Appendices to
"Why were interest-only mortgages so popular during the U.S. housing boom?"
by Gadi Barlevy and Jonas Fisher, Review of Economic Dynamics, 2020

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 Inc's mortgage performance dataset, 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.[28]

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

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

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 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.[29] 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).

[29] 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.

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 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.[30] 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-

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

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 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

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

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 $\dfrac{D}{d} > \dfrac{1-\beta(1-q)}{\beta q}$, 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

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

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

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 $t = \tau + 1, \tau + 2, \ldots$. 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

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

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.

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

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 $\phi_0 / [(1-\phi)\, n]$. 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 $\phi_0/n < \phi_0/[(1-\phi)\, n]$, 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.