Online Appendix
Not For Publication

Liquidity Constraints in the U.S. Housing Market

Denis Gorea Virgiliu Midrigan

Version: April 2018

Contents

This appendix describes in great detail the datasets and procedures that we have used to compute the moments and the assigned parameters in our paper. It also provides more details on the computational routines used to solve our model and presents further robustness checks that were not included in the main text.

1 Income process

We use data from the *Single-year Family Files* of the Panel Study of Income Dynamics to estimate the parameters of the income process. Starting in 1999, the PSID waves are released at a biennial frequency. It's important to note that the year of release is not the year for which the data have been collected. For example, the PSID wave for 1999 would actually report data for 1998. For each wave between 1999 and 2007, we use the nationally representative SRC sample of the PSID to create measures of annual income stemming from labor or transfers. For observations that are not reported at the annual level, we convert the income measures using additional information on the frequency with which each type of income is obtained. We drop observations for which the age, education and marital status of the head is not reported.

After constructing the income variables of interest, we use Taxsim to compute the tax liabilities of each household in our sample. We provide Taxsim with the state of residence for each household in order to get a measure of state taxes that households are liable for. Married household heads are assigned the joint filing status. Unmarried households that have children residing within the household are assigned a household head status in Taxsim. All the other households are considered filing separately. To compute household income before taxes, we specify in Taxsim the wages (net of pension contributions), social security income, taxable pension income, unemployment compensation, workers' compensation, supplemental social security, other welfare, child support, and transfers from relatives for both the head of the household and his/her spouse if married. After running Taxsim, we get a measure of state and federal taxes that we subtract from our measure of total pre-tax household income. We then adjust after-tax income for inflation using the CPI index, where the base year is 1998. Lastly, we apply the OECD equivalence scale to get our final measure of per capita deflated disposable income.

Next, we use the *Cross-year Individual File* of the PSID to generate an unique ID that combines the personal number of each individual ever surveyed in the PSID with the family identifier of the household in which this individual resides in a given year. We need to create this new ID because the *Single-year Family Files* used to compute our measure of disposable income don't record an identifier that could help us track the same household across different PSID waves. For each wave between 1999 and 2007, we keep only the household heads that are still living with their family and merge this data with data on disposable income that we computed earlier. We drop the observations for which the age of the head is inconsistent between the *Cross-year Individual File* and the *Single-year Family Files*. We also drop observations that are splitoffs, namely observations for which the household head moved in or out during the year of survey. We do this in order to make sure that any changes in average household income don't come from households in which the head has changed between two years (i.e., we want to avoid cases in which family income dropped by 50% between two years just because the former head left the family after a divorce). We then pool all the waves together and drop the households that are inconsistent about the age and education of their head across consecutive PSID waves. This data filter guarantees that we keep track of the same household head and their household unit across time (e.g., age of head does not grow by more or less than two years between two survey waves). The result of pooling the waves together is a panel dataset that is used to compute the parameters of the income process required for our model.

We restrict our sample to households for which the household head is between 25 and 85 years old. We also ensure that our panel is balanced by focusing only on the households for which we have all observations for the 1999-2007 PSID waves. Lastly, when computing the variance of log income, the 2-and-4-year autocovariances, as well as the standard deviation and kurtosis of biannual income growth, we exclude the residuals for observations that we qualify as outliers. In this sense, we compute for each household the difference between observed log income and mean log income across time ($\Delta_{it} = \log\left(\text{income}\right)_{it} - \overline{\log\left(\text{income}\right)_{i}}$). When estimating the variance and the autocovariance we exclude the observations that are in the top and bottom percentile of the distribution of these differences (Δ_{it}).

2 Moments

Data from the Survey of Consumer Finances (SCF) is used to compute the majority of moments to which we calibrate our model. The SCF data comes in two formats: (i) *Full Public Data Set* that contains all variables except for the private ones that could identify the reporting household, (ii) *Summary Extract Public Data* which consists of a sample of summary measures computed by the FRB Staff. We combine these two datasets for the purpose of this study, because some data needed to compute the moments used for calibration is not readily available in the *Summary Extract Public Data*.

Most of our moments have income as a common denominator. We compute disposable income using SCF data in a similar way we did for PSID data. Namely, we specify various income components from the Summary Extract Public Data file for each wave and run Taxsim in order to calculate the tax liabilities. We split household wage and business income in two unequal parts, awarding 75% of total income to the head of the household. We then assume that each of earners allocates 6% of their income towards pension contributions. Next, we use the reported measures for social security benefits as well as transfers to proxy for other incomes that the household receives every year. We run Taxsim and subtract the resulting federal tax from pre-tax income.[1] The final output of this routine is a measure of after-tax income. We also use the *Full Public Data Set* to compute the rent paid by households each year. Rent and the after tax income are then equivalized using the OECD equivalence scales.

Aside from income, our moments rely on some measure of household wealth (e.g., net worth, liquid assets, mortgage debt, etc.). Our measure of housing assets is based on the value of the primary residence alone. Liquid assets are the sum of checking accounts, saving accounts, money market deposits, money market mutual fund accounts, certificates of deposit, directly held pooled investment funds, saving bonds, stocks, other residential real estate, nonresidential real estate net of mortgages, and other non-financial assets.[2] Mortgage debt is the sum of all mortgages on the primary residence (including home equity loans and outstanding balances on home equity lines of credit).[3] Liquid debt is the sum of the credit card balances and other mortgage debt on secondary real estate. Housing net worth is the value of the primary residence less any mortgage debt on this residence. Each of these wealth statistics is equivalized using OECD equivalence scales. Lastly, we exclude households whose head is not aged between 25 and 85 years when computing the moments using SCF data, as well as households who have net worth above the 80th percentile of the net worth distribution.

3 Housing turnover

We follow Berger and Vavra (2015) when it comes to computing the share of homes that have been transacted in the total number of homes available in the U.S. Data on existing home sales comes from the National Association of Realtors (NAR). Data on the total housing stock in a given year is sourced from the U.S. Bureau of the Census. We divide total number of homes that have been transacted by total housing stock to obtain a measure of housing turnover.

4 Secondary homes vs. primary residences

We chose to include other residential real estate in our measure of liquid assets, as opposed to having it be part of the housing assets, for the following reasons. First, as described in the

[1]State taxes are not computed by Taxsim because SCF does not report in the publicly available files the state in which each household resides at the time of survey.

[2]Other non-financial assets include oil and gas leases, futures contracts, royalties, proceeds from lawsuits or estates in settlement, and loans made to others.

[3]We include the second lien loans in our calculations of mortgage debt in order to ensure that we capture all household debt. In the 2001 SCF data, second lien loans are held by households whose income is 1.5 times larger than the average income. These households also have a higher wealth relative to aggregate income when compared to all others (2.05 for those who have second liens vs. 1.45 for all households in the sample).

main text, very few homeowners own a secondary home in the data and adding the secondary home to the housing wealth would change our moments by very little.[4] Second, and more importantly, according to the breakdown provided by National Association of Realtors (2014) secondary homes are very different in nature from primary residences:

- A buyer of a secondary home intends to keep it for a median duration of 5-to-6 years, as opposed to a duration of 8 years for buyers of primary residences. As a consequence, the market for secondary residences should have a higher turnover rate.

- On average only 30 percent of secondary home sales are used as vacation properties, while the rest of the secondary home sales are used as investment properties that earn returns. Moreover, at least a quarter of the homeowners who own vacation homes purchase them with an intention to rent to others.

- Between 40 and 50 percent of secondary home owners use cash when purchasing such properties, while the ones who use mortgages finance a lower share of the purchase via debt than homeowners getting a mortgage to buy a primary residence. Furthermore, homeowners that own a secondary home also have higher incomes on average than owners of primary residences.

The arguments above support our conjecture that the market for secondary homes is different from the market for primary residences, and likely to be more liquid.

5 Returns on housing and renting

In our benchmark calibration, we obtain a rental rate of 3.3%, close to the population weighted national average net rental yield of 4.3% reported by Eisfeldt and Demers (2015). These authors use data on large institutional investors to compute yields for rental investments. While institutional investors have different investment horizons compared to U.S. households and are subject to different borrowing constraints, it is encouraging to see that our rental rate is within the 2-10 percent interval obtained by Eisfeldt and Demers (2015).

Our rent to price ratio is also consistent with to the annual rent-price ratio in Sommer et al. (2013), Hedlund (2016) and Hedlund and Garriga (2016). Given that our mortgage rate is set at 2.5% , it implies a 32 percent premium (3.3/2.5) between renting and owning (without taking into account the transaction and refinancing costs which would further bring the premium down). This roughly matches the evidence in Trulia (2016), showing that the premium between renting and owning has varied in the interval of 34-40 percent between 2012 and 2016.

[4] In the 2001 SCF wave, only 11.3 percent of all families owned other residential properties aside from the primary residence, which represented 4.7 percent of the value of total assets.

6 Forbearance programs

The academic literature on the magnitude of forbearance programs in the U.S. mortgage market is rather scarce. While most commercial banks have guidelines implementing such programs (as well as Fannie Mae and Freddie Mac), there is little evidence on how often these programs are used and how successful they are in curing delinquencies, especially for loan renegotiations prior to 2006. The scarce existing evidence presented below does however point to the fact that forbearance programs are rarely offered to delinquent borrowers.

Adelino et al. (2013) use a contract-change algorithm that compares the properties of a given loan across time to infer whether the loan was modified during 2006-2011. They find that in 2006 only 10000 loans per quarter were modifications (i.e. payment, interest or term have changed) in a dataset that covers 60 percent of the U.S. mortgage market. While their algorithm cannot identify forbearance programs clearly, the low number of modifications per quarter indicate that delinquent mortgage loans were rarely renegotiated before the crisis. They also explore changes in the payment size after loan modifications and find that in 2006, households that received payment reductions after modifications were getting only a 10 percent cut to their payments as a result of the modification (on average). This again highlights that very few households are getting significant payment relief even when they become delinquent on their loans and enter into renegotiation programs with loan providers/servicers.

Agarwal et al. (2011) is another study that relies on loan level data to track loan renegotiations. In contrast to Adelino et al. (2013), Agarwal et al. (2011) observe the renegotiation status reported by loan providers in their dataset. Hence they don't need to rely on any algorithm to track mortgages across time and give a more accurate breakdown of renegotiation arrangements by type. The downside is that they focus only on 2008-2009, a period when a lot of government mortgage assistance programs started being implemented. The big take-away for our purposes is that principal deferral were relatively rare (3 percent of all loan modifications). Term extensions were also a rare renegotiation status (about 15 percent of all loan modifications).

Moreover, when tracking renegotiated loans across time, they find that only 2.6 percent of delinquencies entered a repayment plan within the first 6 months after the start of their delinquency. After six more months, more than half of borrowers in their sample were in liquidation mode, about 23 percent of loans have been renegotiated, and about 25 percent had no action. This highlights, that even during a period with a high number of delinquencies, loan renegotiations were an infrequent solution.

The last piece of evidence that we reviewed comes from Orr et al. (2011). These authors study the effects of the Homeowners Emergency Mortgage Assistance Program (HEMAP), a Pennsylvania state policy that provides forebearance programs to borrowers who become delinquent on their mortgages due to short unemployment spells or are subject to financial hardship beyond their control. HEMAP offers temporary loans so that households in need can keep current on their original loan until they get enough income to repay the HEMAP transition loan. This policy was introduced in 1983, and until 2009, 183 thousand borrowers applied for such a loan. Out of these, only 23 percent were approved for HEMAP. Around 80 percent of loan recipients have retained ownership of their residences, repaying the HEMAP loan and getting current on their old mortgage. Despite its success and being self-sustaining (borrowers pay interest on HEMAP loans), the state of Pennsylvania canceled this forbearance program in 2011.

7 Interest rate wedge

We use data from three sources to set our real interest rate on mortgage debt. First, we obtain the average 30-year fixed mortgage rate for 2001 from the FRED database (6.97%). Second, we multiply this rate by (1-0.2391), where 23.91% is the average 2001 marginal tax subsidy on mortgage interest paid as reported by TAXSIM (federal plus state tax rate subsidy, see `http://users.nber.org/~taxsim/marginal-tax-rates/at.html`). Lastly, we subtract the 2001 percent change in CPI (2.8%) from the rate obtained above, to arrive at $6.97 \cdot (1 - 0.2391) - 2.8 = 2.5\%$. We obtained the average annual change in CPI from the FRED database (Consumer Price Index for All Urban Consumers: All Items).

We set our real return on the liquid asset r_l based on the evidence in Davis et al. (2006). They report an after-tax risk-free return of 2.9% for 2001 (see their Table 1, row 6), from which we subtract the 2001 change in CPI to arrive at $r_l = 0.1\%$

8 Equity extraction

Bhutta and Keys (2016) report the fraction of extractors in their sample (see their Table 1, Column 5). Note that their sample is a subsample of the total population, as they include in their analysis only homeowners that haven't moved and that have a mortgage. In order to infer the share of this subsample in the sample of all homeowners, we employ data from the SCF and the PSID. First, we determine the fraction of homeowners who have a mortgage in the 2001 SCF data. Second, we use data from the PSID to determine the fraction of homeowners that did not move in 2001 (based on the 2003 wave, using the answers to the question in which homeowners are asked if they have moved during the previous two years). Lastly, we multiply these two ratios to the share of extractors in Bhutta and Keys (2016), to arrive at the fraction of homeowners that extracted housing equity.

9 Cost of refinancing and closing costs

The average closing costs typically range between 2 to 8 percent of the purchase price (see http://michaelbluejay.com/house/interest.html). These costs include lender-related costs, as well as property-associated expenses. Some of these costs are fixed amounts (e.g., appraisal, credit report, survey and title fees), while others are expressed in percentage terms of the original sale price or mortgage loan size (e.g. origination and insurance fees, property taxes, etc.).

Using the calculator in the link above, we can get a sense of how accurate our estimates of closing and refi costs are. The median sale price of houses sold in 2001 was 173,100 USD according to the FRED database. This would imply lender related costs of 5.2% of sale price or about 33% of average income (computed based on the 2001 SCF data used in our calibration exercise), as well as property related costs of 2.5% of the house price (or 16% of mean 2001 income). These numbers are somewhat higher than the ones assumed by Agarwal et al. (2013) (2000 USD + 1% of mortgage amount closing costs). Our estimates align closer to the ones in Agarwal et al. (2013) and are at the lower end of the parameter values and estimates provided by the Federal Reserve Board (see https://www.federalreserve.gov/pubs/refinancings//#cost).

10 Adjustment of the payment-to-income constraint

In this subsection we describe the adjustment we performed on the payment-to-income ratio reported in Greenwald (2015) in order to account for the fact that mortgage payments in our model do no include tax deductions and inflation.

Let B be the dollar amount of the mortgage that the household has on its home. Their mortgage payment is hence:

$$(1 + i - \gamma) B$$

which we restrict to be consistent with the median PTI for 2001 reported in Greenwald (2015):

$$(1 + i - \gamma) B \leq 0.35PY$$

or,

$$B \leq \frac{0.35}{1 + i - \gamma} PY$$

where PY stands for nominal income. The dollar payments in our model are actually

$$(1 + r_m - \gamma) b$$

because they are after tax and after inflation. So the PTI in the model is

$$\frac{(1 + r_m - \gamma)\, B}{(1 - \tau)\, PY}$$

where the denominator is the after-tax or disposable income.

Since

$$\frac{B}{PY} \leq \frac{0.35}{1 + i - \gamma}$$

the PTI in our model has to be

$$\frac{1 + r_m - \gamma}{1 - \tau} \frac{0.35}{1 + i - \gamma}$$

For our specific numbers: $\gamma = 0.969$, $r_m = 0.025$, $i = 0.0697$, and $\tau = 0.239$ we have a PTI of 0.2558.

However, mortgage payments are only about 60% of all debt payments in the data. Banks compute the total PTI when a household applies for a mortgage loan, in order to reflect all outstanding debt. So, assuming everyone has identical non-mortgage payments equal to a fraction λ of their income, then the mortgage payment and non-mortgage payments are limited by

$$(1 + i - \gamma)\, B + \lambda PY \leq 0.35 PY$$

therefore, the mortgage payments are

$$(1 + i - \gamma)\, B \leq (0.35 - \lambda)\, PY.$$

λ is 0.0767 in the data. This value is based on our SCF sample and is computed as the average non-mortgage payment to average income, where non-mortgage payments include revolving debt payments (other than mortgage debt) and consumer debt payments. Therefore the relevant PTI for us is not 0.35 but rather 0.2733.[5]

So the PTI we use to calibrate our model is

$$\frac{1 + r_m - \gamma}{1 - \tau} \frac{0.35 - \lambda}{1 + i - \gamma}$$

or

$$\frac{1 + 0.025 - 0.969}{1 - 0.239} \frac{0.35 - 0.0767}{1 + 0.0697 - 0.969} = 0.20$$

[5] An alternative would be to assume that mortgage payments are proportional to non-mortgage payments, in which case we have that mortgage debt is about 60% of total debt and so we should use an effective PTI of 0.6*0.35 or 0.21. Our approach is clearly on the conservative side.

References

Adelino, Manuel, Kristopher Gerardi, and Paul S Willen, "Why don't lenders renegotiate more home mortgages? Redefaults, self-cures and securitization," *Journal of Monetary Economics*, 2013, *60* (7), 835–853.

Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, and Douglas D Evanoff, "The role of securitization in mortgage renegotiation," *Journal of Financial Economics*, 2011, *102* (3), 559–578.

_ , John C Driscoll, and David I Laibson, "Optimal Mortgage Refinancing: A Closed-Form Solution," *Journal of Money, Credit and Banking*, 2013, *45* (4), 591–622.

Berger, David and Joseph Vavra, "Consumption Dynamics During Recessions," *Econometrica*, 2015, *83* (1), 101–154.

Bhutta, Neil and Benjamin J. Keys, "Interest Rates and Equity Extraction during the Housing Boom," *American Economic Review*, July 2016, *106* (7), 1742–74.

Davis, Steven J., Felix Kubler, and Paul Willen, "Borrowing Costs and the Demand for Equity over the Life Cycle," *The Review of Economics and Statistics*, May 2006, *88* (2), 348–362.

Eisfeldt, Andrea and Andrew Demers, "Rental Yields and HPA: The Returns to Single Family Rentals," Technical Report, National Bureau of Economic Research 2015.

Greenwald, Daniel L., "Mortgage Refinancing and the Macroeconomy," *Working Paper*, 2015.

Hedlund, Aaron, "Failure to Launch: Housing, Debt Overhang, and the Inflation Option During the Great Recession," 2016.

_ and Carlos Garriga, "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession," 2016.

National Association of Realtors, "Investment And Vacation Home Buyers Survey," 2014.

Orr, James, John Sporn, Joseph S Tracy, and Joe Huang, "Help for Unemployed Borrowers: Lessons from the Pennsylvania Homeowners' Emergency Mortgage Assistance Program," *Current issues in Economics and Finance*, 2011, *17* (2).

Sommer, Kamila, Paul Sullivan, and Randal Verbrugge, "The equilibrium effect of fundamentals on house prices and rents," *Journal of Monetary Economics*, 2013, *60* (7), 854–870.

Trulia, "Rent vs. Buy: Be Haunted by Prices, Not Rates," https://www.trulia.com/blog/trends/rent-vs-buy-oct-2016/ 2016.