the the loan-offer curve both before and after the introduction of the kink. Note that the marginal borrower, whose behavioral change the bunching technique is designed to measure, is one such borrower.

Define borrower i 's observed elasticity as $\hat{e}_i \equiv \frac{\% \Delta D_i}{\% \Delta \mu}$. Note that the observed elasticity can be decomposed into the joint effect of two elasticities, a compensated LTV elasticity of down-payment supply, e_i^c , and an (income-like) initial-loan elasticity of down-payment supply, η_i . We write:

$$\hat{e}_i = e_i^c + \frac{\% \Delta \tilde{B}}{\% \Delta \mu} \eta_i$$

Use the Slutsky-decomposition, $e_i^c = e_i^u - s_i^D \eta_i$, to rewrite the expression. Note that the down-payment share of initial-loan balance can be expressed as, $s_i^D = \frac{D_i^o \mu^o}{\tilde{B}^o}$. Simplify terms.

$$= e_i^c + \frac{\Delta \tilde{B}}{\tilde{B}^o} \frac{\mu^o}{\Delta \mu} \frac{\tilde{B}^o}{D_i^o \mu^o} (e_i^u - e_i^c) = e_i^c + \frac{\Delta \tilde{B} / D_i^o}{\Delta \mu} (e_i^u - e_i^c)$$

Note that $\Delta \tilde{B} = \Delta \mu D_i^* - \Delta \mu D_i^o$ where the first term is due to an effective increase in the initial-loan balance and the second term is the effective fall in the initial-loan balance due to the down-payment funds supplied. Substitute in this expression and simplify:

$$= e_i^c + \frac{D_i^* - D_i^o}{D_i^o} (e_i^u - e_i^c)$$

Finally, note that we can write the change in the average (cf. marginal) slope as: $\Delta a = \frac{D_i^* \mu^o + (z - z^*) \mu'}{D^o} - \mu^o = \frac{D_i^* - D_i^o}{D_i^o} \Delta \mu$. Substitute in this expression and rearrange terms:

$$= e_i^c + \frac{\Delta a}{\Delta \mu} (e_i^u - e_i^c) = \left[1 - \frac{\Delta a}{\Delta \mu} \right] * e_i^c + \left[\frac{\Delta a}{\Delta \mu} \right] * e_i^u$$

Note that when the kink is small and the change in average slope for the marginal buncher is near 0, this simplifies to the compensated elasticity, e_i^c . What is most obviously of interest in this setting, however, is the uncompensated elasticity of housing demand, e_i^u . In the run-up to the crisis, the loan-to-value on the entire loan increased, not just on the portion of the loan above some threshold.

This estimation is helped in a two regards. First, we exploit a fairly large-sized kink, which, by the above proof, will capture a weighted average of compensated and uncompensated housing demand elasticities, so that to the extent that these values differ, we are at least capturing some

of the uncompensated demand. Second, in this setting, it is reasonable to suppose that, although housing demand may respond dollar-for-dollar, down-payment supply does not respond to an increase in the initial available initial loan balance. In this case the compensated and uncompensated elasticities estimated here are close. It is worth noting, finally, that when estimating structural parameters, the bunching literature generally assumes functional forms for which compensated and uncompensated elasticities of earnings supply are identical.

C.4 The Bunching Estimator and Loan-to-Value Elasticity of Housing Demand

The curvature of the indirect indifference curves is responsible for generating bunching. The most natural estimator in the standard bunching framework, as described in the section above, is $\frac{\% \Delta d}{\% \Delta \mu}$, or the value-to-downpayment elasticity of the down-payment. A more natural object of interest for policy purposes is understanding how housing demand responds to the loan-to-value ratio imposed by the lending program. I show how to transform the traditional bunching estimator into one better suited for policy analysis in this application. I proceed somewhat informally, in the hopes of conveying intuition:

We note the fact that $h = \frac{1}{p_h} \frac{d}{1-L_0}$ both at (\hat{d}, \hat{h}) and at $(\hat{d} + \Delta d, \hat{h} + \Delta h)$ and totally differentiate to derive:

$$\frac{dh}{h} = \frac{\frac{dd}{1-L_0}}{\frac{d}{1-L_0}} = \frac{dd}{d}$$

We also note the fact that $\mu = \frac{1}{1-L}$ where L is the loan-to-value ratio of the marginal dollar of housing purchased and totally differentiate to derive:

$$\frac{d\mu}{\mu} = \frac{\frac{1}{(1-L)^2} dL}{\frac{1}{1-L}} = \frac{dL}{L} \frac{L}{1-L}$$

Plugging these expressions into our elasticity of interest, we obtain:

$$\frac{d \ln h}{d \ln L} = \frac{\frac{dh}{h}}{\frac{dL}{L}} = \frac{\frac{dd}{d}}{\frac{1-L}{L} \frac{d\mu}{\mu}} = \frac{L_0}{1-L_0} \frac{d \ln d}{d \ln \mu} = \frac{L}{1-L} \epsilon_d^{\tilde{h}'}$$

C.5 Analogy to Standard Estimators

Kinked budget constraints are a common feature of household choice sets; they are generated by government transfer programs, e.g. income tax schedules, and, in turn, generate bunching behavioral responses from optimizing agents. Bunching estimators commonly exploit this behavior to

estimate real elasticities of fundamental interest, such as the tax elasticity of labor supply. Kinked borrowing constraints are a common feature of government lending programs, e.g. those that restrict both the LTV and loan amount, and, in principle, may generate bunching behavior as well. This behavior may be exploited to estimate financial elasticities of interest, like the elasticity of housing demand with respect to loan-to-value requirements.

References

- A Acolin, J Bricker, P Calem, and S Wachter. Borrowing Constraints and Homeownership. *American Economic Review: Papers and Proceedings*, 106(5):625–629, 2016. doi: 10.1257/aer.p20161084. URL http://repository.upenn.edu/fnce_papers<http://dx.doi.org/10.1257/aer.p20161084>.
- Manuel Adelino, Kristopher Gerardi, and Paul S. Willen. Why don't Lenders renegotiate more home mortgages? Redefaults, self-cures and securitization. *Journal of Monetary Economics*, 60(7):835–853, oct 2013. ISSN 03043932. doi: 10.1016/j.jmoneco.2013.08.002.
- Elliot Anenberg, Aurel Hizmo, Edward Kung, and Raven Molloy. Measuring Mortgage Credit Availability: A Frontier Estimation Approach Measuring Mortgage Credit Availability: A Frontier Estimation Approach *. 2017. URL <https://doi.org/10.17016/FEDS.2017.101>.
- Natalie Bachas, Olivia S. Kim, and Constantine Yannelis. Loan guarantees and credit supply. *Journal of Financial Economics*, aug 2020. ISSN 0304405X. doi: 10.1016/j.jfineco.2020.08.008.
- Michael Bailey, Eduardo Dávila, Theresa Kuchler, and Johannes Stroebel. House Price Beliefs And Mortgage Leverage Choice. *The Review of Economic Studies*, 86(6):2403–2452, nov 2019. ISSN 0034-6527. doi: 10.1093/restud/rdy068. URL <https://academic.oup.com/restud/article/86/6/2403/5194341>.
- Robert Bartlett, Adair Morse, Richard Stanton, and Nancy Wallace. Consumer-Lending Discrimination in the FinTech Era *. *Journal of Financial Economics*, 2019.
- Michael Carlos Best, James Cloyne, Ethan Ilzetzki, and Henrik Kleven. Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches. Technical report, 2018. URL <http://www.nber.org/papers/w24948>.
- Richard Blundell and Hilary W Hoynes. Has 'In-Work' Benefit Reform Helped the Labor Market? In David Card, Richard Blundell, and Richard B. Freeman, editors, *Seeking a Premier Economy*:

- The Economic Effects of British Economic Reforms, 1980-2000*, chapter 10, pages 411–459. University of Chicago Press, 2004. ISBN 0226092844. URL <http://www.nber.org/books/card04-1>.
- Richard Blundell and Andrew Shephard. Employment, hours of work and the optimal taxation of low-income families. *Review of Economic Studies*, 79(2):481–510, apr 2012. ISSN 00346527. doi: 10.1093/restud/rdr034. URL <https://academic.oup.com/restud/article/79/2/481/1532543>.
- Raj Chetty, John N Friedman, Tore Olsen, and Luigi Pistaferri. ADJUSTMENT COSTS, FIRM RESPONSES, AND MICRO VS. MACRO LABOR SUPPLY ELASTICITIES: EVIDENCE FROM DANISH TAX RECORDS *. *Quarterly Journal of Economics*, (126):749–804, 2011. doi: 10.1093/qje/qjr013. URL <https://academic.oup.com/qje/article/126/2/749/1870394>.
- Ramona Dagostino. The Impact of Bank Financing on Municipalities’ Bond Issuance and the Real Economy. Technical report, 2019.
- Anthony A. DeFusco and John Mondragon. No Job, No Money, No Refi: Frictions to Refinancing in a Recession. *Journal of Finance*, 75(5):2327–2376, 2020. ISSN 15406261. doi: 10.1111/jofi.12952.
- Anthony A. DeFusco and Andrew Paciorek. The interest rate elasticity of mortgage demand: Evidence from bunching at the conforming loan limit. *American Economic Journal: Economic Policy*, 9(1):210–240, feb 2017. ISSN 1945774X. doi: 10.1257/pol.20140108. URL <https://doi.org/10.1257/pol.20140108>.
- Mehran Ebrahimian. Student Loans and Social Mobility * †. 2020. URL <https://www.newyorkfed.org/medialibrary/>.
- Giovanni Favara and Jean Imbs. Credit Supply and the Price of Housing †. *American Economic Review*, 105(3):958–992, 2015. doi: 10.1257/aer.20121416. URL <http://dx.doi.org/10.1257/aer.20121416958>.
- FDIC. AFFORDABLE MORTGAGE LENDING GUIDE A Resource for Community Bankers Part I: Federal Agencies and Government Sponsored Enterprises. Technical report, 2018. URL <https://www.fdic.gov/mortgagelending>.
- Min Hwang, Chen L Miller, and Robert Van Order. Effects of FHA loan limit Increases by ESA 2008: Housing Demand and Adverse Selection. Technical report, 2016. URL <https://ssrn.com/abstract=2789620>.

- Dan Immergluck. From minor to major player: The geography of fha lending during the u.s. mortgage crisis. *Journal of Urban Affairs*, 33(1):1–20, 2011. ISSN 07352166. doi: 10.1111/j.1467-9906.2010.00539.x.
- Stephanie Johnson. Mortgage Leverage and House Prices *. 2019.
- Alejandro Justiniano, Giorgio E Primiceri, and Andrea Tambalotti. Credit Supply and the Housing Boom. *Journal of Political Economy*, 127(3):1317–1350, 2019.
- Greg Kaplan, Kurt Mitman, and Giovanni L Violante. The Housing Boom and Bust: Model Meets Evidence. *Journal of Political Economy*, 128(9), 2020.
- Henrik J Kleven and Mazhar Waseem. USING NOTCHES TO UNCOVER OPTIMIZATION FRICTIONS AND STRUCTURAL ELASTICITIES: THEORY AND EVIDENCE FROM PAKISTAN*. *Quarterly Journal of Economics*, 128(2):669–723, 2013. doi: 10.1093/qje/qjt004. URL <https://academic.oup.com/qje/article/128/2/669/1943151>.
- Henrik Jacobsen Kleven. Bunching. *Annual Review of Economics*, 8:435–464, 2016. doi: 10.1146/annurev-economics-080315-015234. URL www.annualreviews.org.
- Tim Landvoigt, Monika Piazzesi, and Martin Schneider. The Housing Market(s) of San Diego †. *American Economic Review*, 105(4):1371–1407, 2015. doi: 10.1257/aer.20111662. URL <http://dx.doi.org/10.1257/aer.20111662>.
- Day Manoli and Andrea Weber. Nonparametric Evidence on the Effects of Financial Incentives on Retirement Decisions †. *American Economic Journal: Economic Policy*, 8(4):160–182, 2016. doi: 10.1257/pol.20140209. URL <http://dx.doi.org/10.1257/pol.20140209>.
- Atif Mian and Amir Sufi. The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis *. *Quarterly Journal of Economics*, 124(4):1449–1496, nov 2009. ISSN 0033-5533. doi: 10.1162/qjec.2009.124.4.1449. URL <https://academic.oup.com/qje/article-lookup/doi/10.1162/qjec.2009.124.4.1449>.
- Atif Mian and Amir Sufi. CREDIT SUPPLY AND HOUSING SPECULATION. 2018. URL <http://www.nber.org/papers/w24823>.
- Kevin A. Park. FHA loan performance and adverse selection in mortgage insurance. *Journal of Housing Economics*, 34:82–97, dec 2016. ISSN 10960791. doi: 10.1016/j.jhe.2016.07.004.
- Wayne Passmore and Shane M Sherlund. Government-Backed Mortgage Insurance, Financial Crisis, and the Recovery from the Great Recession. Technical report, 2016. URL <http://dx.doi.org/10.17016/FEDS.2016.031>.

Jim Puzzanghera. FHA to get \$1.7 billion in its first taxpayer-funded bailout - Los Angeles Times, sep 2013. URL <https://www.latimes.com/business/la-xpm-2013-sep-28-la-fi-0928-fha-bailout-20130928-story.html>.

Shanthi Ramnath. Taxpayers' responses to tax-based incentives for retirement savings: Evidence from the Saver's Credit notch. *Journal of Public Economics*, 101(1):77–93, may 2013. ISSN 00472727. doi: 10.1016/j.jpubeco.2013.02.010.

Emmanuel Saez. Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy*, 2(3):180–202, 2010. doi: 10.1257/pol.2.3.180. URL <http://www.aeaweb.org/articles.php?doi=10.1257/pol.2.3.180>.

James M. Sallee and Joel Slemrod. Car notches: Strategic automaker responses to fuel economy policy. *Journal of Public Economics*, 96(11-12):981–999, dec 2012. ISSN 00472727. doi: 10.1016/j.jpubeco.2012.06.005.

Aaron S Yelowitz. The Medicaid Notch, Labor Supply, and Welfare Participation: Evidence from Eligibility Expansions. Technical Report 4, 1995. URL <https://www.jstor.org/stable/2946644>.