Federal Reserve Bank of Chicago

Why were interest-only mortgages so popular during the U.S. housing boom?

Gadi Barlevy and Jonas D.M. Fisher

REVISED
January 9, 2020

WP 2010-12

Why were interest-only mortgages so popular during the U.S. housing boom?*

Gadi Barlevy
Federal Reserve Bank of Chicago
gbarlevy@frbchi.org

Jonas D.M. Fisher
Federal Reserve Bank of Chicago
jfisher@frbchi.org

January 9, 2020

Abstract

Borrowers in U.S. cities where house prices boomed in the 2000s relied heavily on backloaded interest-only (IO) mortgages that require borrowers to only pay interest for the first few years of the loan. We develop a theory that encompasses common explanations for IO use, and show that while they can account for much of the regional variation in IOs, they cannot explain why IOs were popular in boom cities. We propose a new explanation. In our model, uncertain price appreciation coupled with non-recourse lending can lead to speculation financed with backloaded mortgages. We find evidence that IO borrowers behaved in ways consistent with such speculation, and discuss the policy implications of our findings.

JEL Classification Numbers: E0, O4, R0
Keywords: Housing, house prices, interest-only mortgages, speculation, bubble

*This is a substantial revision to a paper previously circulated under the title "Mortgage choices and housing speculation." We thank Sumit Agarwal, Gene Amromin, David Andolfatto, Marco Bassetto, Martin Eichenbaum, Huberto Ennis, Simon Gilchrist, Lars Hansen, Jonathan Heathcote, Alejandro Justiniano, Benoit Mojon, Alex Monge, Erwan Quintin, Monika Piazzesi, Ned Prescott, Amit Seru, Amir Sufi, and Randy Wright for helpful comments. We are grateful to Albert Saiz for sharing his geographical data. We also thank Keyoung Lee, Ryan Peters, Shani Schechter and Trevor Serrao for their excellent research assistance. The views expressed herein are the authors' and do not necessarily represent the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.

1. Introduction

It is now well established that the U.S. housing boom of the 2000s was associated with a substantial increase in nontraditional mortgages.[1] Many of these mortgages featured backloaded payments in which borrowers were asked to make low payments early on and then substantially higher payments later. The most popular mortgage of this type was the interest-only (IO) mortgage which obliges the borrower to pay just their interest obligation for some initial predetermined period – usually 3, 5, or 10 years – and only then requires borrowers to pay both interest and principal. As we show in this paper, IOs were not just more popular when house prices rose, but they were a key feature of cities with rapid house price growth. The share of IOs is tightly correlated with the rate of house price growth in a city even after controlling for other mortgage characteristics such as the share of subprime loans, the share of privately securitized loans, and the share of mortgages with high leverage.

At the same time, there is no consensus on why IOs surged in popularity, and whether this growth was a cause for concern. Among policymakers, the rise of these mortgages in markets with rapid house price growth raised alarms. In a GAO report to Congress, Williams (2006) questioned the sustainability of these mortgages, arguing that borrowers may not have been aware of all their inherent risks and that some borrowers relied on them to purchase homes for investment purposes and were more likely to default if faced with financial distress. The report suggests policymakers should consider regulating the use of these products.

In contrast, economic research has highlighted various benefits of such mortgages, questioning the need to restrict them. Cocco (2013) finds in United Kingdom data between 2000 and 2008 that IO loans were primarily used by borrowers whose income was predictably set to grow and who could use such loans to relax their liquidity constraints. Chiang and Sa-Aadu (2014) argue that IO mortgages can benefit more mobile households. Dokko, Keys, and Relihan (2019) find that IOs and other non-traditional mortgages facilitated home purchases in areas where housing became expensive. Similarly, Bäckman and Lutz (2017) argue that affordability was behind the surge in IOs in Denmark. LaCour-Little and Yang (2010) and Brueckner, Calem, and Nakamura (2016) argue that expected future house price growth

[1] See, for example, Amromin, Huang, Sialm, and Zhong (2018).

led lenders to offer mortgages with backloaded payments that borrowers prefer but would be too expensive in the absence of expected house price growth. If mortgages with backloaded payments confer the benefits identified in these papers, interfering with the ability to enter such contracts could make agents worse off.

Whether the rise in IO use demands a policy response depends on why such mortgages were popular in cities with rapid house price growth. To address this question, we construct a theoretical framework that features endogenous house prices and mortgage choice. We use this model to guide our empirical exploration of the various explanations for the appeal of IOs. Our evidence suggests these explanations can account for much of the variation in IO use across cities. At the same time, we find that they cannot account for the popularity of IOs in cities with rapid house price appreciation. Essentially, cities where house prices took off most quickly did not feature systematically higher expected income growth, more mobile populations, less affordable housing, or evidence of a preference for backloading, but borrowers in these cities did show a systematic penchant for using IOs.

Since the residual variation in IOs after controlling for these factors remains strongly correlated with house price appreciation, we turn to a different and novel explanation for the concentration of IOs in cities with rapid price appreciation. In our model, if mortgages are not subject to recourse, meaning a lender cannot seize the borrower's income in case of default, agents may have an incentive to speculate when house price growth is risky, profiting if house prices rise but defaulting if they fall. Demand by speculators can push house prices above the expected present value of housing services of the marginal house. This overvaluation would encourage IO use. Speculators find IOs attractive since they can shift more risk to lenders, but lenders also like them because they encourage borrowers to sell overvalued properties earlier than they would otherwise, exposing themselves to less risk. Consistent with this, we find that for the cities with high price growth, IOs were more popular in non-recourse states. We also find that the borrowers who used IOs behaved in a manner consistent with speculation: They were more likely to sell their property when house prices rose and more likely to default when house prices fell than those with traditional mortgages.

To the extent that the rise of IOs in cities with rapid house price growth is driven by speculation, this would suggest a role for policy. Speculators in our model take on too much leverage because they fail to internalize the costs of default that their lenders incur, so there is scope for intervention. At the same time, our analysis suggests IOs might be a symptom rather than a source of the underlying problem, and restricting IOs may be inappropriate. First, we do find that much of the variation in IO use across cities can be explained by factors that imply borrowers benefit from IOs, suggesting restricting these contracts, even if only when we see a rise in their use in booming cities, would likely harm many borrowers. Second, in our model, it is not the introduction of IOs that causes speculation, but the arrival of opportunities for speculation that lead agents to use IOs. Preventing agents from using backloaded mortgages would not discourage speculation, and agents would have an incentive to speculate even with traditional mortgages. A more effective intervention would grant lenders recourse or impose down payment requirements to directly discourage speculation.

Our analysis also suggests that the rise in backloaded mortgages may be a useful beacon for alerting policymakers to the presence of speculation, and so an additional reason not to prevent their use is that it would destroy information useful to policymakers. However, our model also cautions against blindly focusing on the overall use in IOs as a measure of speculation, since we find they appeal to borrowers for other reasons. Sullivan (2018) reports a recent resurgence in IOs and cites it as a cause for concern. But if this latest rise in IOs were concentrated in recourse states, as extending our data to recent years suggests, this increase may not be a sign of rising speculation but a more benign indicator.

2. Evidence on mortgage choices across cities

In this section, we review the evidence on nontraditional mortgages, particularly backloaded mortgages, and argue that IOs robustly predict rapid price appreciation. The two most popular backloaded mortgages at this time were IOs, which we described above, and Option-ARMs, which give borrowers the option to pay less than the required interest and so increase their principal obligation, up to some amount, for a predetermined period.

Table 1 reports summary statistics for different types of for-purchase first-lien mortgages in 2003, before the surge in house prices, and 2006, when house prices peaked. Our data comes from Black Knight Inc, previously known as the McDash dataset. It is compiled by mortgage servicers, and covers 9 out of the top 10 servicers and about 60% of the total mort-

Table 1: Mortgages originated for purchase in 2003 and 2006

				Share of Mortgage Type			
Type	Year	Share of total	Mean Amount	Sub-Prime	Private Securitized	Pre Pay Penalty	Investor
IO	2003	2.2	339.2	12.2	71.0	11.9	7.2
	2006	19.0	328.4	9.5	61.4	17.6	11.3
Option ARM	2003	3.1	280.9	10.1	24.0	41.5	3.3
	2006	5.0	341.5	24.5	69.6	65.7	13.5
Fixed (Not backloaded)	2003	73.0	157.0	1.9	12.8	4.9	5.2
	2006	56.0	184.5	2.9	15.7	2.7	7.4
ARM (Not backloaded)	2003	13.1	219.0	11.3	22.4	17.1	5.6
	2006	11.3	230.6	28.5	67.3	28.9	9.1
Other	2003	8.6	146.8	1.0	7.2	1.4	13.5
	2006	8.6	191.3	14.1	35.1	18.8	14.3
All	2003	100.0	172.1	3.5	15.2	7.7	6.1
	2006	100.0	225.5	9.1	35.0	12.8	9.3

Note: Shares are in percentage points and mean amount is in thousands of current dollars.

gage market during the period we study. We describe our data in more detail in Appendix A. Given our interest in backloading, we report statistics for IOs and Option-ARMs and then divide the remaining mortgages into fixed rate, adjustable rate (ARM), and a remaining catch-all "Other" category that accounts for less than 9% of the sample.

Table 1 highlights the rapid growth in backloading during this period. In 2003, IOs represented 2.2% of mortgages in our dataset, compared to 19% in 2006. Option-ARMs also grew, but not nearly as dramatically.[2] Although both mortgages feature backloaded payments, they differ in several ways. First, borrowers who took out IOs were primarily prime borrowers with high credit rating; fewer than 10% were subprime in 2006, compared with 24% of borrowers with Option-ARMs. The two kinds of mortgages also differ in terms of the frequency with which they featured prepayment penalties. Less than 20% of IOs in 2006 featured prepayment penalties, compared to nearly two thirds of Option-ARMs. Option-ARMs were much less likely to be privately securitized than IOs in 2003 but more likely in 2006. Option ARMs thus seem to involve riskier borrowers, more restrictions, and were increasingly more likely to be securitized over time.[3] We will focus primarily on IOs. Not only were they more popular, but, as we show later on, they also seem more closely related to house price growth across cities.

We now establish our finding that during the housing boom IO use was strongly correlated with rapid house price appreciation at the city level even after controlling for other mortgage characteristics. To measure house price appreciation, we first identified in each city the quarter between 2003q1 and 2008q4 in which real house prices peaked. We then computed the maximum 4-quarter log real price growth between 2003q1 and the quarter when house prices peaked in each city. This captures cities where house prices grew a lot over a short period of time, and seems to assign a high rank to the cities most commonly associated with a property boom.[4] As a measure of IO use we computed the maximum IO share among all

[2] The fact that some IOs and Option-ARMs were offered in 2003 underscores the fact that, while rare before the boom, neither product was strictly new. For example, Golden West Financial in California offered option-ARM mortgages back in 1981.

[3] See Amromin et al. (2018) for additional comparisons of borrowers with IOs and option-ARMs.

[4] For example, the two cities with the highest maximum 4-quarter price growth in our data are Las Vegas and Phoenix, respectively, yet these cities rank only 53rd and 57th in terms of their average price appreciation from 2000 to their peak, respectively. Although we focus on house price growth during the boom, it is closely correlated with rapid house price decline during the bust. The correlation between our measure and the

first-lien for-purchase mortgages per quarter between 2003q1 and 2008q4.[5]

Figure 1: House prices and IO shares in cities with elastic and in-elastic housing supply

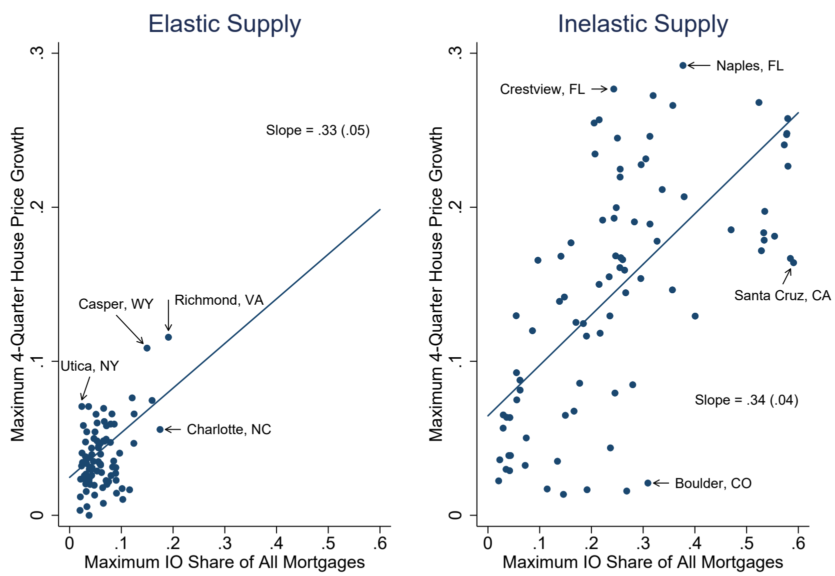

Figure 1 plots our measures of peak house price growth and peak IO use. It is instructive to separate the sample into cities with relatively elastic housing supply and those with relatively inelastic house supply. The left panel includes only cities that rank in the bottom half of all cities according to the share of undevelopable land as compiled by Saiz (2010) and the housing regulation index compiled by Gyourko, Saiz, and Summers (2008). The right panel includes only cities that rank in the top half of all cities according to these two rankings. Unsurprisingly, cities with relatively elastic housing supply exhibit low rates of house price appreciation. What is clear from the figure is that these cities also largely avoided IOs; the share of IOs is below 10% for most of these cities and below 20% for all of these cities. In cities with relatively inelastic housing supply, there is wide variation in both house price appreciation and IO usage, but the two are highly correlated. IOs were used mainly in cities with inelastic housing supply, and then more in cities where house prices surged.

To assess the relationship between the two variables quantitatively, we regressed the maximum 4-quarter house price appreciation on the maximum share of IOs, both with and

largest 4-quarter price decline between when prices peaked and the end of 2010 is .75.

[5] We also considered the share of IOs weighted by loan size, but the results were similar. Even though the average IO loan in Table 1 is larger, IOs were more common in more expensive cities. Within cities, IOs do not appear to be systematically larger or smaller.

Table 2: IOs are correlated with rapid price appreciation

Independent Variable	Dependent variable House price growth (1)	(2)
IO	0.43***	0.45***
	(0.04)	(0.04)
Controls	No	Yes
Adjusted R^2	0.59	0.77
Observations	240	240

Note: OLS regressions of Maximum 4 quarter price appreciation on peak IO share with and without control variables, weighted by number of mortgages. See the text for a description of the variables included in the regressions. Robust standard errors are in parenthesis. The superscripts ***, **, and * denote statistical significance at the .1, 1 and 5 percent levels respectively.

without controls. Our data consists of 240 cities for which we have a full set of controls. We stress that our regression is not meant to be interpreted causally. Instead, it captures whether IO use in a city helps predict house price growth in that city, and hence the extent to which IO use is a distinguishing feature of cities with rapid house price growth. Our results are in Table 2. The first column uses only data on the peak IO share. The coefficient on the IO share is statistically significant at the 0.1% level. The adjusted R^2 for this regression shows that IO usage alone can account for 59% of the variation in house price appreciation. The coefficient of 0.43 on the IO share implies that the maximum 4-quarter house price appreciation in the city with the largest share of IO mortgages (0.59) should exceed house price appreciation in the city with the smallest share (0.016) by 28 percentage points. In the second column, we add various controls to see how well IOs predict residual house price growth. Following previous studies, we include log levels and growth rates of population and per capita income, levels and changes of the unemployment rate and the median property tax rate, and the two housing supply variables mentioned above.[6] Adding the controls improves the overall fit in terms of adjusted R^2 but hardly affects the coefficient on IO share.[7] We will refer to the regression in this second column as our *baseline regression*.

[6]For levels, we average variables from 2003 to the peak-house-price date except property taxes which are the median in 2000. The growth rates and changes are calculated over the 2003 to peak-house-price date period, except for property taxes which is the change between 2000 and the peak-house-price date.

[7]We also added as a control the city-specific house price growth between 1985 and 1989, which has no effect on the coefficient on the peak IO share.

Table 3: Controlling for Other Mortgage Attributes

Independent Variable	Dependent variable: House price growth (1)	(2)	(3)	(4)	(5)	(6)	(7)
IO	0.45***	0.46***	0.38***	0.38***	0.35***	0.45***	0.24*
	(0.04)	(0.04)	(0.08)	(0.06)	(0.08)	(0.04)	(0.11)
Sub prime		-0.14					-0.33
		(0.13)					(0.27)
Securitized			0.08				0.08
			(0.08)				(0.13)
High CLTV				-0.51*			-0.27
				(0.20)			(0.28)
Long term					0.40		0.52*
					(0.21)		(0.22)
Investor						0.44***	0.36**
						(0.13)	(0.12)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.77	0.77	0.77	0.78	0.78	0.79	0.81
Observations	240	240	240	240	240	240	240

Note: OLS regressions of Maximum 4 quarter price appreciation on indicated variables, weighted by number of mortgages. See the text for a description of the variables included in the regressions. Robust standard errors are in parenthesis. The superscripts ***, **, and * denote statistical significance at the .1, 1 and 5 percent levels respectively.

Next, we examine whether it is really the share of IOs that distinguishes cities with rapid house price growth or some other aspect of mortgage markets that might be correlated with IO use. For example, Justiniano, Primiceri, and Tambalotti (2017) argue that when the Fed stopped cutting rates in 2003, the prior boom in refinancing gave way to a boom in private securitization. Although many of these private label securities involved subprime loans, Justiniano et al. (2017) report that a third were IOs. Justiniano, Primiceri, and Tambalotti (2019) show that easing lending can lead to a significant rise in house prices. This suggests the distinguishing feature of cities with rapid house price might be securitization and subprime lending rather than IOs. Alternatively, cities with rapid house price growth might be distinguished not by backloading but a demand for more affordable mortgages, including those with a smaller down payment and a longer maturity. Still another possibility is that investment properties were more common in cities with rapid house price growth, and investors tend to rely more on IOs. To test whether IOs might be proxying for these alternatives, we computed for each city the maximum share of subprime loans, loans that were privately securitized within a year of origination, loans with a maturity of more than 30 years, and loans for investment. We also computed the average share of mortgages with a CLTV in excess of 80%. We then added these variables to our baseline regression.

As is evident from Table 3, including these variables does not diminish the role of a city's IO share in predicting rapid house price growth. The coefficient on IOs remains statistically indistinguishable from our baseline regression in each case. When we control for all the mortgage variables together in the final column, the coefficient on the share of IOs does fall. However, it remains statistically significant at the 5% level. Moreover, since the standard error on the IO share more than doubles when we include all the mortgage variables, we cannot reject the hypothesis that the coefficient is the same as in our baseline regression.

Our key takeaway from Table 3 is that the use of IOs is a distinguishing feature of cities with rapid house price growth distinct from other patterns in mortgages. Once we control for the IO share in a city, some but not all other mortgage attributes are significant. Interestingly, a higher share of subprime mortgages, all else equal, does not imply faster house price growth in that city.[8] Since our analysis compares house price growth *across* cities, this

[8]We find similar results when we use two alternative measures of subprime mortgages, one based on the

finding does not contradict previous work which showed that *within* cities, poorer areas with more subprime borrowers tended to be those with the fastest house price appreciation.[9] More private securitization, longer maturities, and more purchases of investment properties do seem to predict, other things equal, faster house price. But these features cannot explain the popularity of IOs in such cities. House price growth also seems to be lower in cities where borrowers were highly leveraged, i.e. where a larger share of first-liens involved a CLTV of over 80%. One possibility is that lenders sought to protect themselves in markets with rapid house price growth. We will return to this point in Section 8.

3. Theoretical framework

We now develop a theoretical framework with endogenous house prices and mortgage choice to help us understand why IO use might be correlated with rapid house price growth. Our discussion begins by assuming agents have ample wealth so that they have no need to borrow to focus on equilibrium house prices. We then introduce a motive for borrowing and mortgage choice. We use this model in subsequent sections to sort through various explanations for why backloaded mortgages might be concentrated in cities with rapid house price appreciation.

3.1. Setup

Consider a city with a fixed stock of identical houses normalized to 1.[10] We assume the mass of potential residents in the city exceeds the housing stock. Agents who do not reside in the city must reside in a catch-all outside location. All potential residents are infinitely-lived and discount the future at the rate β. They derive utility from a consumption good and from residing in the city. Agents can differ in how they value housing services relative to consumption, and their utility from housing services depends on whether they own or rent, reflecting the benefits of ownership such as the ability to customize one's dwelling.

share of mortgages in subprime mortgage pools and another based on the share of mortgages issued by lenders identified as subprime lenders. See Appendix A.1.1 for more details.

[9] See, for example, Mian and Sufi (2009), Landvoigt, Piazzesi, and Schneider (2015) and Guerrieri, Hartley, and Hurst (2013). Subprime lending also matters in ways that are not directly related to prices. For example, Chambers, Garriga, and Schlagenhauf (2009) and Corbae and Quintin (2016) find that subprime lending played an important role in the rise in home ownership during the boom and foreclosures during the crash.

[10] Assuming a fixed housing stock is convenient but not essential; we could have allowed for some construction as long as capacity constraints would bind eventually.

Regardless of how many houses an agent owns, she can occupy at most one house per period and derives housing services only from that house. As we alluded to above, we begin by assuming agents have deep pockets so they can afford to buy a house without borrowing.

Each agent i's utility is given by

$$\sum_{t=1}^{\infty} \beta^t \left(c_{it} + v_i h_{it}\right), \tag{1}$$

where c_{it} denotes agent i's consumption at date t and h_{it} is an indicator equal to 1 if i resides in the city at date t and zero otherwise. The first thing to note about this specification is that it assumes agents are risk-neutral. We invoke this assumption for its tractability, but comment below on the consequences of allowing for risk-aversion. Second, this specification implies each agent i is willing to trade v_i consumption goods in order to reside in the city for one period. We assume all agents are willing to give up the same amount of consumption goods per period to rent a house, $v_i = \left(\beta^{-1} - 1\right) d$, but differ in how much they value owning the house they occupy. *Low types* value owning the same as renting, i.e. $v_i = \left(\beta^{-1} - 1\right) d$. *High types* value owning their house at $v_i = \left(\beta^{-1} - 1\right) D$ per period, where $D > d > 0$.

If the population of potential residents was fixed, house prices would be constant. In particular, if the mass of high types was less than the stock of houses, all high types would own at least one house so they can occupy it. All of the remaining houses would be occupied by low types, either as owners or renters. The price of a house would be d, the value of buying a house and renting it out indefinitely. If the price were lower, demand for housing would exceed its supply. If the price exceeded d, no agent would own a house not occupied by a high type unless they expected to gain by selling it for an even higher price after discounting. If this were the case house prices would have to grow indefinitely at a rate of at least $\beta^{-1} - 1$, which violates the transversality condition that the value of any asset at date t discounted to the present tends to 0 as $t \to \infty$. If the mass of high types instead exceeded the housing stock, all housing would be owned and occupied by high types. By the same logic, the price of a house would be D. When the population of potential residents is fixed, then, house prices equal the present discounted value of the housing services society obtains from the marginal house. We refer to this valuation as the *fundamental value* of a house.

3.2. House Price Appreciation

For house prices to rise over time, we need the population of high types to rise as well. We will refer to a migration wave, although in principle one could think of new demand for housing from those previously shut out of credit markets who can now buy homes because of an increase in credit supply. Mian and Sufi (2009), Justiniano et al. (2017), and Mian and Sufi (2018) provide evidence of such increases in credit supply starting in 2003.

We assume there are more houses than high types at date 0, but that migration allows the mass of high types to eventually exceed the stock of housing. We model migration as a steady flow of newcomers that continues for a potentially random length of time. Suppose the initial mass of high types at date 0 is $1 - \phi_0 < 1$ so that initially there are fewer high types than houses. A mass $n < \phi_0$ of new potential residents begins to arrive at date 0. As long as a mass of n migrants keeps arriving, there is a constant probability q per period that the migration wave will end and no more migrants will arrive as of the following period. We assume a fraction $1 - \phi$ of each new cohort are high types and ϕ are low types. When $q = 1$, the migration wave lasts for one period, and the mass of high types will fall short of the housing stock with certainty. When $q = 0$, the migration wave will last indefinitely, and the mass of high types will exceed the mass of housing after $\frac{\phi_0}{(1-\phi)n}$ periods with certainty. For $0 < q < 1$, whether the mass of high types will ultimately exceed the stock of houses or not is initially uncertain. With each new cohort of arrivals, the odds that high types will ultimately exceed the housing stock rises given that fewer additional periods of migration are needed for high types to outnumber houses. If new buyers stop arriving before $\frac{\phi_0}{(1-\phi)n}$ periods pass, the mass of high types will forever fall short of the housing stock, and from that point on house prices would equal d. Zeira (1999) previously showed that this type of uncertainty can generate sustained asset price booms (or a crash if arrivals stop), and Burnside, Eichenbaum, and Rebelo (2016) applied this insight to the housing market.

We assume $\frac{\phi_0}{(1-\phi)n}$ is not an integer, and define t^* as the smallest integer exceeding $\frac{\phi_0}{(1-\phi)n}$. By date t^* it will be known whether the mass of high types exceeds the mass of housing or not. Hence, the price of housing from date t^* on is either D if migrants arrive through this date or d if not. Let p_t denote the price of a house for $t < t^*$ if migrants have arrived through date t. By construction, the housing stock still exceeds the number of high types at date t, so some houses must be occupied by low types. Agents can always buy a house at date t, rent it out for $(\beta^{-1} - 1)\, d$ at date $t+1$ and then sell it. In equilibrium this strategy cannot offer strictly positive profits, and so the price at date t must satisfy

$$p_t \geq \beta \left[\left(\beta^{-1} - 1 \right) d + E\left(p_{t+1}\right) \right].$$

Appealing to the same transversality argument as before, it follows that the price p_t cannot exceed the right hand side of this inequality without growing faster than the discount rate indefinitely. Hence, we have

$$p_t = (1-\beta)\, d + \beta \left[qd + (1-q)\, p_{t+1} \right] \tag{2}$$

for $t \leq t^*$. This is a difference equation with terminal condition $p_{t^*} = D$.

When $q = 1$, we have $p_t = d$. That is, if migration stops after one period before the mass of high types exceeds the housing stock, the price will equal the valuation of low types. When $q = 0$, the price at each date t is just the present discounted value of earning rents $(1-\beta)d$ until date t^* and generating housing services $(1-\beta)D$ from t^* on. House prices rise over time, but only because of discounting. In line with this, house price growth p_{t+1}/p_t is at most $1/\beta$, the risk free rate. Intuitively, if house prices grew faster, agents could buy houses, rent them out, and earn more than the risk-free rate with certainty.

In the more interesting case where $0 < q < 1$, up until date t^* house prices rise towards D with each arrival but fall to d if the migration wave stops. These dynamics are illustrated graphically in Figure 2. Intuitively, each arrival means high types are more likely to eventually exceed the stock of housing, which means a higher expected discounted value of the services from the marginal house given it is more likely to eventually serve a high type. As house prices rise, the potential capital loss to owning a house rises, since the price falls to a constant d if the migration wave stops. To ensure people are willing to own housing at each date, there must be a growing potential capital gain to offset this risk. Hence, house prices grow at an increasing rate as migrants keep arriving.

The risk of a collapse in house prices allows for realized house price growth that exceeds

Figure 2: House Prices when $0 < q < 1$

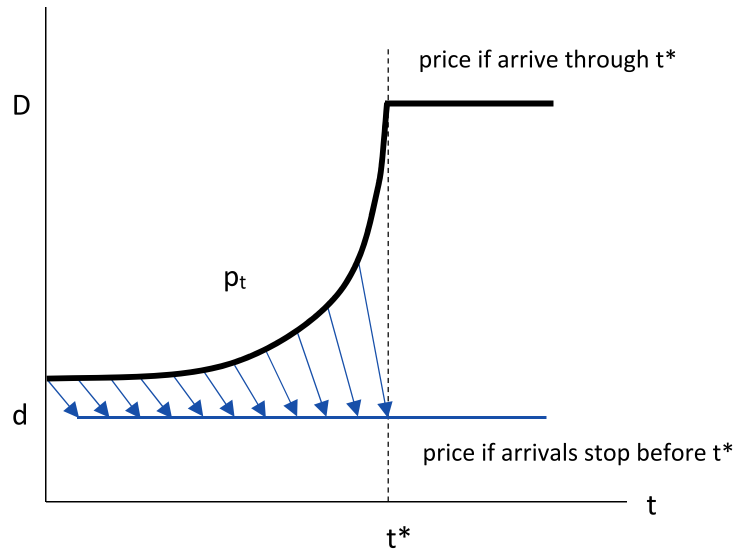

the risk free rate $1/\beta$ if more migrants arrive, in line with the type of house price growth observed in the data. Since the price will at most grow from d to D, house price appreciation can only exceed $1/\beta$ if $D/d > 1/\beta$. The exact condition, proven in Appendix B.1, is:

Proposition 1: For $q > 0$, if

$$\frac{D}{d} > 1 + \frac{1-\beta}{\beta q} \tag{3}$$

then there exists a date $t \leq t^*$ such that if traders arrive through date t, then $p_t/p_{t-1} > 1/\beta$.

As $q \to 0$, it becomes impossible to satisfy condition (3). Only when $0 < q < 1$ and it is uncertain if the mass of high types will exceed the stock of housing can house prices growth exceed $1/\beta$. To relate this to the empirical cross-sectional variation in house price growth, we can view cities with low price growth as those with $q = 1$, i.e., where increased demand for housing was limited or could have been met by additional supply. Cities where house prices grew were those where demand from high types was set to exceed the housing stock.

Our next result shows that even if realized house price growth can exceed the risk-free rate, the *expected* rate of house price appreciation cannot exceed the risk free rate $1/\beta$.

Proposition 2: For $t < t^*$, if traders arrive through date t, then $1 < E\left[p_{t+1}/p_t\right] < 1/\beta$.

Case, Shiller, and Thompson (2013) find that in cities where realized house price growth was far above the risk-free rate in the past, survey respondents expect equally high house price appreciation in the future. This cannot happen in our model, but only because we assume agents are risk-neutral. We could get expected house price growth that exceeds the risk free rate if agents were risk-averse. For example, suppose we replaced the per period utility flow in (1) with $u(c_{it} + v_i h_{it})$, where $u(\cdot)$ is concave. In Appendix B.2, we confirm this can generate expected house price appreciation above $1/\beta$. Intuitively, risk-averse agents require compensation to hold a house when buying a home is risky. When the number of high types falls short of the number of houses, the marginal buyer is necessarily taking on risk. We prefer to work with risk-neutral preferences for analytical tractability.[11]

3.3. Incorporating Mortgage Lending

So far, we have assumed all agents have enough resources to buy a house. We now assume some agents are liquidity constrained and must borrow to buy their homes. They will borrow from those with deep pockets, who may already own homes themselves. The latter would be willing to lend if the interest rate on loans was sufficiently attractive. We begin with a formulation in which agents are indifferent about how their mortgage payments are structured and then show how a small modification leads to a preference for mortgages in which agents repay their debt as quickly as their income allows. In the sections that follow we discuss how this preference for fast repayment can relate to the use of backloaded mortgages.

To simplify things, assume only agents who arrive from date 0 on will need to borrow. Let $\{\omega_{it}^{\tau}\}_{t=\tau}^{\infty}$ denote the income stream of the cohort that arrives at date τ. That is, ω_{it}^{τ} denotes the income that agent i who arrived in the city at date τ earns at date $t \geq \tau$. For now, suppose this income stream is certain. We assume the present discounted value of income for each cohort is such that agents can always afford to buy a house, but that agents do not earn enough to pay for it in full upon arrival. Specifically,

$$\omega_{it}^{\tau} = 0 \text{ for } t = \tau; \tag{4}$$

[11] Assuming agents are risk-neutral also rules out the possibility of shocks to risk-aversion as a source of house price growth. But a shock to risk-aversion would imply either high realized price growth and low expected future price growth or low realized growth and high expected future price growth. By contrast, survey evidence suggests expected future price growth is higher in cities with high realized growth. This suggests price growth is driven by an increase in the amount of risk rather than a change in risk tolerance.

$$\omega_{it}^{\tau} \;\; > \;\; \left(\beta^{-1}-1\right) D \text{ for all } t > \tau. \tag{5}$$

Condition (4) implies an agent cannot offer any resources towards buying a house upon arrival. Condition (5) ensures her subsequent income can cover the required interest payment on a loan in the amount of the price of the house that charges the risk-free rate. It implies

$$\sum_{t=\tau+1}^{\infty} \beta^{t-\tau}\omega_{it}^{\tau} \geq D \tag{6}$$

which means the present discount value of an agent's income evaluated at the risk-free rate always exceeds the price of a house.

For now, we assume lenders have full recourse to seize a borrower's income. We will eventually relax this assumption and consider non-recourse mortgages where lenders can only go after a house but not an agent's income. Full-recourse and condition (6) together imply mortgage lending is risk-free: A lender can always grab the borrower's income and ensure full repayment at the risk-free rate. When loans are risk-free, agents will be indifferent about how mortgages payments are structured: Any principal that the borrower fails to pay down today can be saved at the risk-free rate, and the interest income can be used to exactly cover the borrower's larger obligation. The riskless case is a natural benchmark because house price dynamics are identical to the case where all home buyers have deep pockets. This is because home buyers are indifferent between buying a house outright or borrowing to pay for it. There is no additional cost to borrowing to buy a home, so demand for housing will be the same as if agents were not liquidity constrained.

To break the indeterminacy in mortgage choice, loans must be risky for lenders. We therefore introduce a small probability ε per period that a borrower loses her ability to earn income, earning 0 instead of ω_{it} from this date on. In this event, the borrower will be unable to repay her remaining debt. She could sell the house, but the price may fall short of her obligation. We assume that when the borrower cannot pay back her obligation, the lender recovers only a fraction $\theta < 1$ of the sale price. Lenders can thus incur losses. These losses will be greater if migration stops and house prices fall between when the borrower purchased her house and when she lost her ability to earn. To minimize these agency costs, it is optimal for lenders and borrowers to structure mortgages to ensure the fastest possible repayment path which is achieved by devoting all available income to paying down the principal. This decreases the probability that the agent will be unable to repay their obligation by selling the house and thus reduces the agency costs of borrowing.[12]

The interest rate on the optimal mortgage exceeds the risk-free rate given there are expected losses from default. Since the interest rate exceeds the risk-free rate, low types will strictly prefer renting to buying. If ε is small enough, high types would still want to buy a house immediately upon arrival given the utility they get from ownership. We can formalize this by taking the limit as $\varepsilon \to 0$. In this case, the equlibrium path for house prices converges to the equilibrium price path when there is no risk of default. By taking the limit rather than assuming $\varepsilon = 0$, we select a particular equilibrium in the riskless economy in which only high types buy houses, and use a particular type of mortgage even though absent any risk they are indifferent among mortgages.

The optimal contract features a repayment plan in which households use all of their income each period to pay back their mortgage. While this is far-fetched, we could assume households face exogenous necessary expenditures as in Piskorski and Tchistyi (2011) and Mayer, Piskorski, and Tchistyi (2013) and reinterpret ω_{it} as discretionary income. That is, agents choose the mortgages that use up all of their discretionary income even as they continue to consume while paying off their loans. There is some evidence that borrowers do prefer to repay their mortgages quickly. For example, Dhillon, Shilling, and Sirmans (1990) show that mortgages have shorter terms in regions where incomes are high relative to house prices or high in absolute terms. The fact that backloaded mortgages were rare in the US except in cities with fast house price growth is also consistent with a desire for faster repayment. Thus, while obviously a simplification, a theory of mortgage choice based on a preference for fast repayment seems like a reasonable benchmark.

[12] Piskorski and Tchistyi (2011) study optimal contracts in a setting with several common features to ours, and also argue that the optimal mortgage while house prices grow is the fastest repaying one. Hart and Moore (1994) also describe a model in which borrowers repay as fast as possible, not because that minimizes transaction costs but because a large outstanding obligation would give the borrower bargaining power over the lender and allow them to reduce their obligation.

3.4. Summary

To recap, our framework features two key elements: persistent uncertainty about long-term housing demand, which allows for sustained house price growth in excess of the risk free rate in some states of the world, and a preference to repay mortgages as fast as possible. In what follows, we use this framework to explore the implications of different – but not mutually exclusive – explanations for the popularity of backloading in cities with high price growth.

4. Income growth with liquidity constraints

The first explanation for the popularity of IOs focuses on the role of income growth. In our model, agents want to pay down their mortgages as quickly as possible. Hence, agents whose incomes are set to grow, so ω_{it} rises with t, would seek mortgages with rising payments. They prefer these to equally affordable fixed payment mortgages, since they want to make larger payments as their incomes rise in order to keep their interest charges low.[13]

By this logic, IOs should be more popular in cities where incomes are set to grow more. If cities with high house price growth also experienced higher income growth, this could explain why IOs are concentrated in cities with rising house prices. Although we model income as exogenous, it is not hard to see why cities with house price growth might also feature rising incomes. Consider a city that experiences stochastic productivity growth. In particular, suppose city-level productivity grew over time but with constant probability growth would stop. If labor markets in the city were competitive, wages would grow with productivity. Higher wages would attract increasingly more migration, including by some who value home ownership. This would lead to the kind of house price dynamics we assumed.[14]

We now examine whether IOs were more common in cities with faster income growth, and whether this can explain the concentration of IOs in cities with high price growth. Our approach follows Cocco (2013) who uses individual-level data from the UK to test if

[13] Cocco (2013) argues that households with rising incomes may also buy larger homes and would be unable to afford the initial payments if they had to also pay for principal. Our framework abstracts from this channel since we assume all houses are identical.

[14] Since wage growth is stochastic in this setup, borrowers might have to refinance their mortgages if wage and house price growth stopped earlier than expected. Below we report that borrowers with IOs were somewhat less likely to refinance when house prices fell than traditional mortgage borrowers.

income growth is related to a preference for IOs. He proposes two tests. First, he looks at whether households who took out IOs had faster income growth than those who took out traditional mortgages. If those households correctly expected faster income growth, their incomes should grow faster on average. Second, he uses information on borrowers from the time of purchase, such as their age and occupation, to forecast their future income, and looks at whether households with predictable income growth were more likely to use IOs.

Table 4: The role of income growth

	Dependent variable		
Independent Variable	IO (1)	IO (2)	Price (3)
IO			0.42***
			(0.04)
Expected income growth	1.38***		0.21
	(0.29)		(0.11)
Realized income growth		3.81**	1.12**
		(1.22)	(0.42)
Controls	No	No	Yes
Adjusted R^2	0.17	0.08	0.77
Observations	240	240	240

Note: OLS regressions of Maximum 4 quarter price appreciation or maximum IO share on the indicated variables, weighted by number of mortgages. See the text for a description of the variables included in the regressions. Robust standard errors are in parenthesis. The superscripts ***, **, and * denote statistical significance at the .1, 1 and 5 percent levels respectively.

We do not have individual-level data on income. Furthermore, our focus is on patterns across cities rather than households. We therefore look at whether IOs were more common in cities with higher realized income growth as well as cities whose industry mix in 2005, when IO usage was at or near its peak, predicted higher city income growth. For realized income growth, we use the average annual per capita income growth between 2003 and the year in which house prices peaked. To calculate expected income growth, we regressed average annual per capita income growth for the 2001-2005 period on 2001 employment shares, and used the coefficients from this regression and employment shares in 2005 to predict income growth for each city over the 2005-2009 period.[15]

[15] See Appendix A.3 for the list of industries we use. The adjusted R^2 of the in-sample regression is .77,

The first two columns of Table 4 show how the peak share of IOs in each city is related to these variables. In line with the model, cities with higher actual income growth and higher expected income growth use IOs more heavily. This is also in line with what Cocco (2013) finds for individual-level data on UK borrowers between 2000-2008.[16] The last column of Table 4 indicates whether the popularity of IOs in cities with high realized and expected income growth can account for the popularity of IOs in cities with high price appreciation. That is, we include realized and expected income growth in our baseline regression of house price growth on the share of IOs to see if it can explain the correlation we find. Including these measures does nothing to moderate our finding that a high IO share predicts faster house price appreciation. Indeed, we already included realized income growth in our baseline regression in Table 2. Cities with high IO use and high house price appreciation did not systematically feature high income growth.

5. Demographics and heterogeneity in buyer characteristics

The next explanation we consider for why IOs were concentrated in cities with rapid house price growth concerns demographics. Recall that we earlier observed that if households face necessary expenditures, they would choose mortgages that exhaust their disposable income rather than actual income. Suppose households in two cities had the same income growth, but households in one city were younger and had larger necessary expenditures early on due to young children. Then we would expect to see more backloaded mortgages in the city with younger households. The logic is the same as in the previous section on income growth.

Chiang and Sa-Aadu (2014) argue that higher mobility might also give rise to a preference for backloaded mortgages. Their argument relies on agents being risk-averse, and so does not emerge in our benchmark setup. Still, we can use our framework to offer a sense of their argument. Consider two high type households, one who intends to live in the city forever and the other who will have to leave after one period. Both value owning their house while living in the city, and so both may opt to buy housing. The household that intends to remain

but it predicts income growth in 2005-2009 poorly because of the Great Recession.

[16] Cocco (2013) finds expected income growth does not help to explain IO use in the UK in 1991-2000 when IOs were less regulated and more prevalent. By contrast, we find a role for expected income growth in the US at a time when IOs were prevalent and regulation was relatively lax.

in the city will not care about the possibility that house prices will collapse, since it will not sell its house. By contrast, the household that will move will care about house prices. Repaying principal is equivalent to investing in housing, which leaves the household more exposed to house price risk. Under some conditions, the mover will prefer to minimize her exposure to this risk by not repaying principal.

Table 5: The role of heterogenous buyer characteristics

Independent Variable	IO (1)	IO (2)	IO (3)	Price (4)
	Dependent variable			
IO				0.26***
				(0.05)
Median age	-0.01		-0.00	0.00
	(0.01)		(0.00)	(0.00)
Gross migration		2.97*	-0.54	0.84**
		(1.24)	(0.44)	(0.30)
Controls	No	No	No	Yes
State fixed effects	No	No	Yes	Yes
Adjusted R^2	0.00	0.06	0.84	0.94
Observations	240	240	240	240

Note: OLS regressions of Maximum 4 quarter price appreciation or maximum IO share on the indicated variables, weighted by number of mortgages. See the text for a description of the variables included in the regressions. Robust standard errors are in parenthesis. The superscripts ***, **, and * denote statistical significance at the .1, 1 and 5 percent levels respectively.

These considerations suggest that IOs should be more popular in cities where households are younger and more mobile. Moreover, if cities with younger and more mobile populations are associated with faster house price appreciation, this logic could explain why IOs were concentrated in cities with rising house prices. To explore this possibility, we compiled data on the median age for each city from the 2000 Census and on each city's average gross migration rate between 2000 and 2005. The first two columns of Table 5 show how the maximum share of IOs in each city is related to these measures. Cities whose residents are younger are more likely to use IOs, although the coefficient is not statistically significant. IOs were also more common in cities with higher turnover rates, a pattern that is statistically significant at the 5% level.

Of course, there can be a host of other demographic differences across cities that will be related to mortgage choice, e.g. patience and risk-aversion. While it is impossible to control for all of these factors, to the extent that populations are more homogeneous within states than across states, we can try to control for these differences using state fixed-effects. This is what we do in column (3) of Table 5. The adjusted R^2 rises to 84% in this case, suggesting differences in IO usage are largely between rather than within states. Of course, this controls for any differences across states beyond just demographics. Adding state fixed effects eliminates a role for gross mobility. For cities within the same state, the coefficient on gross migration turns negative and statistically insignificant. This suggests gross migration is correlated with some other factor associated with IO usage that varies across states.

In the last column of Table 5 we examine whether the popularity of IOs in cities with low median age, high gross migration, or in particular states explain why IO usage is a good indicator of house price appreciation. Controlling for these variables, particularly state fixed effects, does reduce the coefficient on the share of IOs relative to the baseline regression in Table 2. But the coefficient remains highly statistically significant. When we restrict attention to variation among cities within the same state, then, a high IO share remains a strong indicator of rapid price appreciation even after we control for that city's median age and gross migration rate. Neither age, mobility, nor time invariant state-level variation can explain why IOs were so popular in cities with rapid house price growth.

6. Affordability

The explanations we have considered so far focused on features that make IOs appealing to certain borrowers but which are not directly related to house prices. And, indeed, we found that while these explanation can account for why IOs are more popular in some cities than others, they cannot explain why IOs were more popular in cities with rapid house price growth. We now turn to explanations that relate directly to house prices. The first is that as housing became more expensive, borrowers were drawn to IOs for affordability reasons.

To illustrate this argument, suppose income were constant across all cohorts and over time, i.e., $\omega_{it}^{\tau} = \omega$ for all i, τ and $t > \tau$. Since we assume agents have no resources when they arrive, they must borrow the full value of the house upon their arrival, p_{τ}. The fastest repayment path would oblige borrowers to pay ω until they pay off their loan. That is, with constant incomes, mortgages will feature fixed payments for a set maturity. To determine this maturity, note that as $\varepsilon \to 0$, the interest rate tends to $1/\beta$. Using the familiar formula for the payment on a fixed-payment mortgage of size p_τ, maturity T, and interest rate $1/\beta$, the maturity T consistent with a payment of ω must solve

$$\frac{(\beta^{-1} - 1)}{1 - \beta^T} p_\tau = \omega.$$

If p_τ rises relative to ω, borrowers would be forced to take out longer maturity loans given it takes more time for the same fixed payment to pay off a larger loan at the same interest rate. This logic implies borrowers should extend the maturity of their loans rather than shift towards backloaded mortgages. Indeed, Table 3 shows that mortgages with maturities of over 30 years were more common in cities with rapid house price growth. However, there may be a limit on how much lenders are willing to extend the maturity on their loans. If that were the case, households could still approximate a long maturity mortgage by repeatedly refinancing a backloaded mortgage. Although such mortgages offer only a temporary reprieve from affordability concerns, this relief can be extended through refinancing. IOs may have been the best available contract that constrained borrowers had access to.

To the extent that affordability generates a preference for IOs, we should observe that IOs were more popular in cities where house prices were high relative to income. This is not the same as our finding that IOs were common in cities with rapid house price growth. Although rapid house price growth will be associated with higher house prices, initial house prices in these cities may have been low. In addition, incomes could grow alongside house prices, so cities with rapid house price growth need not be those where housing is least affordable. Indeed, our findings below will suggest they were not.

To explore the role of affordability, we compiled Census Bureau data on median house prices in 2000 by city. We then used our house price index to arrive at a level of house prices for each city and each year in our sample. We divided this price by the average nominal income per capita, and took the maximum value of this ratio between 2003 and 2008 to arrive at a peak ratio of price to income for each city. The first column in Table 6 shows

Table 6: The role of affordability

Independent Variable	Dependent variable IO (1)	Price (2)
IO		0.40***
		(0.10)
Peak house price to income ratio	0.06***	0.00
	(0.00)	(0.01)
Controls	No	Yes
Adjusted R^2	0.79	0.77
Observations	240	240

Note: OLS regressions of Maximum 4 quarter price appreciation or maximum IO share on the indicated variables, weighted by number of mortgages. See the text for a description of the variables included in the regressions. Robust standard errors are in parenthesis. The superscripts ***, **, and * denote statistical significance at the .1, 1 and 5 percent levels respectively.

that the maximum share of IOs in each city is highly related to this affordability measure. This one variable alone accounts for almost 80% of the cross-city variation in IO use. There is a clear preference for IOs in cities where house prices are high relative to average income per capita.[17] At the same time, the second column in Table 6 shows that including this affordability measure when we regress house price growth on the maximum share of IOs has little impact on the extent to which the IO share predicts the rate of house price growth. Cities with rapid house price growth and high IO use were not necessarily cities where house prices were particularly high relative to income.

Our model suggests another way to explore affordability. Since households in our model want to pay down their debt as quickly as possible, they should be reluctant to use backloaded mortgages until prices rise. That is, house price growth should lead the use of IOs. Since we do not require the same set of city-level covariates to explore this issue, we expand the sample to all 376 cities for which we have both house price and mortgage data and use data for 2000q1 through 2006q4 to focus on the period of rising house prices.

Figure 3 shows the correlation between the change in the log IO share of mortgages in quarter $t+j$ with the change in the log real house prices at quarter t. For the 35 cities with

[17] Bäckman and Lutz (2017) argue that affordability also played an essential role in explaining the popularity of IOs in Denmark, although their analysis does not use regional variation as we do.

Figure 3: Growth of IO share leads house price growth

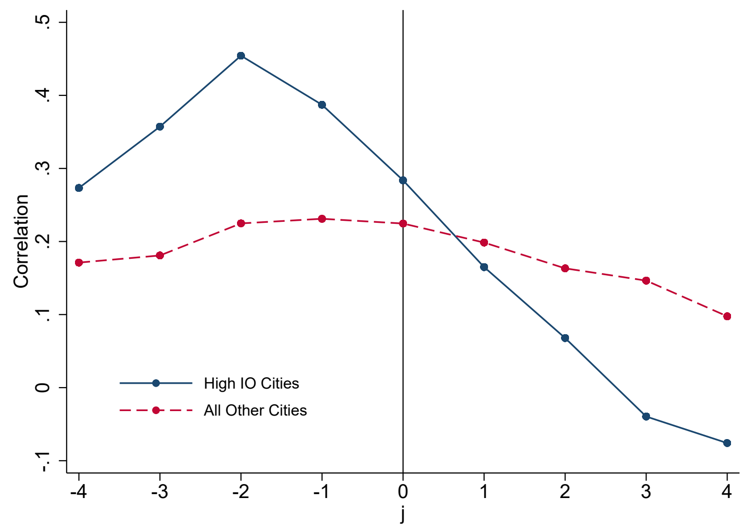

Note: Figure displays cross correlations of change in log IO share at quarter $t + j$ with change in log real house price at quarter t. High IO cities are those cities with a maximum share of IOs in excess of 1/3.

the highest IO share (with a peak IO share of at least 33%), IO growth clearly leads house price growth. In the remaining cities, IO growth only slightly leads price growth.

Table 7: House prices do not predict IOs in high IO cities

	All Cities (1)	High IO Cities (2)
Sum of Δio coefficients	0.58***	0.61***
	(0.03)	(0.06)
Sum of Δhp coefficients	0.09*	-0.11
	(0.04)	(0.10)
F-statistic for exclusion of Δhp coefficients (p-value)	3.43	1.83
	(0.02)	(0.14)
Adjusted R^2	0.40	0.46
Lags	3	3
Number of cities	376	35
Number of quarters	24	24

Note: Panel regressions of log changes in IO share (Δio) on its own lags and lags of log changes in house prices (Δhp.) The number of lags is the lowest value such that we could not reject the hypothesis of no-autocorrelation in the residuals up to that value using the Arellano and Bond (1991) test. Standard errors are robust to correlation in the residuals across cities. The superscripts ***, **, and * denote statistical significance at the .1, 1 and 5 percent levels respectively.

Table 7 reports Granger causality regressions of log house price growth, Δhp, on log IO share growth, Δio. The first column reports the results for all 376 cities. Here, we find that house price growth Granger causes growth in IOs. The effect is statistically significant but small. The second column reports results for the 35 cities with the highest peak IO shares. For these cities we cannot reject the hypothesis of no Granger causality, and the point estimate is negative. Once again, although we find evidence consistent with the affordability hypothesis in the broad cross section, the cities with the highest IO share, which also feature high price growth, do not conform with the affordability hypothesis. In those cities, IO use seems to pre-date house price growth.[18]

7. Expected house price appreciation

The previous section examined whether past house price growth led households to use IOs as houses became less affordable. However, recall that survey evidence suggests that in cities with past house price growth, expected future house price growth was also higher. Higher expected house price growth may separately lead households to prefer IOs.

The argument for why expected future house price growth might encourage IOs is laid out in LaCour-Little and Yang (2010) and Brueckner et al. (2016). These papers assume borrowers want to postpone repaying their debt, in contrast to our setup. We can capture this scenario in our framework as well by assuming that some high types are impatient, i.e. their discount rate is below β. These impatient agents would want to delay their repayment and enjoy earlier consumption. Lenders would charge them higher interest rates to cover the risk that the borrower may lose her earning ability before repaying her obligation, which will limit the extent to which these impatient agents will choose to backload. The two papers argue that all else equal, higher expected house price growth should mitigate default concerns and allow borrowers to pay less for the privilege of backloading.

A problem with this logic is that when prices are endogenous, something must change to allow for higher expected house price growth. Recall that in our model, realized house

[18] Dokko et al. (2019) also compare the timing of house price growth and mortgage choice. They look for break-points in house price growth to identify the start of the boom in house prices for each city, and find that different types of alternative mortgage products rose after this date. However, for IOs their data shows an increase that predates the boom, in line with what we find.

price growth can be high only when there is also a risk that house prices can fall. While a rise in house prices will mitigate losses from default, the fall in housing prices if it should occur would exacerbate them. The possibility of house price growth in the good state of the world may lead lenders to be more willing to offer backloaded mortgages. But lenders must benefit sufficiently when house prices rise to cover their expected losses if house prices fall.

In principle, lenders can charge a higher interest rate on backloaded loans than traditional loans, although the rate might be lower than what they would charge when house prices are constant. This creates a situation in which borrowers have an incentive to refinance their loan if house prices rise. Indeed, our model suggests agents should be able to do this. House prices rise only if migration continues, and each arrival lowers the odds that house prices will fall. Since this makes lending less risky, new lenders would be willing to refinance existing loans at a lower rate. This would destroy the incentive to make such loans in the first place. However, as Gorton (2008) and Mayer et al. (2013) discuss, these incentives can be overcome by stipulating a prepayment penalty on borrowers that will either prevent them from refinancing or allow lenders to collect fee income if borrowers do refinance. Gorton (2008) argues this was precisely the business model lenders were using.

Is there any evidence that backloading was tied to prepayment penalties in line with this explanation? Table 1 suggests that IOs were somewhat more likely to feature prepayment penatlies than all mortgages, although Option-ARMs were dramatically more likely to feature them. This is consistent with our hypothesis, since impatient households who want to backload may prefer an Option-ARM that can be even more backloaded than an IO. To the extent that these loans compensated lenders not with high interest rates but from fee income, in line with evidence that backloaded mortgages did not command much higher interest rates, we should see that backloaded mortgages with prepayment penalties were more likely to be refinanced when house prices grew than other mortgages, despite the fact that they carry a penalty. Since the Black Knight mortgage data we use only records whether a loan was paid off early and not whether it was refinanced, we merged our data with recorder of deeds data on individual properties from DataQuick to distinguish between sales and refinancing. We discuss the DataQuick database and our matching algorithm in Appendix A.1.2. Since Option-ARMs involve adjustable rates, as did most of the IOs in our sample, the most natural comparison group for these loans is non-backloaded ARMs. The incentives to refinance fixed rate mortgages are different.

We estimate the propensity of holders of first-lien IO, Option-ARM and non-backloaded ARM mortgages originated for purchase in 2005 and 2006 to refinance by the end of 2010. In estimating the propensity to refinance, we follow Elul, Souleles, Chomsisengphet, Glennon, and Hunt (2010) in using a linear probability model with a 6^{th} order polynomial in the number of months since origination as a substitute for a proportional hazard model. In each case the dependent variable is a dummy equal to 1 if a mortgage refinances in that quarter and 0 otherwise. A mortgage is included in our sample until it is repaid or the borrower defaults. We include a large number of mortgage characteristics as controls (the full list of controls is described in Appendix A.1.1) as well as city-quarter dummies that allow for cyclical patterns to differ across cities. Given these controls, in any given quarter we are comparing mortgages with similar characteristics that originated in the same city at different dates and thus face different house appreciation, after controlling for the typical tendency to refinance as a function of time since origination. We are interested in the impact of the cumulative price change since origination and its interaction with mortgage type. We define Δ^+ House Price $\equiv \max(\Delta\text{House Price}, 0)$ as the change in house prices if house prices rose relative to the origination year and Δ^- House Price $\equiv -\min(\Delta\text{House Price}, 0)$ as the change in house prices if house prices fell relative to the origination year. The interaction between price changes and mortgage type tells us whether a particular type of mortgage was more likely to refinance when house prices appreciated or depreciated.

Columns (1) and (4) in Table 8 show that all of the borrowers in our sample were more likely to refinance when house prices rose, since the coefficients on Δ^+ House Price are positive. They were also less likely to refinance when house prices fell given the coefficients on Δ^- House Price are negative. Furthermore, the coefficients on the interaction terms suggests that the propensity to refinance was more responsive to house prices among borrowers with IOs, and even more responsive among borrowers with Option-ARMs. However, the greater price sensitivity of refinancing among borrowers with Option-ARMs appears to be related to the fact that these mortgages disproportionately imposed prepayment penalties (PPP). When we compare IOs and Option-ARMs that either both stipulate PPPs, columns (2)

Table 8: Propensity of backloaded mortgages to refinance

	Mortgage Type					
	IO			Option-ARM		
Independent Variable	All (1)	PPP (2)	No PPP (3)	All (4)	PPP (5)	No PPP (6)
Type	-0.003***	-0.017***	-0.001	-0.003***	-0.007***	0.004***
	(0.000)	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)
Δ^+ House Price	0.065***	0.044***	0.069***	0.029***	0.023***	0.043***
	(0.003)	(0.004)	(0.003)	(0.004)	(0.004)	(0.004)
Δ^- House Price	-0.032***	-0.029***	-0.027***	-0.010	-0.006	-0.019**
	(0.005)	(0.006)	(0.005)	(0.006)	(0.006)	(0.006)
Type $\times$ Δ^+ House Price	0.020***	0.180***	-0.015***	0.072***	0.119***	-0.014**
	(0.003)	(0.005)	(0.003)	(0.003)	(0.004)	(0.005)
Type $\times$ Δ^- House Price	-0.007*	0.013**	-0.012***	-0.016***	-0.013***	-0.019***
	(0.003)	(0.004)	(0.003)	(0.003)	(0.004)	(0.004)
Adjusted R^2	0.02	0.02	0.01	0.02	0.02	0.01
Observations	1,321,531	722,083	1,175,742	1,025,361	863,907	737,748
Mortgages	102,642	59,633	89,238	84,201	70,557	59.873

Note: Linear probability models of the propensity of holders of first lien mortgages originated for purchase in 2005 and 2006 to subsequently refinance by quarter before the end of 2010. The OLS regressions include controls for mortgage characteristics (see Appendix A.1.1), a 6^{th} order polynomial in months since origination and city-quarter dummies. 'Δ^+ House Price' denotes the maximum of zero and the cumulative price change since origination. 'Δ^- House Price' denotes the absolute value of the minimum of zero and the cumulative price change since origination. 'Type' is a dummy indicating whether the mortgage is of the type indicated in the column header, where 'PPP' or 'no PPP' indicates whether or not the mortgage has a pre-payment penalty. In addition to IOs (Option-ARMs) the regressions include all ARM originations excluding Option-ARMs (IOs). Estimates are based on a random 15% sample of our merged Black-Knight/DataQuick dataset. The superscripts ***, **, and * denote statistical significance at the .1, 1 and 5 percent levels respectively.

and (5), or both omit these penalties, columns (3) and (6), the two mortgages look similar. Refinancing is most responsive to house price growth among borrowers with backloaded mortgages with PPPs, followed by borrowers with non-backloaded mortgages, and is least responsive among borrowers with backloaded mortgages without PPPs. Thus, refinancing behavior is consistent with the notion that backloaded mortgages were structured in a way that allowed backloading in exchange for rewarding lenders with fee income when house prices grow.[19]

Can the ease of backloading in places where house prices were expected to grow explain the popularity of IOs in cities with high price apprecaition? Even though IOs did not feature prepayment penalties as often as Option-ARMs, it could be that IOs with prepayment penalties were precisely those that prevailed in cities with rapid house price growth. If this were true, controlling for the share of mortgages with prepayment penalties or for indicators of the presence of impatient borrowers, such as the share of Option-ARMs, would reduce the ability of the share of IOs in a city to predict house price growth. Columns (1) and (2) of Table 9 show that IOs were indeed more popular in cities where prepayment penalties were more common as well in cities with a preference towards even more backloaded mortgages that allow for negative equity. However the last column of Table 9 shows that including these variables has virtually no impact on how IO use predicts house price growth. IOs were not more popular in cities with rapid house price growth because lenders in those cities were more willing to stipulate prepayment penalties or highly backloaded loans.

8. Speculative home buyers

On their own, each of the explanations for the popularity of IOs we have considered so far can help account for some of the variation in IO use across cities but not for the popularity of IOs in cities with rapid house price growth. We now confirm this remains true when we consider these explanations jointly. The first two columns of Table 10 regress the peak IO share in each city on all of the covariates we have considered so far, both with and without

[19] Note that not all borrowers who faced a prepayment penalty would have incurred these penalties when they refinanced. In practice, prepayment penalties were typically in place for only a fixed period, and some of those who refinanced waited until after the penalty expired to refinance.

Table 9: The role of pre-payment penalties and option-ARMs

Independent Variable	Dependent variable IO (1)	IO (2)	Price (3)
IO			0.45***
			(0.07)
PPP	1.22***		-0.01
	(0.17)		(0.15)
Option-ARM		1.75***	-0.00
		(0.16)	(0.24)
Controls	No	No	Yes
Adjusted R^2	0.47	0.70	0.77
Observations	240	240	240

Note: OLS regressions of Maximum 4 quarter price appreciation or maximum IO share on the indicated variables, weighted by number of mortgages. 'PPP' and 'Option ARM' denote maximum shares of mortgages with these attributes. See the text for a description of the variables included in the regressions. Robust standard errors are in parenthesis. The superscripts ***, **, and * denote statistical significance at the .1, 1 and 5 percent levels respectively.

state fixed effects. Even without state fixed effects, the adjusted R^2 is 84%. Adding state fixed effects increases the adjusted R^2 to 94%. These variables can thus explain nearly all of the variation in IOs across cities. But the last two columns of Table 10 show that when we include both the predicted IO share from these regressions as well as the residual IO share, both are significant in predicting the maximum 4-quarter house price growth. That is, the residual variation in IOs not spanned by the factors we control for remains strongly and significantly predictive of rapid house price growth. We need another factor to explain why IOs were so popular in these cities.[20]

The explanations we have considered so far assume agents buy housing for the services it provides and rely on backloaded mortgages because of their cash flow, impatience or likelihood of moving. Lenders are willing to accommodate this preference for backloading, so backloaded contracts offer a gain from trade. In line with this, all mortgage transactions are carried out by high types. Low types do not buy housing, since borrowing to buy housing they could equally rent is costly. We now consider an explanation for IO use in which low

[20] We also added state fixed effects to the last two columns in Table 10. The coefficients on residual IO share remain highly significant.

Table 10: Comparing Predicted and Residual IO Share

Independent Variable	IO (1)	IO (2)	Price (3)	Price (4)
	Dependent variable			
Predicted IO share			0.46***	0.47***
			(0.04)	(0.04)
Residual IO share			0.39**	0.29**
			(0.12)	(0.09)
Expected income growth	0.49*	0.08		
	(0.22)	(0.11)		
Realized income growth	0.12	0.87*		
	(0.60)	(0.39)		
Median age	-0.00	-0.00		
	(0.00)	(0.00)		
Gross migration	0.43	-0.73*		
	(0.44)	(0.37)		
Peak house price to income ratio	0.04***	0.04***		
	(0.01)	(0.00)		
PPP	-0.01	-0.08		
	(0.18)	(0.15)		
Option ARM	0.46	0.44**		
	(0.26)	(0.16)		
Controls	No	No	Yes	Yes
State fixed effects	No	Yes	No	No
Adjusted R^2	0.84	0.94	0.77	0.77
Observations	240	240	240	240

Note: OLS regressions of Maximum 4 quarter price appreciation or maximum IO share on the indicated variables, weighted by number of mortgages. 'Predicted IO share' and 'Residual IO share' in columns (3) and (4) correspond to the fitted values and residuals from the IO regressions in columns (1) and (2), respectively. See the text for a description of the variables included in the regressions. Robust standard errors are in parenthesis. The superscripts ***, **, and * denote statistical significance at the .1, 1 and 5 percent levels respectively.

types also demand housing, not for its services but as a way to gamble on housing at the expense of creditors. This motive emerges once we modify our model to allow for non-recourse mortgages, i.e. loans in which the lender cannot seize the borrower's income in case of default. Several US states restrict mortgages to be non-recourse, and recourse is not always exercised even in states that allow it.

Limits on recourse induce low types to buy housing by making strategic default profitable. If lenders have recourse, assumption (6) makes strategic default pointless, since lenders can recover their obligation fully from the borrower's income. But if lenders have no recourse, low types can buy housing and gamble at their lender's expense: They buy a house hoping more migrants come and prices rise, but will default if migrants fail to arrive.[21] We assume lenders cannot tell the two types apart. Although lenders will want to avoid low types, if the latter have the same income and buy the same housing as high types, they will look indistinguishable to lenders. If lenders cannot offer contracts that screen out low types, e.g. contracts that tie payments to house prices (which may be costly to enforce) or pay speculators to walk away (which may attract non-speculators), they will try to minimize their expected losses from speculators. IOs can serve that purpose if they encourage speculators to sell their houses earlier than they would otherwise. IOs no longer serve to accommodate the needs of borrowers, but are a sign of lenders responding to ongoing speculation.

We defer the formal analysis of non-recourse mortgages to Appendix B.3. Here, we just sketch the equilibrium that emerges. During the migration wave, both high and low types borrow to buy houses. High types buy because they value housing. Low types buy to speculate. As in Allen and Gale (2000), demand by speculators gives rise to a *bubble* in housing, meaning house prices exceed the fundamental value of a house, defined in Section 3.1 as the expected present value of the housing services from the marginal house. Intuitively, low types value houses not because of the services they yield, but because of the profits they can earn if migrants come. This means they value houses only for their upside potential, and would buy a house even if its price exceeded the expected value of housing services it could provide. The option value of default drives a wedge between price and fundamental value.

In addition, in equilibrium lenders will offer two types of contracts during the migration wave. The contract that low types take is designed to encourage them to sell the asset early. The intuition for this, as observed in Barlevy (2014), is that since houses are overvalued, lenders and borrowers would be jointly better off selling the house rather than holding on to it. One way to accomplish this is to encourage the borrower to repay his loan quickly, e.g.

[21] High types may also have an incentive to default when house prices fall to buy a house at a lower price. In practice default restricts access to credit, and so they may not be able to buy another house once they default. We proceed as if high types never default.

via a balloon payment. Since the balloon is more constraining than a traditional mortgage, it must offer other features to draw in speculators. One such draw is low initial payments. These low initial payments effectively act as a bribe to induce speculators to sell quickly. In practice, anti-predatory lending laws likely restricted the use of balloon payments, which are quite rare in the data.[22] IOs may have served as a substitute by encouraging borrowers to sell the house before their payment reset even without an explicit balloon payment while still rewarding borrowers with low payments initially. Note that in this explanation, IOs are a response to, rather than a cause of, speculation. If we restricted lenders to only offer traditional mortgages, low types would still have an incentive to speculate. Since high types still prefer to repay their debt as fast as possible, lenders will also offer them mortgages with payments that match their income. In equilibrium, low types are indifferent between the two mortgages and would use the faster repayment mortgage if that were the only one offered. Since lenders cannot avoid low types and prefer that they use IOs, they will offer both types of mortgages and not just those aimed at high types.

In contrast to the previous explanations, where lenders were willing to offer the backloaded mortgages that borrowers wanted, here backloaded contracts are associated with a trade that is not mutually beneficial. Lenders would rather not lend to speculators, even if they prefer giving them backloaded mortgages to traditional ones. In fact, these trades make society worse off, since borrowers fail to internalize the default costs they incur on lenders, who recall only recover a fraction θ of the value of the house they lend against. An intervention that reduces the amount agents borrow to speculate can improve welfare.[23]

Unfortunately, there is no measure of speculation that we can similarly use to see if it can account for the popularity of IOs in cities with high price growth. However, our model offers other testable implications we can look for to see if there is any evidence consistent with speculation. First, it implies we should observe more IOs in cities with high price growth in non-recourse states than recourse states. It also implies that lenders offer IOs to encourage speculators to sell houses faster than they would with traditional

[22] For a discussion of anti-predatory lending laws at the state and federal level, see Bostic, Chomsisengphet, Engel, McCoy, Pennington-Cross, and Wachter (2012).

[23] As Allen, Barlevy, and Gale (2019) point out, agents would be better off if lenders purchased the house directly and gave speculators a transfer than lending to the speculator and potentially incur default cost.

mortgages. We cannot observe how borrowers who use IOs would have behaved if they took out traditional mortgages. However, in the model those who choose traditional mortgages are not speculators, and should be less likely to sell. Thus, we should see that borrowers with IOs are more likely to sell their homes than borrowers with traditional mortgages when house prices rise. By the same logic, when house prices fall, we should see that borrowers with IOs are more likely to default than borrowers with traditional mortgages. We now investigate these implications of speculation in our model.

Figure 4 reproduces the scatter plot of maximum IO share and maximum 4-quarter price appreciation for cities with inelastic housing supply from Figure 1 but now distinguishes between cities in states with recourse mortgages (open circles) and those with non-recourse (solid circles), using the classification in Ghent and Kudlyak (2011).[24] As is clear from the figure, the cities with the highest IO shares are those in states without recourse. More generally, among cities with a sufficiently high rate of house price appreciation, the maximum IO share is considerably higher in non-recourse states than in recourse states. This is consistent with the view that non-recourse encourages the use of IOs.

Figure 4: Impact of recourse in cities with in-elastic housing supply

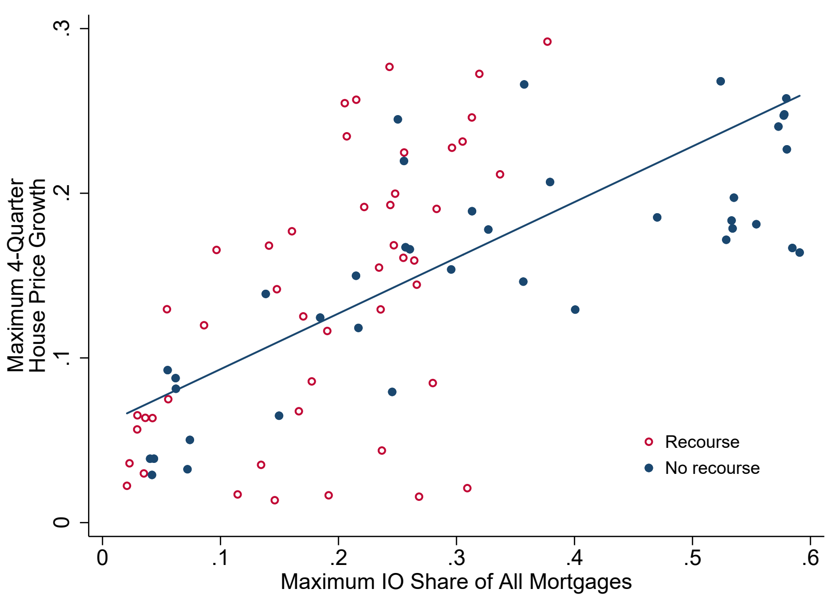

[24]As Ghent and Kudlyak (2011) point out, California requires that the original loan against a property be non-recourse, but does not require a refinanced mortgage to be non-recourse. Since we focus on for-purchase mortgages, we follow them and classify California loans in our sample as non-recourse.

Next, we consider our model's predictions for the propensity of borrowers with IO mortgages to sell and default. To do so we exploit the same individual mortgage data, sample period, and regression framework we used to examine refinancing behavior in Table 8. Since our predictions concern the behavior of borrowers with IOs as compared to borrowers with traditional mortgages, we use all non-backloaded mortgages as our comparison group. Table 11 shows our results. We find that the propensity of a borrower to sell his house in response to a price increase is twice as large for borrowers with IOs than for borrowers with traditional mortgages. Likewise, the propensity of a borrower to default in response to a price fall is twice as large for borrowers with IOs.

Our findings for default rates are consistent with results in LaCour-Little and Yang (2010), Brueckner et al. (2016), and Amromin et al. (2018). However none of these papers considered differences in the propensity of these mortgages to sell in response to a house price increase. Interestingly, Cocco (2013) finds that households who took out IOs in the UK were less likely to move than other type of borrowers, even between 1991 and 2000 when IOs were relatively unregulated. By contrast, in the US these borrowers seem less attached to their homes, at least when house prices rise.[25]

The evidence we have presented is thus consistent with the possibility that it is the use of IOs by speculators that might account for their residual popularity in cities with rapid house price growth. While it is not surprising that borrowers in non-recourse states would be willing to take out such loans in order to gamble on house prices, the question is why lenders would be willing to extend such loans given the risk it exposes them to. Our model suggests lenders might have been willing to go along because they wanted to encourage borrowers to sell more quickly. Is there any support for this view? Although lenders may not have realized that house prices were set to fall broadly and significantly, recall that IOs were largely confined to certain geographic areas. And in those areas, there was certainly concern about the possibility of a bubble. For example, Fed Chairman Alan Greenspan testified on

[25]For robustness, we considered restricting our comparison group to only non-backloaded ARMs. Borrowers with IOs remain more likely to default in response to a fall in house prices. However the propensity to sell in response to a house price increase is sensitive to including hybrid ARMs known as 2/28s and 3/27s in which the interest rate is fixed for 2 or 3 years before the first reset. These mortgages should have appealed to borrowers who expected to move within 2 to 3 years and who would have otherwise taken out fixed rate loans. When we excluded these hybrid loans, IOs were more likely to sell in response to a rise in house prices.

Table 11: Propensity of IOs to Sell or Default

Independent Variable	Dependent Variable	
	Sale	Default
IO	0.000*	0.005***
	(0.000)	(0.000)
Δ^+ House Price	0.018***	-0.013***
	(0.001)	(0.001)
Δ^- House Price	-0.007***	0.068***
	(0.001)	(0.002)
IO $\times$ Δ^+ House Price	0.019***	-0.012***
	(0.001)	(0.002)
IO $\times$ Δ^- House Price	-0.002	0.072***
	(0.001)	(0.001)
Adjusted R^2	0.01	0.03
Observations	3,888,597	3,888,597
Mortgages	262,718	262,718

Note: Linear probability models of propensity of holders of first lien mortgages originated for purchase in 2005 and 2006 to subsequently sell or default before the end of 2010. Other than the type of transaction being considered the regressions are formulated as in Table 8. See that table for definitions of the regressors. The regressions include all mortgages except for option-ARMs. Estimates are based on a random 15% sample of our merged Black-Knight/DataQuick dataset. The superscripts ***, **, and * denote statistical significance at the .1, 1 and 5 percent levels respectively.

June 9, 2005 that "although a 'bubble' in home prices for the nation as a whole does not appear likely, there do appear to be, at a minimum, signs of froth in some local markets where home prices seem to have risen to unsustainable levels."[26] Even before the crisis, IOs attracted attention as a vehicle for speculators. For example, Williams (2006) reports that some of the borrowers attracted to alternative mortgage products included those who bought homes for investment purposes. There are other signs that lenders in hot markets were wary. Recall that in Table 3 we found cities where house prices surged had less rather than more leverage. Ono, Uchida, Udell, and Uesugi (2013) similarly find that when land prices surged in Japan, down payment requirements for business loans increased. This suggests lenders in boom markets took steps to protect themselves against default risk.

Interestingly, Minsky (1982) argued that speculative booms seem to be associated with

[26] Speech available online at http://www.federalreserve.gov/boarddocs/testimony/2005/200506092/default.htm.

loans in which borrowers need only pay the interest on what they borrow. His interpretation of this pattern is somewhat different from ours. Rather than use these loans to encourage speculators to sell, he argued lenders get swept up in the euphoria surrounding speculative manias and speculate themselves. Our data does not allow us to divine the motivations of lenders. But according to both views, these contracts are a sign of speculation, and some of the policy implications we discuss in the Conclusion would be the same under either view.

9. Conclusion

We have gone beyond the well known fact that the boom in US housing markets during the mid 2000s was associated with a rise in the use of nontraditional mortgages to show that the use of interest-only mortgages in particular was a distinguishing feature of US cities where house prices surged. When we sort through the various explanations that have been advanced for the popularity of these mortgages, we find that they can indeed account for a large fraction of the variation in IOs across cities, but not for their popularity in cities with rapid house price growth. We then argue that if lenders have limited recourse, IOs may arise because of speculation. In line with this explanation, we find that rapid house price growth was associated with more IO use in non-recourse states than in recourse states. We also find that borrowers who took out IOs behaved in ways that are consistent with speculation. This suggests the popularity of IOs in cities with rapid house price growth may be indicative of speculative demand for housing in those cities.

To the extent that the popularity of IOs in cities with rapid house price growth is evidence of speculation in those cities, there may be a role for policymakers to intervene in order to discourage speculative borrowing in those cities that involves default costs that speculators likely fail to internalize. However, even if such an intervention is desirable, it need not translate into the type of restrictions on IOs some policymakers advocated, and certainly not a blanket restriction on the use of such loans.

In our model, banning IOs would not curb speculation. Recall that in our setup, uncertain migration leads to risky house price growth which encourages agents to speculate and to use IOs, rather than the introduction of IOs that induces agents to speculate. They would continue to speculate even if restricted to traditional mortgages. We expect the same would be true if lenders offered IOs because they were swept up in a speculative fervor, as Minsky (1982) argues, as restricting contract choice would not forestall the underlying euphoria. By contrast, our results suggest allowing recourse may be more effective in curbing speculation. Recall that we found IO shares were higher in states without recourse. This argument is also consistent with findings by Bäckman and Lutz (2017) that in Denmark, where loans are recourse, IOs were highly popular yet do not appear to have been used for speculation, and default rates remained low even when house prices fell. Of course, determining whether allowing for recourse is desirable requires a more rigorous setup that can speak to the costs and benefits of this approach. Other possible interventions include leverage restrictions that make speculation less profitable, although these have to be balanced against the harm done to more liquidity constrained home buyers who would also be affected.

In addition, we find that a significant amount of the variation in IOs across cities are consistent with perfectly benign reasons for using such mortgages, namely demand by households for mortgages that match their income flows or their relative impatience. Even though there have been some work lately pointing to the recent return of IOs as a source of concern, e.g. Sullivan (2018), updating our mortgage data to 2018 revealed that IO use was highest in New York and Connecticut, both states with full recourse, and not the non-recourse states where IOs surged back in the mid 2000s. This suggests the recent rise in IOs is probably not due to speculation as in the mid 2000s. One could replicate our remaining analysis to check whether the recent pattern in IO use is concentrated in cities with higher expected income growth or less affordable housing.

Finally, we caution that while our findings suggest IO use in cities with rapid house price growth indicated ongoing speculation, it is not obvious that speculation will always be manifested in IO use. To the extent that speculation leads to overvaluation, lenders will want to reward speculators for committing to selling their assets faster than they would otherwise. IO contracts are one way of achieving this, but any feature that creates an incentive for borrowers to sell their asset could in principle achieve the same outcome. In this regard, the fact that IO use in 2018 is below 1% when we look at the same Black Knight data we use to look at the mid 2000s does not mean policymakers can rest easy, as there may be other types of mortgages that could indicate speculation.

References

Allen, F., G. Barlevy, and D. Gale (2019). Asset price booms and macroeconomic policy: a risk-shifting approach. Federal Reserve Bank of Chicago Mimeo.

Allen, F. and D. Gale (2000). Bubbles and crises. *The Economic Journal 110*(460), 236–255.

Amromin, G., J. Huang, C. Sialm, and E. Zhong (2018). Complex mortgages. *Review of Finance*.

Arellano, M. and S. Bond (1991). Some tests of specifiation for panel data: Monte carlo evidence and an application to employment equations. *Review of Economic Studies 58*, 277–297.

Bäckman, C. and C. Lutz (2017). The impact of interst-only loans on affordability. https://ssrn.com/abstract=3055240. Accessed: 2018-09-30.

Barlevy, G. (2014). A leverage-based model of speculative bubbles. *Journal of Economic Theory 153*(September), 459–505.

Bostic, R., S. Chomsisengphet, K. Engel, P. McCoy, A. Pennington-Cross, and S. Wachter (2012). Mortgage product substitution and state anti-predatory lending laws: Better loans and better borrowers. *Atlantic Economic Journal 40*, 273–294.

Brueckner, J., P. Calem, and L. Nakamura (2016). House price expectations, alternative mortgage products and default. *Journal of Money, Credit and Banking 48*(1), 81–112.

Burnside, C., M. Eichenbaum, and S. Rebelo (2016). Understanding booms and busts in housing markets. *Journal of Political Economy 124*(4), 1088–1147.

Case, K., R. Shiller, and A. Thompson (2013). What have they been thinking? Homebuyer behavior in hot and cold markets. *Brookings Papers on Economic Activity 45*(2), 265–315.

Chambers, M., C. Garriga, and D. E. Schlagenhauf (2009). Accounting for changes in the homeownership rate. *International Economic Review 50*(3), 677–726.

Chiang, Y.-M. and J. Sa-Aadu (2014). Optimal mortgage contract choice decision in the presence of pay-option adjustable rate mortgage and the balloon mortgage. *Journal of Real Estate Finance and Economics 48*, 709–753.

Cocco, J. (2013). Evidence on the benefits of alternative mortgage products. *Journal of Finance 68*(4), 1663–1690.

Corbae, D. and E. Quintin (2016). Leverage and the foreclosure crisis. *Journal of Political Economy 123*(1), 1–65.

Davis, M., J. Fisher, and M. Veracierto (2015). Gross migration, housing and urban population dynamics. Chicago Fed working paper.

Dhillon, U., J. Shilling, and C. Sirmans (1990). The mortgage maturity decision: The choice between 15-year and 30-year FRMs. *Southern Economic Journal 56*(4), 1103–1116.

Dokko, J., B. Keys, and L. Relihan (2019). Affordability, financial innovation and the start of the housing boom. Federal Reserve Bank of Chicago Working Paper 2019-1.

Elul, R., N. Souleles, S. Chomsisengphet, D. Glennon, and R. Hunt (2010). What "triggers" mortgage default. *American Economic Review 100*(2), 490–494.

Foote, C., K. Gerardi, L. Goette, and P. Willen (2010). Reducing foreclosures: No easy answers. In D. Acemoglu, K. Rogoff, and M. Woodford (Eds.), *NBER Macroeconomics Annual 2010*, Volume 204, pp. 89–138.

Ghent, A. and M. Kudlyak (2011). Recourse and residential mortage default: Evidence from US states. *The Review of Financial Studies 24*(9), 3139–3186.

Gorton, G. (2008). The Panic of 2007. In *Economic Policy Symposium*, pp. 131–262. Federal Reserve Bank of Kansas City.

Guerrieri, V., D. Hartley, and E. Hurst (2013). Endgoenous gentrification and housing price dynamics. *Journal of Pubic Economics 100*(C), 45–60.

Gyourko, J., A. Saiz, and A. Summers (2008). A new measure of the local regulatory environment for housing markets. *Urban Studies 45*(3), 693–729.

Hart, O. and J. Moore (1994). A theory of debt based on the inalienability of human capital. *The Quarterly Journal of Economics 109*(4), 841–ï¿½879.

Justiniano, A., G. Primiceri, and A. Tambalotti (2017). The mortgage rate conundrum. NBER Working Paper No. 23784.

Justiniano, A., G. Primiceri, and A. Tambalotti (2019). Credit supply and the housing boom. *Journal of Political Economy*.

LaCour-Little, M. and J. Yang (2010). Pay me now or pay me later: Alternative mortgage products and the mortgage crisis. *Real Estate Economics 38*(4), 687–732.

Landvoigt, T., M. Piazzesi, and M. Schneider (2015). The housing market(s) of San Diego. *American Economic Review 105*(4), 1371–1407.

Mayer, C., T. Piskorski, and A. Tchistyi (2013). The inefficiency of refinancing: Why prepayment penalties are good for risky borrowers. *Journal of Financial Economics 107*(3), 694–714.

Mian, A. and A. Sufi (2009). The consequences of mortgage credit expansion: evidence from the US mortgage default crisis. *Quarterly Journal of Economics 124*(4), 1449–1496.

Mian, A. and A. Sufi (2018). Credit supply and housing speculation. NBER Working Paper No. 24823.

Minsky, H. P. (1982). *Can "It" Happen Again? Essays on Instability and Finance*. Armonk, New York: M. E. Sharpe, Inc.

Ono, A., H. Uchida, G. Udell, and I. Uesugi (2013). Lending pro-cyclicality and macro-prudential policy: Evidence from Japanese LTV ratios. http://dx.doi.org/10.2139/ssrn.2262575.

Piskorski, T. and A. Tchistyi (2011). Stochastic house appreciation and optimal mortgage lending. *Review of Financial Studies 24*(5), 1407–1446.

Saiz, A. (2010). The geographic determinants of housing supply. *Quarterly Journal of Economics 125*(3), 198–217.

Sullivan, P. (2018, December). The return of the risky loan. *New York Times*, B5.

Williams, O. (2006). Alternative mortgage products: Impact on defaults remains unclear, but disclosure of risks to borrowers could be improved. G.A.O. Report to the Chairman, Subcommittee on Housing and Transportation, Committee on Banking, Housing, and Urban Affairs, U.S. Senate.

Zeira, J. (1999). Informational overshooting, booms, and crashes. *Journal of Monetary Economics 43*, 237–257.

For online publication only

A. Data

This appendix provides a detailed description of our data construction.

A.1. Mortgage Data

Our mortgage data is primarily drawn from Black Knight Financial Services Services (BK) mortgage performance data, formerly known as McDash. We describe this data below, as well as other mortgage datasets we either merged with or used to supplement the BK data.

A.1.1. Black Knight Financial Services Mortgage Data (McDash)

The BK data is reported at a monthly frequency, based on data provided by various mortgage servicers on the loans they service. From these monthly reports we construct a single "static" file that includes a single record on each loan ever observed based on information on the loan at origination as well as over the life of the loan. We use this information to calculate the mortgage shares as well as identify the types of the loans included in the linear probability regressions.

An issue with BK is that different servicers begin to report data to BK at different dates. When a servicer joins, they typically only report on their outstanding loans. Mortgages that originated before the servicer began reporting to BK are therefore disclosed if they survive into the reporting period. This leads to survivorship bias in the data. Some have argued for only using data from 2005 on, once a significant number of servicers already began to report their data. For our purposes, this would throw out the most relevant period in which IOs and prices were rising. An alternative approach argues for only including mortgages that are reported in BK shortly after their origination date, and throwing out any mortgage the servicer reports but which we know was originated some time ago. This would preserve mortgages issued before 2005 by servicers who had already contracted with BK. This approach is used, for example, by Foote, Gerardi, Goette, and Willen (2010). The problem with this approach is that it may lead to biases due to heterogeneity across servicers. As mentioned by Williams (2006) (p11), IOs were highly concentrated, with a small number of lenders accounting for the bulk of such mortgages. If such servicers happened to report late to BK, as indeed appears to be the case, restricting attention to mortgages that are reported soon after origination would make it seem as if IOs surged later than they actually did. For this reason, we chose to include all mortgages in BK, regardless of whether they were originated before the servicer began to report to BK.

To get a sense of the extent of survivorship bias in the data we use, we compared mortgage counts in BK to those reported under the Home Mortgage Disclosure Act (HMDA). The latter dataset is reported to regulators in real time and should be free of any survivorship bias. Up until early 2003, the ratio of the number of mortgages we observe in BK to the number of mortgages reported in HMDA exhibits a rising trend, suggesting data in the early years of our sample undercounts mortgages. But we see no clear trends from 2003 on; the share of all for-purchase mortgages in BK to all mortgages in HMDA in 2003q1 (62%) is no different from 2010q4 (61%). This is consistent with the fact that interest rates bottomed out in 2003, so refinancing should be a lesser issue for mortgages originating after this date.

We also looked for evidence of survivorship bias in IOs. In particular, we looked at the number of IOs in the LoanPeformance (LP) database that we discuss below, which provides data on all mortgages securitized in private-label mortgages pools of nonprime mortgages. Again, this data should not be subject to survivorship bias. When we compare the ratio of the number of IO mortgages in LP to the number of IO mortgages in the BK data that are identified as privately securitized, we find no trend between from 2003q1 on, with the ratio fairly stable around 60%. This suggests suvivorship bias is not an important problem from 2003 on. Moreover, since backloaded mortgages are more likely to pay early, to the extent there is survivorship bias, it should cause our series on the use of such mortgage to lag true usage earlier in our sample. As such, it should make the use of IOs appear to lag house price growth, biasing against our findings.

Using the BK data, we identify first-lien (LIEN_TYPE), for-purchase mortgages originating in each CBSA (PURPOSE_TYPE_LP) and each quarter (ORIG_DT) and classify them using criteria we describe below. Mortgages in the data are reported by zip code, which we then aggregate to the CBSA level. For zip codes that do not fall entirely within a single CBSA we assigned all mortgages for that zip code to the CBSA with the largest share of houses for that zip code.

It is possible to double count mortgages if a mortgage was transferred between servicers in BK, i.e. if one servicer sold the mortgage to another. To avoid this, we matched all loans in our data on zip code, origination amount, appraisal amount, interest type, subprime status, level of documentation, the identity of the private mortgage insurance provider if relevant, payment frequency, indexed interest rate, balloon payment indicator, term, indicators for VA or FHA loans, margin rate, and an indicator of whether the loan was for purchase or refinance, treating missing and unknown values as wildcards. Loans that matched on all these criteria were treated as duplicates, and we kept only one record in such cases.

To identify IO mortgages, we use the IO flag (IO_FLG) reported by BK, which in turn is based on payment frequency type (PMT_FREQ_TYPE). We do not include any VA and FHA residential loans in our count of IOs. Since BK only started classifying loans as IO in 2005, for mortgages that originated and terminated before 2005, we looked at whether the initial scheduled payment in the first month (MTH_PI_PAY_AMT) was equal to the interest rate on the mortgage that month (CUR_INT_RATE) times the initial amount of the loan (ORIG_AMT). Using mortgages that survive past 2005 revealed that in a small but non-negligible number of IOs, the scheduled payment was not equal to but exactly twice the monthly interest rate times the initial loan amount, perhaps because of a quirk in the reporting convention of some servicers. Experimenting with the post-2005 data led us to classify as IOs those mortgages where the ratio of the scheduled payment to the interest rate times original loan amount was in either [0.985, 1.0006] or [1.97, 2.0012]. This approach correctly identified 98.5% of IOs while falsely identifying about 1.5% of non-IOs as IO in the post-2005 period.

Subprime mortgages in BK are mortgages whose mortgage type (MORT_TYPE) is coded as Grade 'B' or 'C', following Foote et al. (2010). For robustness, we experimented with two other measures of the share of subprime mortgages using other datasets. One is the ratio of the number of mortgages in subprime mortgage pools for each CBSA and quarter reported by LoanPerformance (see description below) to the number of mortgages in each CBSA and quarter. The other is the share of mortgages in HMDA in each CBSA and quarter that were originated by lenders identified in HMDA as subprime lenders.

Other mortgage classifications are constructed as follows. Long-term mortgages are mortgages with an amortization term (TERM_NMON) strictly greater than 360 months. Option-ARM mortgages are mortgages with code OPTIONARM_FLG set to 'Yes'. Fixed rate and ARM mortgages are identified with code PROD_TYPE set to '10' and '20' respectively. All mortgages that have a known type (PROD_TYPE not 'Unknown' or missing) but are not classified above as IO, Option-ARM, Fixed or ARM are classified as 'Other'.

To assign whether a mortgage is privately securitized we make use of the fact that while it takes some time for mortgages to be either privately securitized or purchased by government sponsored enterprises (GSEs), the turnover rate for these mortgages once they end up securitized is quite low. For mortgages in which the INVESTOR_TYPE field is avaliable 12 months after origination, we assign the investor type to the value of INVESTOR_TYPE for this month. However, investor type may be unavailable 12 months after origination: the mortgage may terminated before 12 months, the mortgage may have lasted longer than 12 months but the servicer did not report the loan to BK until at least 12 months after the loan originated, or because the data may be missing. For mortgages where investor type is not available 12 months after origination, we proceed as follows. If the loan was first reported to BK at least 7 months or more after origination, we use the first investor type reported. This would give us the investor type at least 6 months after origination, which tends to be highly persistent over the life of the mortgages in our sample. If the loan was first reported to BK within 6 months of origination, we use the last investor type reported if the loan terminated before 12 months. Otherwise, we treat investor type as missing. Privately securitized mortgages have an investor type which corresponds to a privately securitized mortgage pool. GSE mortgages are mortgages whose investor type is GNMA, FNMA, or FHLMC. Portfolio mortgages are those with a portfolio investor type.

Mortgages are classified as having a pre-payment penalty if the PP_PEN_FLG is set to 'Yes'. In cases where this variable is 'Unknown' or missing we exclude the loan from the sample in our regressions that distinguish between loans with and without pre-payment penalties.

Hybrid ARMs (2/28, 3/27) are determined using the FIRST_RATE_NMON variable which states the number of months the initial interest rate is fixed. We exclude these ARMs by excluding mortgages with FIRST_RATE_NMON greater than or equal to 24 and less than or equal to 36 months.

We also identify mortgages for properties that are reported by the borrower as purchased for investment rather than with the intent of the owner to occupy the property. Such mortgages are identified with an occupancy status (OCCUPANCY_TYPE) given as "Non-owner/Investment."

To be included in the counts in Table 1 a loan must be known to be sub-prime, privately securitized, have a pre-payment penalty and be taken out by an owner that does not intend to occupy the property.

For the share variables, we computed the maximum share of each mortgage type among all first-lien for-purchase mortgages in each CBSA between 2003q1 and 2008q4 that have known attributes.

All of these maximum shares are computed at the CBSA level, and used in our cross-sectional analysis where the primary unit is a CBSA. When we look at the propensity of different mortgages to default or terminate due to either a sale or a refinance, we instead analyze data at the level of individual mortgages. Here, we use the set of controls suggested by Elul et al. (2010) to analyze the propensity for default. These variables include the first observed interest rate (CUR_INT_RATE from the loan's first appearance in the dynamic data), FICO score (FICO_ORIG) and its square, the log of the original loan amount (ORIG_AMT), two indicators of private mortgage insurance, one by loan type (LOAN_TYPE) and the other as reported by the mortgage insurance company (PMI_CODE_TYPE), privately securitized flag as defined above, GSE investor type as defined above, FHA loan type flag (LOAN_TYPE), condo property type flag (PROP_TYPE), a flag for low/no documentation (DOCUMENT_TYPE), flags for 15-year and 40-year mortgages (identifed from the variable TERM_NMON), and the loan to value ratio (LTV_RATIO).

A.1.2. Recorder of Deeds data (DataQuick)

The BK mortgage data described above does not include information on the nature of property transactions. As a result, we cannot distinguish between a mortgage that paid early because the property was sold and one where the mortgage paid early because the loan was refinanced. To distinguish between the two, we matched the mortgage data from BK with data from the DataQuick (DQ) dataset on recorder of deeds, which was subsequently purchased by and renamed as CoreLogic. The DQ dataset contains information on transactions by property. To the extent we can match this to the BK dataset, we can determine how the mortgage paid off early.[27]

To match records betweek BK and DQ, we initially matched records using static data available from the time that the mortgage is originated. Broadly, we matched records with the same zip code and closing date in which the price of the house (house value) and the amount of the loan (origination amount) were the same in the two databases. We also required the mortgage be identified as for-purchase in both datasets. We could also identify from the DQ data whether a mortgage has a fixed or adjustable rate schedule. However, there was some evidence that this variable was not always recorded correctly in DQ, and so we also matched mortgages that matched along other characteristics regardless of whether they were classified as fixed or adjustable in DQ.

We considered several different criteria for matching. If at least one of these criteria were satisfied we considered the mortgage matched. While BK records a single origination date for each mortgage, DQ lists two dates for each mortgage, the filing date when in the mortgage was filed with the recorder of deeds and the transfer date in which ownership was transferred. From our match, it appears that the filing date typically occurs on or before the origination date, while the transfer date typically occurs on or after the origination date. Our strictest criterion matched records in BK and DQ from the same zip code with identical house values and origination amounts and where the origination date exactly matches either the filing date or transfer date. Another criterion we used required that either the filing date fell between 21 days before to 7 days after the origination date or the transfer date fell

[27] A complete documentation of our matching algorithm is available upon request.

either between 5 days before the origination date to 21 days after the origination date or else between 30 and 37 days after the origination date. Since some servicers in BK reported the closing date for a mortgage as the 1st of the month in which the mortgage originated regardless of the true origination date, we considered mortgages where the house value and origination amount matched exactly, and either the filing date or transfer date occurred in the same month as in DQ, but the mortgage was listed as originating on the 1st of the month in BK. For the same date criteria, we also considered mortgages where the origination amount matched exactly, where the house value in BK and DQ were with $1000 of one another, or where the origination amount matched exactly while the house value in DQ was listed as $0, or where the house value in both datasets matched exactly but the origination amount was equal to $0 in DQ. We found that variation in origination amount tended to produce many false matches, i.e. this variable was recorded with more care than house value. Across all criteria and cities, we were able to match about 40% of the mortgages in BK to DQ rising to about 50% of the mortgages from 2009.

Once loans from the two datasets have been matched on these static criteria, we took all the mortgages that we matched and were identified in BK that were paid in full and which included a paid-in-full date in BK. For these mortgages, we looked at whether there was either a sale or refinance on the same property in DQ anywhere from 30 days before to 30 days after the paid-in-full date. To avoid spurious matches in which the transaction in DQ that occurred within 30 days of the paid-in-full date in BK did not represent a sale or refinance of the original mortgage we matched between BK and DQ, we required that either the name on the original mortgage transaction was the same as the one when the mortgage terminated, or else that there was a chain of transactions in which the name on the original mortgage transferred ownership with no money changing hands to another entity, who transferred ownership with no money changing hands to some other entity, and so on, until the name associated with the final transaction. Among mortgages in which another transaction was recorded in DQ within 30 days of the paid-in-full date in BK, we could connect the name on the original mortgage to the final transaction, and identify the later transaction as either a sale or refinance, in 95% of the cases.

Our sample for the linear probability regressions included all BQ mortgages that were of a given type and for which we had a full set of controls that never pay in full or default plus all mortgages that meet this criteria but do pay in full and we are able to match with DQ and determine whether the property was sold or the mortgage was refinanced.

A.1.3. Other Mortgage Data

We also used data from the LoanPerformance (LP) data to supplement the BK data. The LP data reports mortgages in private-label mortgages pools of nonprime mortgages, meaning Alt-A and subprime mortgages. We used this data in two ways. First, we used it together with HMDA data to obtain an alternative estimate of the share of mortgages in each CBSA and each quarter that were subprime. Second, in contrast to the BK data, the LP data matches all liens against a property and reports the combined loan-to-value (CLTV) ratio for each property, including second liens. This allowed us to compute the fraction of first-lien for-purchase mortgages in LP in which the cumulative LTV of all loans at origination exceeded 80%.

We used HMDA data to construct another measure of the share of subprime mortgages in a city. In particular we used the share of loans issued by known subprime lenders as classified by HMDA.

A.2. House Price Data

Our primary measure of house price appreciation is the Federal Housing Finance Agency (FHFA) house price index, computed at the level of Core-based Statistical Areas (CBSAs) as defined in 2009 by the Office of Management and Budget. CBSAs with populations greater than 2.5 million are divided into Metropolitan Divisions.[28] For these CBSAs, FHFA reports data for each Division rather than the CBSA as a whole. We follow this convention throughout, using Metropolitan Divisions in lieu of the CBSA where applicable.

For robustness, we also constructed CBSA-level price indices based on the CoreLogic House Price Index and the Zillow Home Value Index. CoreLogic reports prices at the same CBSA level as FHFA. Zillow was available to us at the CBSA level as well, but for fewer cities.

For each price series, we construct our price variables as follows. We first convert each house price series into a real series by dividing through by the Consumer Price Index for urban consumers as reported by the Bureau of Labor Statistics.

For each CBSA, we identify the quarter within the period 2003q1-2008q4 in which the real price peaks in each respective CBSA. If we denote the price in a given city at date t by p_t, then the quarter in which price peaks is just

$$t^* = \arg\max_t \{p_t\}_{t=2003q1}^{2008q4}$$

Our summary measure of real house price appreciation in each city is the highest 4-quarter growth in real house prices between $t = 2003q1$ and t^*, i.e.

$$\text{Max4QGrowth} = \max\{\ln(p_t/p_{t-4})\}_{t=2003q1}^{t^*}$$

A.3. Other Data

We compiled additional cross-sectional data on CBSAs on non-mortgage data from various sources. Where necessary, we used translation tables from MABLE/Geocorr2K, the Geographic Correspondence Engine based on the 2000 Census from the Missouri Census Data Center, to convert data to the CBSA level. For each variable and except where noted below, we calculated both the average level and the average change between 2003q1 and the quarter in which real houses peak in that CBSA (or 2003 and the year in which the peak occurs for annual variables).

Population for each CBSA comes from the Census Bureau's Current Population Reports, P-60, at an annual frequency. All of our averages use log average annual population. We

[28] The CBSAs that are divided into Metropolitan Divisions are: Boston-Cambridge-Quincy, MA-NH; Chicago-Naperville-Joliet, IL-IN-WI; Dallas-Fort Worth- Arlington, TX; Detroit-Warren-Livonia, MI; Los Angeles-Long Beach-Santa Ana, CA; Miami-Fort Lauderdale-Miami Beach, FL; New York-Northern New Jersey-Long Island, NY-NJ-PA; Philadelphia-Camden-Wilmington, PA-NJ-DE-MD; San Francisco-Oakland-Fremont, CA; Seattle-Tacoma-Bellevue, WA; and Washington-Arlington-Alexandria, DC-VA-MD-WV.

work with log population in 1000s.

Real per capita personal income for each CBSA comes from the Bureau of Economic Analysis at an annual frequency. We work with log real per capita personal income in 1000s of 2005 dollars.

Unemployment rates for each CBSA come from the Bureau of Labor Statistics (BLS), and are available at a monthly frequency. We aggregate up to a quarterly frequency by averaging the months in each quarter and then compute quarterly averages. We work with unemployment rates in percentage points. Property tax rates for each CBSA are constructed from data in the American Community Survey (ACS) from the US Census Bureau, using an extract request from IPUMS USA (available at http://usa.ipums.org). In particular, we took data on annual property taxes paid (PROPTX99) and house value (VALUEH). Since PROPTX99 is a categorical variable, we set the tax amount to the midpoint of each respective range. Thus, a tax in the range of $7,001-$8,000 is coded as $7,500. Anything above $10,000 is coded as $10,000. For each household, we estimate the tax rate as the ratio of taxes paid divided by the value of the house. We then compute the median tax rate across all households in the survey in each CBSA in each year. Focusing on the median mitigates the top-coding in taxes paid. Since the ACS has its own definition of metro areas, we need to use the IPUMS metro area-to-MSA/PMSA translation table and then use a MSA/PMSA-to-CBSA table from GEOCORR2K. We also weight households by household weight (HHWT). The number of cities for which the property tax variable is available is substantially lower in 2003 compared to other survey years. For this reason we use tax rates in percentage points in 2000 for the level and changes between tax rates in percentage points between 2000 and the year of the peak real house price.

Industry employment shares for the years 1995 and 2000 use county-level employment counts from the Bureau of Economic Analysis, available an annual frequency. The county-level data are aggregated into CBSAs using Office of Management and Budget's 2009 definitions, then we calculate each industry's share of total nonfarm employment. The industries include NAICS two digit codes 11, 21, 22, 23, 31-33, 42, 44-45, 48-49, 51, 52, 53, 54, 55, 56, 61, 62, 71, 72, 81, 92.

The gross migration data is from Davis, Fisher, and Veracierto (2015) who construct city-level in- and out-migration rates based on the IRS county-to-county migration matrices.

The maximum price-to-income ratio variable was calculated as follows. First we obtain the median level of house prices in each US county in 2000 from the Census Bureau. These medians are calculated by the Census Bureau using the micro data from the 2000 Decennial Census of Housing.[29] The median house price for a city is identified as the housing-unit-share-weighted sum of median prices in the counties comprising a city where the housing units are based on the same underlying 2000 Decennial Census of Housing and are calculated by the Census Bureau. We identify median house prices for each year after 2000 by applying the growth of the FHFA price index for that city to the 2000 median price over the relevant period. The price-to-income ratio in each year is obtained by dividing the nominal price by nominal per capita personal income for the city described above. The maximum price-to-income ratio is the maximum value of this last variable over the 2003-2008 period.

The median age variable is population-weighted average of median ages in each county

[29] These house prices are self reported and correspond to single family homes only.

comprising a city. The median ages and population by county are obtained from the Census Bureau and are based on the 2000 Decennial Census.

B. Theory

B.1. Proofs of Propositions

Proof of Proposition 1: Rearranging (2) implies

$$\frac{p_{t+1}}{p_t} = \frac{1}{\beta(1-q)}\left(1-\frac{d}{p_t}\right)+\frac{d}{p_t}$$

This expression will exceed $1/\beta$ if

$$\frac{1}{\beta(1-q)}\left(1-\frac{d}{p_t}\right)+\frac{d}{p_t} > \frac{1}{\beta}$$

or, upon rearranging,

$$\frac{d}{p_t} < \frac{q}{1-\beta(1-q)} \tag{7}$$

The highest price we can observe before t^* occurs at date t^*-1, when from (2) we know that

$$p_{t^*-1} = (1-\beta)\,d+\beta qd+\beta(1-q)\,D$$

Substituting this for p_t implies

$$\frac{d}{(1-\beta(1-q))\,d+\beta(1-q)\,D} < \frac{q}{1-\beta(1-q)}$$

which can be rearranged to yield

$$\frac{1-\beta(1-q)}{\beta q} < \frac{D}{d}$$

Hence, if $\frac{D}{d} > \frac{1-\beta(1-q)}{\beta q}$, the growth rate p_{t^*}/p_{t^*-1} will exceed $1/\beta$.

Proof of Proposition 2: Recall that if cohorts arrive through date t, then

$$p_t = (1-\beta)\,d+\beta E_t\left[p_{t+1}\right]$$

Rearranging, we get

$$\begin{aligned} E_t\left[\frac{p_{t+1}}{p_t}\right] &= \frac{1}{\beta}-\frac{1-\beta}{\beta}\frac{d}{p_t} \\ &= \frac{1}{\beta}\left(1-\frac{d}{p_t}\right)+\frac{d}{p_t} \end{aligned}$$

Since $d < p_t < D$, then as long as $t < t^*$, we have

$$\frac{d}{p_t} < 1$$

and so

$$1 < E\left[\frac{p_{t+1}}{p_t}\right] < \frac{1}{\beta}$$

as claimed.

B.2. Risk Averse Borrowers

This section considers the case where home buyers are risk averse. In particular, we replace the utility in (1) for cohort τ with

$$\sum_{t=\tau+1}^{\infty} \beta^t u\left(c_{it} + v_i h_{it}\right) \tag{8}$$

where $u''(\cdot) \leq 0$. The utility function in (8) assumes consumption and housing services are perfect substitutes. It also assumes individuals from cohort τ only care about consumption from date $\tau + 1$. This assumption substantially simplifies the analysis but is not essential.

For tractability, we further assume individuals earn a constant income stream that begins to accrue one period after they arrive. That is, $\omega_{i\tau}^{\tau} = 0$ and $\omega_{it}^{\tau} = \omega$ for $t = \tau + 1, \tau + 2, ...$ We continue to maintain condition (6) by requiring

$$\frac{\beta}{1-\beta}\omega > D$$

This ensures homeowners could repay a loan of D at an interest rate of $1/\beta$.

In addition, we assume a large group of risk-neutral lenders with ample wealth who discount at rate β. This ensures the gross equilibrium risk-free rate will equal $1/\beta$.

Finally, we assume $t^* = 2$, meaning uncertainty is resolved at date 2. This implies

$$p_t = \begin{cases} d \text{ for all } t \geq 2 & \text{w/prob } q \\ D \text{ for all } t \geq 2 & \text{w/prob } 1-q \end{cases}$$

Hence, we only need to solve for house prices at date 1.

To determine the equilibrium price at date 1, we work backwards from date 2, when all uncertainty is resolved. If $p_2 = d$, low types will be indifferent about owning a house, since they value a house at d, while high types will strictly prefer to own a house, since they value a house at D. If instead $p_2 = D$, low types will prefer to sell a house rather than own it, while high types will be indifferent about owning a house. Hence, we can assume without loss of generality that a high type will own a house from date 2 on, i.e. he will either keep a house he bought earlier or buy a house if he hasn't yet, and that a low type will not own a house from date 2 on, i.e. he will either sell a house he bought earlier or not a buy a house if he hasn't yet.

For a low type who arrives at date 1, if he does not buy a house at date 1, then since he will not buy a house at date 2 he will consume his permanent income from date 2 on, i.e. $c_{it} = \omega$ for $t \geq 2$. The expected utility as of date 1 from this strategy is

$$\frac{\beta}{1-\beta} u(\omega) \tag{9}$$

Alternatively, a low type can buy a house at date 1. Since low types are indifferent between occupying a house or renting it out, we can proceed as if a house purchase is purely an investment. That is, the low type will be able to rent out his house for $(\beta^{-1} - 1)\, d$ at date 2, then sell the house at price p_2, against which he will owe p_1/β from borrowing to buy the house at date 1. Hence his wealth W_2 at date 2 is given by

$$W_2 = \begin{cases} \dfrac{\omega}{1-\beta} + (\beta^{-1} - 1)\, d + d - \dfrac{p_1}{\beta} & \text{with probability } q \\ \dfrac{\omega}{1-\beta} + (\beta^{-1} - 1)\, d + D - \dfrac{p_1}{\beta} & \text{with probability } 1-q \end{cases}$$

Since $u(\cdot)$ is concave, a low type agent will choose a constant consumption c_{it} from date $t = 2$ on. This means $c_{it} = (1-\beta)\, W_2$, or

$$c_{it} = \begin{cases} \omega + (\beta^{-1} - 1)\,(d - p_1) & \text{with probability } q \\ \omega + (\beta^{-1} - 1)\,(d - p_1) + (1-\beta)(D-d) & \text{with probability } 1-q \end{cases}$$

The expected utility for the low type is then

$$\frac{\beta}{1-\beta} \left[\begin{array}{c} (1-q)\, u\left(\omega + (\beta^{-1} - 1)\,(d - p_1) + (1-\beta)\,(D-d)\right) \\ + qu\left(\omega + (\beta^{-1} - 1)\,(d - p_1)\right) \end{array} \right]$$

Hence, a low type will buy a house at date 1 iff

$$(1-q) u \left(\omega + \frac{1-\beta}{\beta}\,(d - p_1) + (1-\beta)(D-d)\right) + qu\left(\omega + \frac{1-\beta}{\beta}\,(d-p_1)\right) \geq u(\omega) \tag{10}$$

Next, we consider high type agents. Suppose a high type buys a house at date 1. Since he will keep the house at date 2, purchasing the house immediately involves no uncertainty. At date 2 he will have $\omega/(1-\beta) - p_1/\beta$ in non-housing wealth, D in terms of housing wealth, and $(\beta^{-}1 - 1)D$ worth of housing services they can consume. With concave utility, the optimal plan would be to set the sum of his consumption and the consumption equivalent value of housing services each period to a fraction $1-\beta$ of these resources, i.e.

$$c_{it} + (\beta^{-}1 - 1) D h_{it} = (1-\beta) + \frac{D - p_1}{\beta} + D$$

for all $t \geq 2$. This will yield an expected utility of

$$\frac{\beta}{1-\beta} u\left(\omega + \left(\beta^{-1} - 1\right)(D - p_1)\right)$$

Alternatively, a high type could wait to buy a house at date 2. In this case, his non-housing wealth at date 2 will equal $\omega/(1-\beta) - p_2$ and his housing wealth will equal D. From date 3 on he will be able to consume housing services. Again, the optimal plan would be to set the sum of his consumption and the consumption equivalent value of housing services each period to a fraction $1-\beta$ of his wealth, i.e.

$$c_{it} + (\beta^{-1} - 1)Dh_{it} = (1-\beta)\left(\frac{\omega}{1-\beta} + D - p_2\right)$$

Since housing services only become available at date 3, the agent will enjoy more consumption at date 2 and then reduce his consumption intake but increase his intake of housing services from date 3. Since p_2 is equal to d with probability q and D with probability $1-q$, the expected utility as of date 1 from waiting to buy a house at date 2 is given by

$$\frac{\beta}{1-\beta}\left[\begin{array}{l} (1-q)\,u(\omega) \\ \qquad + qu(\omega + (1-\beta)(D-d)) \end{array}\right]$$

Hence, a high type will buy a house at date 1 iff

$$u\left(\omega + \frac{1-\beta}{\beta}(D-p_1)\right) \geq (1-q)\,u(\omega) + qu(\omega + (1-\beta)(D-d)) \tag{11}$$

We now show that if low types are willing to buy a house, high types will as well. A low type will buy iff (10) holds. Since $u(\cdot)$ is concave, the utility of the expectation of the arguments on the RHS exceeds the expression on the RHS. Since $u(\cdot)$ is increasing, this inequality implies the expectation of the arguments on the RHS exceeds ω, i.e.

$$\omega + (\beta^{-1} - 1)(d - p_1) + (1-\beta)(1-q)(D-d) \geq \omega$$

We can rearrange this condition to get

$$p_1 \leq (1-\beta)d + \beta(qd + (1-q)D)$$

That is, low types will only buy the house at date 1 if the price at date 1 is less than what they expect to earn from buying the house, renting it out in period 2 and then selling it. Since $D > d$, it follows that

$$p_1 < (1-\beta)D + \beta(qd + (1-q)D)$$

With a little algebra, we can rearrange this inequality to obtain

$$(1-\beta)q(D-d) > (\beta^{-1} - 1)(D - p_1)$$

Appealing to the fact that $u(\cdot)$ is concave, this implies condition (11). The latter condition ensures high types will want to buy housing at date 1. It follows that in equilibrium, high types will strictly prefer to buy houses. Since they buy houses from low types, low types must in equilibrium be indifferent between buying housing at date 1. This means that in equilibrium, condition (10) must hold with equality, i.e.

$$(1-q)u\left(\omega+\frac{1-\beta}{\beta}(d-p_1)+(1-\beta)(D-d)\right)+qu\left(\omega+\frac{1-\beta}{\beta}(d-p_1)\right)=u(\omega) \quad (12)$$

This equation determines p_1. From this equation, it follows that we can drive p_1 arbitrarily close to d by making $u(\cdot)$ sufficiently concave. To see this, define $x=\omega+(\beta^{-1}-1)(d-p_1)$. Then we can rewrite (12) as

$$(1-q)u(x+(1-\beta)(D-d))+qu(x)=u\left(x+\frac{1-\beta}{\beta}(p_1-d)\right)$$

As we make $u(\cdot)$ more concave, the only way for this equality to hold is if $x+(\beta^{-1}-1)(p_1-d)$ tends to x, or, alternatively, if p_1 tends to d. The expected house price growth, $E_1[p_2/p_1]$ will then tend to $q+(1-q)D/d$. If D/d is sufficiently large, specifically if it exceeds $(1/\beta-q)\,/\,(1-q)$, the expected growth rate $E_1[p_2/p_1]$ will exceed $1/\beta$.

B.3. Non-Recourse Lending

This section analyzes the case of non-recourse mortgages, i.e. where lenders can only go after the house a borrower purchased but not his income.

As a first step, we work with a restricted environment to make the analysis easier. We assume agents are risk-neutral as in the paper. We assume all houses at date $t=0$ are owned outright by agents and so there are no outstanding debt obligations against houses at this time. New cohorts start arriving at date 1, but since $\omega_{it}^t=0$ each cohort has no resources to pay for a house upon arrival. To avoid getting into questions about the optimal timing of purchases, we assume individuals must buy a house when they arrive or give up on the option to buy a house. We also assume that once an agent secures a loan to buy a house, they cannot subsequently refinance their loans. This avoids analyzing both refinancing and the constraints it imposes on the initial loan.

To simplify the analysis, we assume all agents earn a constant income ω for the first t^* periods of their life, starting with the first period after they arrive. We further assume that ω is small enough that agents will have a debt obligation that exceeds d for the first t^* periods of their loan. If this is true for the first cohort of home buyers who buy at date 1, it will necessarily be true for subsequent cohorts who will face even higher house prices. We do need to make sure that the income agents earn after t^* is high enough to ensure they can afford to buy a house when they borrow at the risk-free rate. Since default is already an issue, we can set the probability an agent becomes disabled ε to zero.

Since low types are the ones who engage in speculation, we want the fraction of low types to be positive, i.e. $\phi>0$. At the same time, we don't want this fraction to be too large. We can therefore think of setting $\phi\approx 0$, so lenders can expect to recover most of their loans even if low types borrow and then default, and the interest rate on loans will be close to the risk-free rate.

We also assume that the size of each arriving cohort, n, is small. To motivate this assumption, note that there are two relevant threshold dates in our economy. The first is t^*, the smallest integer that exceeds $\phi_0/\left[(1-\phi)\,n\right]$. At date t^*, high types will buy out the entire stock of housing. The second threshold, which we denote t^{**}, is the smallest integer that exceeds ϕ_0/n. At date t^{**}, the original non-high types who own houses at date 0 would get rid of their housing if they sold one unit to each agent who arrived between date 1 and t^{**}. Since $\phi_0/n < \phi_0/\left[(1-\phi)\,n\right]$, it follows that $t^{**} \leq t^*$. We want this inequality to be strict, meaning the original owners can sell their housing *before* high types occupy all available houses. For any value of ϕ, we can always set n to be sufficiently small that this condition will be satisfied. Formally, our analysis involves letting both ϕ and n tend to zero, but with n tending to 0 quickly enough to ensure $t^{**} < t^*$.

During the migration wave, all members of each cohort will want to buy a house. High types want to buy because they derive a large service flow from housing and intend to live in their house indefinitely. Low types want to buy because they want to speculate, i.e. to sell the house at a higher price if more high types arrive later and default otherwise without losing any of their income. The absence of recourse makes such speculation profitable.

Lenders will of course try to avoid lending to speculators. For example, they will refuse to lend to an agent to buy more than one house, since an agent only derives high utility from one house and will default on any other house if house prices fell. Any agent who borrows to buy two houses will not value at least one of them for its housing services. To avoid lenders trying to screen out low types by offering contracts that would discourage speculation, we assume lenders can only offer debt contracts with non-contingent and non-random repayment schedules. This is certainly consistent with the data, but here it is important in that it prevents lenders from issuing the type of contracts that only high types might agree to. In particular, we are ruling out contracts that stipulate payments that rise and fall with house prices which high types would be fine with but would eliminate speculative profits for low types, and we are ruling out the option of paying low types not to borrow, which is a lower cost way for lenders to provide low types with information rents than letting them borrow to buy an asset.

To anticipate the equilibrium we describe, all agents in each cohort will borrow to buy a house. New agents buy houses from those who own houses that have the lowest reservation price, which will be those with the lowest debt obligation against the house. Thus, until date t^{**}, agents will buy houses from those who already own these houses at date 0. From date t^{**} on, new buyers will buy houses from low types with the longest tenure as homeowners, i.e. from low types who were the first to arrive among current low type homeowners. Once we establish this is indeed what happens in equilibrium, we discuss the types of mortgages that the different types will use.

Let us first argue that high types want to buy houses before date t^*. Let p_t denote the price of houses if new cohorts arrive through date t. Appealing to the transversality condition as before, we know the price of housing cannot exceed D. Now, suppose $p_t = D$ for some $t < t^*$. Then agents who don't value housing services at D would want to sell at this date: Waiting will not yield a higher price that offers profits, but there is a risk the price will fall to d if the migration wave stops. Since there would not be enough high type agents to buy all of the housing that low types want to sell before t^*, this cannot be an equilibrium, and $p_t < D$ for all $t < t^*$. For small ϕ, the interest rate that ensures lenders earn zero profits will be close to the risk free rate. High types would then be strictly better off buying when they arrive than not buying at all. At date t^*, house prices must equal D if migration continued. Any remaining non-high type owners would prefer to sell at this date, and so in equilibrium high types buy these houses. Without any remaining uncertainty, high types are indifferent about borrowing at the risk-free rate to buy housing at this price.

Next, we argue that low types will also want to buy houses before date t^*. Before date t^{**}, there must be some original owners of houses who hold on to their houses given each agent can buy at most one house. But we now argue that there must also be original owners who are willing to sell. The only other agents who might sell are low types who previously bought the house. But we assumed their income ω was low enough that their debt obligation exceeds d for the first t^{**} periods of their loan. If they sell at date t, they will earn $p_t - L_t$, where L_t denotes their debt obligation at date t and, which by assumption, exceeds d. If they wait to sell next period, they will earn $\beta\left(\left(\beta^{-1}-1\right)d + E_t\left[p_{t+1}\right]\right)$, and if charged the risk free rate they will owe L_t/β. Thus, without default, they would be indifferent between selling today and waiting. But if they default when prices fall, they will benefit because the house they sell is worth less than their debt obligation, so they get an additional benefit from waiting and partly default on their obligation. Since this option to default is not available to original owners who have no debt against the house, original owners will be more willing to sell. By the same logic, those who have a larger debt obligation will prefer to wait to sell even when agents with a smaller debt obligation are indifferent.

The above discussion implies that for $t < t^{**}$, original owners will both hold and sell houses, meaning they must be indifferent between the two. This means

$$\beta\left(\left(\beta^{-1}-1\right)d + E_t\left[p_{t+1}\right]\right) = p_t$$

If migration continues into date $t+1$, the price of housing p_{t+1} will exceed $E_t\left[p_{t+1}\right]$, and low types could earn a return above $1/\beta$ by buying a house at price p_t, then renting it out and selling it at date $t+1$. Since agents can always default and continue consuming their income, low types will want to buy houses as long as ϕ is small and they are charged an interest rate close to $1/\beta$.

From date t^{**} on, house prices will be determined by the indifference of a marginal owner who has some debt obligation against the house. This marginal owner will suffer less if migration stops next period than an agent with no outstanding debt, so house prices don't grow as much as they do in the first t^{**} periods. But it will still be the case that the price must rise if migration continues and falls to d if migration stops.

Next, we argue that the price of housing p_t exceeds the value of housing services from the last unit of housing. Denote this value by f_t. We can characterize f_t recursively as follows. Before date t^*, the value of the last house is that it can be used to provide housing services to a low type. Thus, it generates a flow value of $\left(\beta^{-1}-1\right)d$. If migration continues through date t^*, the house can be used to provide housing services to a high type. If migration stops before t^*, the house can only ever be used to provide housing services to a low type. Hence, f_t is given by

$$f_t = \beta\left(\beta^{-1}-1\right)d + \beta\left(qd + (1-q)\,f_{t+1}\right)$$

Note the similarity with (2), which characterizes prices in the economy with recourse. We now argue that without recourse, the price p_t will exceed f_t. It cannot fall below this level, or else the lenders who finance new home buyers would simply buy a house themselves. Suppose, then, that $p_t = f_t$. We focus on $t = t^{**}$. At this date, those who buy houses will buy them from a low type agent who arrived earlier and still has a debt obligation that exceeds d. But $p_t = f_t$, such an agent would prefer not to sell. Selling today yields a gain of

$$p_t - L_t = f_t - L_t$$

while waiting one period yields a gain of

$$\beta (1 - q) \left(\left(\beta^{-1} - 1 \right) d + p_{t+1} - \beta^{-1} L_t \right) + \beta q \cdot 0$$

Since $L_t > d$, the latter expression is strictly higher than

$$\beta (1 - q) \left(q \left(d - \beta^{-1} L_t \right) + \left(\beta^{-1} - 1 \right) d + p_{t+1} - \beta^{-1} L_t \right)$$

and since $p_t \geq f_t$ for all t, this last expression is bounded below by

$$f_t - L_t$$

Hence, we cannot have an equilibrium in which $p_t = f_t$ at date t^{**}, since if we did then no agents would be willing to sell houses at date t^{**}. But if $p_t > f_t$ at date t^{**}, it must also be higher before date t^{**} from the fact that agents with no debt obligation must be indifferent between selling and waiting at these dates.

Finally, we turn to the type of mortgages that would trade during the migration wave. Suppose only one type of mortgage were offered. Then this mortgage would stipulate the fastest possible repayment. For if the only mortgage offered did not stipulate the fastest possible repayment, a lender could offer a contract with faster repayment and a slightly lower interest rate. High types would prefer this mortgage since a lower interest rate and faster payments means they have to pay less in total. But as long as the interest rate was only a little lower, low types would still prefer the original mortgage since slower repayment increases the option value to default. This way, a lender would only attract high types but earn slighly lower profits on each, meaning it can increase its total profits. The original contract with less than the fastest repayment could thus not have been an equilibrium.

Next, suppose the only contract offered in equilibrium was the fastest repayment mortgage. Some of these mortgages will not be repaid by date t^{**}, and so some agents who borrowed to buy a house after date 1 must wait until date t^{*} to sell. A lender could then do better by offering two contracts instead of one: the one with the fastest repayment, and another that forces the borrower to sell the asset by date t^{**}. Recall that at this date, any agents who have no debt obligation would have sold off their asset. In particular, there exist agents with a positive debt obligation who are indifferent about selling a house. That means that agents without a debt obligation would strictly prefer to sell, since they don't receive the benefit from the option to default. If we consider a lender and low type borrower together, their total value for the asset is the same as an agent who has no debt obligation against the asset. Thus, we can make the two of them jointly better off than under the original contract. This requires a contract that induces the borrower to sell the asset at date t^{**}. To make sure the borrower is willing to take this contract, the borrower's utility must be weakly higher than the contract with the fastest repayment path. But since both agents are collectively better off, it will be possible to do this and leave the lender better off. Hence, a single contract cannot be an equilibrium. The equilibrium must feature separation.

As discussed in Barlevy (2014), the equilibrium is a Spence-Miyazaki-Wilson separating equilibrium in which high types cross-subsidize low types. Lenders will offer high types contracts that require the fastest possible repayment, since these are least desirable to low types, and will offer low types contracts that encourage early repayment in exchange for some reward to the borrower for accepting a contract with more constraints. The interest rate on the contracts offered high types will offset the losses expected from lending to low types. Intuitively, since the asset is a bubble, selling it is better than holding it indefinitely. By the same logic, it is also better selling the asset than holding it for too long.

There are various ways to induce the borrower to sell the house they buy more quickly in exchange for some reward. One example is a mortgage with a balloon payment that is paired with a low interest rate. Another is an IO with low initial payments that becomes unaffordable. Thus, our framework does not imply speculation must result specifically in an IO mortgage rather than a balloon mortgage. However, if we considered an environment in which the only mortgages lenders could offer were those that pay as fast as possible and IOs, then in equilibrium both would be used. By contrast, if lenders had full recourse, the same framework would imply lenders would only offer mortgages with fast repayment even if they could offer IOs. In that sense, non-recourse encourages the use of IOs when house prices are uncertain.

Once we restrict lenders to IOs and fast repayment mortgages, we can further argue that the interest rates on these two types of mortgages must be such that low types are exactly indifferent between the two types of contracts. If that weren't the case, then a lender could offer only the mortgage with fast repayment and attract only high types, earning a profit. But this contradicts the fact that we know both contracts are offered in equilibrium. Hence, low types are indifferent. In addition, lenders must believe that low types will continue to borrow from them if lenders only offered the contract with the fastest possible repayment to sustain this equilibrium. An equilibrium is thus both separating and involves particular beliefs on the part of lenders.

Working Paper Series

A series of research studies on regional economic issues relating to the Seventh Federal Reserve District, and on financial and economic topics.

Working Paper Series *(continued)*

Establishments Dynamics and Matching Frictions in Classical Competitive Equilibrium **WP-07-16**
Marcelo Veracierto

Technology's Edge: The Educational Benefits of Computer-Aided Instruction **WP-07-17**
Lisa Barrow, Lisa Markman, and Cecilia Elena Rouse

The Widow's Offering: Inheritance, Family Structure, and the Charitable Gifts of Women **WP-07-18**
Leslie McGranahan

Incomplete Information and the Timing to Adjust Labor: Evidence from the
Lead-Lag Relationship between Temporary Help Employment and Permanent Employment **WP-07-19**
Sainan Jin, Yukako Ono, and Qinghua Zhang

A Conversation with 590 Nascent Entrepreneurs **WP-07-20**
Jeffrey R. Campbell and Mariacristina De Nardi

Cyclical Dumping and US Antidumping Protection: 1980-2001 **WP-07-21**
Meredith A. Crowley

Health Capital and the Prenatal Environment:
The Effect of Maternal Fasting During Pregnancy **WP-07-22**
Douglas Almond and Bhashkar Mazumder

The Spending and Debt Response to Minimum Wage Hikes **WP-07-23**
Daniel Aaronson, Sumit Agarwal, and Eric French

The Impact of Mexican Immigrants on U.S. Wage Structure **WP-07-24**
Maude Toussaint-Comeau

A Leverage-based Model of Speculative Bubbles **WP-08-01**
Gadi Barlevy

Displacement, Asymmetric Information and Heterogeneous Human Capital **WP-08-02**
Luojia Hu and Christopher Taber

BankCaR (Bank Capital-at-Risk): A credit risk model for US commercial bank charge-offs **WP-08-03**
Jon Frye and Eduard Pelz

Bank Lending, Financing Constraints and SME Investment **WP-08-04**
Santiago Carbó-Valverde, Francisco Rodríguez-Fernández, and Gregory F. Udell

Global Inflation **WP-08-05**
Matteo Ciccarelli and Benoît Mojon

Scale and the Origins of Structural Change **WP-08-06**
Francisco J. Buera and Joseph P. Kaboski

Inventories, Lumpy Trade, and Large Devaluations **WP-08-07**
George Alessandria, Joseph P. Kaboski, and Virgiliu Midrigan

Working Paper Series *(continued)*

Working Paper Series *(continued)*

Working Paper Series *(continued)*

Working Paper Series *(continued)*

Public Pensions and Labor Supply Over the Life Cycle **WP-10-09**
Eric French and John Jones

Explaining Asset Pricing Puzzles Associated with the 1987 Market Crash **WP-10-10**
Luca Benzoni, Pierre Collin-Dufresne, and Robert S. Goldstein

Does Prenatal Sex Selection Improve Girls' Well-Being? Evidence from India **WP-10-11**
Luojia Hu and Analía Schlosser

Mortgage Choices During the U.S. Housing Boom **WP-10-12**
Gadi Barlevy and Jonas D.M. Fisher