

Mortgage Market Institutions and Housing Market Outcomes: A Dynamic General Equilibrium Approach*

Edward Kung[†]

October 9, 2014

Abstract

I develop a general equilibrium model of housing and mortgage markets where house prices, mortgage interest rates, leverage ratios, and contract availability are all endogenous. A housing demand curve is created by solving the dynamic decision problem of potential buyers whose housing decisions are affected by the types of mortgage contracts available to them. Equilibrium house prices adjust so that demand clears with supply created by existing owners who sell their homes. Mortgages are provided by lenders who set interest rates such that expected returns equal the opportunity cost of funds. The model is calibrated using data from Los Angeles from 2003 to 2010, and is used to study the impact of the disappearing market for non-agency mortgages in 2008, as well as various government responses.

*Preliminary and incomplete. Please do not cite without author's permission. I thank Patrick Bayer for outstanding advice as well as for providing the DataQuick data. I thank Stuart Gabriel for generously providing the BlackBox data. Additional thanks to Maurizio Mazzocco, John Asker, Till von Wachter, Simon Board, Peter Arcidiacono, Jimmy Roberts, and Alan Collard-Wexler for helpful comments.

[†]Department of Economics, UCLA. E-mail: ekung@econ.ucla.edu.

1 Introduction

For most people, especially younger or poorer households with little wealth, the size and quality of a house they are able to purchase (indeed if they are able to purchase at all) is greatly restricted by the amount of down-payment they can afford and the amount of mortgage financing they are able to obtain. The price and availability of mortgage credit should therefore have strong impacts on housing market outcomes. Because of this, mortgage market interventions are the primary policy lever through which the U.S. government seeks to influence housing market outcomes. It is therefore important to understand the interrelationship between mortgage and housing markets—not only for the policy relevance—but also because it speaks to broader questions about how credit markets interact with real markets.

The goal of this paper is to study the role of mortgage credit availability in determining housing market outcomes, and vice versa, using a dynamic general equilibrium framework with rational utility-maximizing households. In the model, demand for housing is created by potential buyers, who decide which homes to purchase by solving a dynamic optimization problem that takes into account current and future house prices, and the current and future set of available mortgage contracts. Supply is created by existing homeowners who exit the market, and equilibrium prices are such that housing demand and supply are equal in each state of the world. Mortgages are provided by competitive lenders who set interest rates such that expected returns are equal to their opportunity cost of funds. The availability of certain mortgage contracts may be restricted due to government restrictions, and others may be unavailable if the lenders are unable to make positive net returns.

By modeling both the housing and mortgage markets, house prices, mortgage interest rates, contract availability, and leverage ratios are all fully endogenous in the model. It is important that both markets are endogenous because we are primarily interested in counterfactuals that change the set of mortgage contracts available (this is indeed how government policy towards the mortgage market is handled). When the set of available mortgage con-

tracts changes, the house price process will change—but when house prices change, so too will the set of available mortgage contracts and the interest rates at which they are available.

Part of the challenge in writing down models of the housing market is that housing is a very difficult asset to model. Houses are a highly illiquid asset purchased in a decentralized market by small individuals with low wealth, borrowing constraints, and heterogeneous preferences over the “dividend” that a unit of housing pays.¹ This is already a great departure from standard asset pricing models. To make matters more complicated, housing units themselves can be both vertically and horizontally differentiated, making it difficult to speak of a single, thick market for housing.

I begin by developing a dynamic general equilibrium model of housing that incorporates most of the above features. In the model, there is a housing market with two vertically differentiated housing types, each with fixed supply. Each unit of housing is owned by a single household, or homeowner, and each homeowner has a mortgage contract. The dynamic decision problem of the homeowner follows in the same vein of Campbell and Cocco (2003, 2014). In each period, homeowners can either be hit with a moving shock or not. If they are not hit by a moving shock, they have the option of continuing in their current mortgage by making the required payment, or they can refinance. After the refinance decision, they solve an intertemporal consumption-savings problem. Homeowners hit by a moving shock will either sell their home or default on their mortgage, which results in foreclosure. The default, refinance, and consumption-savings decisions are all affected by house prices, by the types of contracts available in the mortgage market, and by the features of the owner’s own mortgage contract.

At the same time as homeowners are selling their homes, potential buyers

¹In a purely financial model of housing, the dividend that a unit of housing pays is the rental price of an equivalent unit. However, there can be many non-pecuniary benefits to owning a particular unit of housing, and different households can have very different preferences over said non-pecuniary benefits. Non-pecuniary benefits to owning housing are not arbitrated out of the price because of the illiquidity of housing and the small size of the market participants.

are entering the market. The potential buyers are heterogeneous in their income, wealth, and preference for housing. Because the potential buyers have limited wealth, (some of them) need to borrow in order to purchase a home. They cannot borrow freely, and must instead use mortgage contracts where the debt is collateralized against their homes. Different types of mortgage contracts are offered, each by a risk-neutral competitive lender who takes house prices as given and offers mortgage contracts that generate an expected return equal to the opportunity cost of funds. Given the set of offered contracts, potential buyers make their home purchasing decisions to maximize their expected utility as future homeowners. House prices are in equilibrium when the demand for homes from potential buyers equals the supply put on the market by existing owners. The distribution of buyer heterogeneity can change over time, which introduces stochastic variability to house prices. Importantly, the demand for homes from potential buyers is affected by the kinds of mortgage contracts offered by lenders, and the kinds of mortgage contracts offered by lenders depends on housing market conditions.

The model is calibrated using housing and mortgage market data from Los Angeles for the period 2003 through 2010. The model is then used to study the effect of non-agency mortgage credit availability.² Non-agency mortgages are modeled as a different kind of contract compared to agency loans, with the key difference being that agency loans carry no default risk but are restricted in size. The period of time 2003-2010 is of special interest because non-agency mortgage financing became unavailable in 2008 (see Figure 1, which shows that securitization of non-agency mortgages all but disappeared in 2008). I use the model to ask: “What would have happened to the housing market

²“Agency” loans refer to loans that are securitized by one of the three government sponsored enterprises (GSEs): Freddie Mac, Fannie Mae and Ginnie Mae. The GSEs carry either explicit or implicit guarantees on their credit obligations, and therefore agency loans are usually available at lower interest rates than non-agency loans. However, in order for a loan to be securitizable by the GSEs, they have to meet certain regulatory guidelines; for example, the size of the loan needs to be below the “conforming loan limit”—a level set by the Federal Housing Finance Agency. Because of the limitations on agency loans, the presence of a non-agency market greatly increases the set of available contracts. Generally speaking, if a borrower wished to obtain a large loan, either in terms of size or in terms of leverage ratio, the borrower would have to go through the non-agency market.

in Los Angeles if non-agency mortgage financing did not disappear in 2008?” I find that the disappearance of non-agency mortgage financing had a very large effect on house prices, on the order of a 30% increase. The model is rich enough to take into account the endogeneity between non-agency availability and house prices, and I show that even if the non-agency market did not exogenously disappear in 2008, some high leverage non-agency contracts would continue to be unavailable, due to expected house price declines.

I then use the model to study the effectiveness of various government responses to the housing crisis in 2008. There are two main responses to consider. First, conforming loan limits were increased in various cities in 2008 as part of the Economic Stimulus Act of 2008. In Los Angeles, the nominal conforming loan limit went from \$417,000 in 2007 to \$729,750 in 2008. Second, interest rates were reduced. The interest rate on a 10-year treasury bond was 4.6% in 2007 and 3.6% in 2008.

I find that in each case, the government response had non-trivial effects. Increasing the conforming loan limit increased the price of high-valued homes in 2008 by about 11 percent, and increased the price of low-valued homes by about 6 percent. The differential effect between high versus low-valued homes was about 5 percent, which is consistent with difference-in-difference estimates of the same policy effect reported in Kung (2014). In 2009 and 2010, the effect of the higher conforming-loan limits were still positive, but smaller than they were in 2008. This is also expected as the absolute level of house prices decreased in 2009, so the conforming loan limit was binding in fewer instances.

In contrast, reducing interest rates had a larger effect in 2009 and 2010 than in 2008. Interest rates therefore appear to have a larger effect when house prices are low than when they are high. When house prices are high, the *availability* of mortgage credit and the ability for borrowers to overcome down-payment constraints becomes a relatively more important consideration than changes in cash-flow caused by reduced interest rates. On the other hand, when prices are low, more buyers are able to overcome down-payment constraints, and so cash-flow considerations become a relatively more important driver of

aggregate house prices.

Related Literature

The relationship between housing markets and mortgage markets is a topic that has attracted renewed interest in the aftermath of the financial and housing crisis in 2008. A number of papers, such as Himmelberg et al. (2005); Mian and Sufi (2009); Glaeser et al. (2010); Favara and Imbs (2010); Adelino et al. (2014); Kung (2014), have used reduced-form and descriptive methods to demonstrate a variety of ways in which mortgage markets impact housing markets. The reduced form approach is valuable and robust to the stronger assumptions that model-based approaches require, but is unable to speak to counterfactuals which are far from the data or from the natural experiment being used.

In terms of model-based approaches, a number of important papers before this one have considered equilibrium models of housing markets impacted by mortgage markets. Ortalo-Magné and Rady (2006) study the role of income shocks and credit constraints; Favilukis et al. (2010) study the macroeconomic implications of housing wealth and limited risk-sharing; Landvoigt et al. (2014) study an assignment model of housing which can explain cross-sectional variation in capital returns. In these models, mortgage markets are either treated as exogenous, or else the representation of the mortgage market is highly stylized, with borrowing coming in the form of single-period bonds at exogenous interest rate or collateral constraint. A number of papers have also considered equilibrium mortgage markets, including Campbell and Cocco (2014), who study a model of default with endogenous mortgage contracts, and Corbae and Quintin (2014) who study the role of leverage in default. In these papers, the housing market and house price process is taken as exogenous.

To my knowledge, this is the first paper to fully endogenize both housing and mortgage market conditions, while simultaneously allowing for specific details of mortgage contracts and mortgage market institutions to be studied. As I discussed above, the endogeneity of both housing and mortgage markets is not simply an intellectual exercise with no practical applications. General

equilibrium responses from either market can be very important, depending on the counterfactual studied. The main contribution of the paper is therefore the development of a model that allows the study of a wide variety of counterfactuals in a setting where both housing and mortgage markets can respond to the counterfactual.

The goal of this paper is not to study optimal mortgage design or optimal housing policy. However, the model can be easily extended to study the effects of introducing a variety of alternative mortgage products or government policies. Therefore, the model can be used to study insights into optimal mortgage design, as in Piskorski and Tchisty (2010), in an empirically realistic setting. One counterfactual of interest may be, for example, the effect of introducing mortgage indexed to house prices, as discussed in Shiller (2008); Mian and Sufi (2014). Such extensions of the model are left to future research.

The rest of the paper proceeds as follows. In section 2, I set up the model, which builds upon the dynamic model of Campbell and Cocco (2014) and closes it with a computationally tractable demand system. In this section, I leave many features of the model as abstract objects, in order to highlight the modular design of the model. In section 3, I discuss the details in the implementation of the model, making concrete any areas of abstraction from section 2. In section 4, I discuss the Los Angeles housing and mortgage market to which the model will be calibrated. In section 5, I discuss the details of the calibration and discuss how well the model can fit the data. In section 6, I study the role of non-agency mortgage credit availability and the effectiveness of government response. Section 7 concludes.

2 Model

2.1 Preliminaries

Time is discrete. There is a housing market with two house types, $h = 0$ and $h = 1$, which can be thought of as low and high quality homes, respectively. There is a fixed stock μ of each house type, for a total of 2μ housing units alto-

gether. Each unit of housing can be occupied by one and only one household, and each household can occupy only one unit of housing. Let s_t be an abstract aggregate state variable (which I will specify in more detail later). The price of house type h in state s_t is given by $p_h(s_t)$. For now, no restrictions are placed on the aggregate state variable except that it evolves over time according to a first-order Markov process.

2.2 Mortgage contracts

A mortgage contract is represented by the vector $z_t = (a_t, r_t, b_t, m_t)$, which includes the mortgage's current age a_t , interest rate r_t , remaining balance b_t , and product type m_t . There are M product types, including $m = 0$ which denotes "no mortgage". A mortgage's product type determines its behavior outside of differences in interest rate and balance. In addition to the variables in z_t , a mortgage's behavior may also depend on the type of house it is collateralized with, h , and the current aggregate state s_t . Mortgage behavior is defined by three functions. First, let $pay_h(z_t, s_t)$ define the payment required to stay current on the mortgage, when the mortgage is described by z_t , the collateral type is h , and the current state is s_t . Second, let $\psi_h(z_t, s_t)$ define the amount that the lender is able to recover in the event of a default. Third, let $\zeta_h(z_t, s_t, s_{t+1})$ define the transition rule for potentially time-varying mortgage characteristics. The function ζ is defined such that $z_{t+1} = \zeta_h(z_t, s_t, s_{t+1})$ holds.³

2.3 Homeowners

Homeowners are modeled in a vein similar to Campbell and Cocco (2003, 2014). The key differences are that in this model, homeowners are infinitely lived and are hit with exogenous moving shocks, and that there is a reduced emphasis on modeling income and inflation risk.

Homeowners are infinitely lived but have a probability λ of having to move out of the housing market each period. Homeowners do not re-enter the hous-

³I give examples of pay_h , ζ_h , and ψ_h for various kinds of contracts in the Appendix.

ing market after they move, so moving is treated as a terminal state.⁴ Homeowners care about flows of consumption and about their final wealth at the time of a move. If the amount consumed in a period is c , the flow utility received is $u(\theta^h c)$, where h is the type of house that the homeowner lives in. $u(c)$ is a CRRA utility function with risk-aversion parameter γ . $\theta > 1$ is a parameter that determines the relative attractiveness of living in high vs. low quality homes. If the amount of wealth at the time of moving is w , then the terminal utility received is $\beta u(c)$, where β is simply a scale parameter. Homeowners are expected utility maximizers and discount future utility flows using discount parameter $\delta < 1$.

Moving

Homeowners enter each period being described by their income y_i (which is constant over time), their liquid wealth w_{it} , their mortgage contract z_{it} , and the type of house they own h_i . For notational convenience, let $x_{it} = (y_i, w_{it}, h_i)$ be the characteristics of the homeowner separate from their mortgage contract.

If the homeowner is required to move, it can either pay off its remaining mortgage balance and sell the house, or it can default (in which case the lender repossesses and sells the house). If the homeowner chooses to sell, it simultaneously pays off its remaining mortgage balance b_{it} and receives the current price of the home, $p_{h_i}(s_t)$. The homeowner's utility over final wealth

⁴The requirement to move out of the housing market can most easily be thought of as death of the homeowner, but can also be thought of as an exogenous job relocation shock. Within-market moves are not considered, although this only matters to the extent that within-market movers make large trades in terms of housing value (within-market moves between homes of the same value create no net supply or demand in the market). The extent to which such within-market moves is empirically important is explored in the Appendix. The rationale for ignoring within-market moves is that doing so significantly reduces the state space of the model. When within-market moves are considered, the state space needs to include the distribution of wealth of existing owners, which is highly history dependent. In section 3, I discuss a “patch” to the model which proxies for the effect of within-market movers—but the extent to which this patch will be affected by general equilibrium considerations under a counterfactual exercise is unknown. Understanding the general equilibrium effects of homeowners trading up and trading down within a housing market is admittedly important, but left to future research.

at the time of the move is therefore:

$$V^{sell}(x_{it}, z_{it}, s_t) = \beta u(y_i + w_{it} + p_{h_i}(s_t) - b_{it})$$

On the other hand, if the homeowner chooses to default, it does not have to pay off its existing mortgage balance, but it also forfeits any potential proceeds from the sale of the house. In addition, the homeowner must pay a default cost, c_D . The utility to defaulting is therefore:

$$V^{def}(x_{it}, z_{it}, s_t, \epsilon_{it}^D) = \beta u(y_i + w_{it} - c_D)$$

If the owner is behaving optimally, it will default if and only if:

$$V^{def}(x_{it}, z_{it}, s_t) > V^{sell}(x_{it}, s_t)$$

Finally, we write:

$$V^{move}(x_{it}, z_{it}, s_t, \epsilon_{it}^D) = \max \{V^{sell}(x_{it}, z_{it}, s_t), V^{def}(x_{it}, z_{it}, s_t)\}$$

and we let $\tau_{it} = \tau(x_{it}, z_{it}, s_t)$ be an indicator function for whether the homeowner chooses to default.

Staying

If the homeowner is not required to move, then it can either continue in its current mortgage contract by making the required payment, or it can refinance into a new contract. After that, the homeowner decides how much to consume and how much to save at risk-free rate rfr_t , which is part of s_t .

If the homeowner chooses to continue in its current mortgage, then it must make the required payment. The homeowner's budget constraint is therefore:

$$c_{it} + \frac{1}{1 + rfr_t} w_{it+1} = y_i + w_{it} - pay_{h_i}(z_{it}, s_t)$$

and the contract it carries into next period is:

$$z_{it+1} = \zeta_{h_i}(z_{it}, s_t, s_{t+1})$$

If the homeowner instead chooses to refinance, it first pays a financial cost of refinancing c_R . Then, the homeowner simultaneously pays off the existing mortgage balance b_{it} and chooses a new product type m' and new loan amount $b' \leq \bar{b}_{m'}(x_{it}, s_t)$. The new loan amount must be less than $\bar{b}_m(x_{it}, s_t)$, which is an equilibrium object that determines the maximum borrowing amount for given product type, borrower characteristics (including collateral type) and aggregate state.⁵

Let $r_m(b, x_{it}, s_t)$ be the equilibrium function that maps a new origination of product m and loan amount b to an interest rate, when the borrower's characteristics (including collateral type) are x_{it} and the aggregate state is s_t . The new interest rate when refinancing into mortgage m' , b' is therefore $r' = r_{m'}(b', x_{it}, s_t)$ and the new mortgage is described by vector $z' = (0, r', b', m')$. The budget constraint is therefore:

$$c_{it} + \frac{1}{1 + rfr_t} w_{it+1} = y_i + w_{it} - b_{it} + b' - pay_{h_i}(z', s_t) - c_R$$

and the next period contract is:

$$z_{it+1} = \zeta_{h_i}(z', s_t, s_{t+1})$$

Homeowner decision problem

Let $\rho_{it} = \rho(x_{it}, z_{it}, s_t)$ be a policy rule that determines whether or not a homeowner who does not have to move refinances. Let $m'_{it} = m'(x_{it}, z_{it}, s_t)$ and $b_{it} = b'(x_{it}, z_{it}, s_t)$ be the policy rules that determine the contract that a homeowner refinances into. Finally, let $c_{it} = c(x_{it}, z_{it}, s_t)$ be the policy rule

⁵Limitations on the maximum borrowing amount may arise, for example, from regulations prohibiting the origination of agency loans greater than the conforming loan limit. The maximum borrowing amount may also arise endogenously if lenders are not able to profitably originate large loans to certain borrowers in certain situations.

that determines consumption.

From the perspective of a homeowner not having to move at time t , the homeowner chooses policy rules d , ρ , m' , b' , and c to maximize:

$$u(\theta^{h_i} c_{it}) + E \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \left\{ (1-\lambda)^{s-t} [u(\theta^{h_i} c_{is})] + (1-\lambda)^{s-t-1} \lambda V_{is}^{move} \right\} \middle| x_{it}, z_{it}, s_t \right] \quad (1)$$

subject to the following constraints for all $s = t, \dots, \infty$

$$z'_{is} = (0, r_{m'_{is}}(b'_{is}, x_{is}, s_s), b'_{is}, m'_{is}) \quad (2)$$

$$z_{is+1} = \begin{cases} \zeta_{h_i}(z'_{is}, s_s, s_{s+1}) & \text{if } \rho_{is} = 1 \\ \zeta_{h_i}(z_{is}, s_s, s_{s+1}) & \text{if } \rho_{is} = 0 \end{cases} \quad (3)$$

$$pay_h(z'_{is}, s_s) \leq 0.5y_i \quad (4)$$

$$c_{is} + \frac{1}{1+rfr_s} w_{is+1} = \begin{cases} y_i + w_{is} - b_{is} + b'_{is} - pay_{h_i}(z'_{is}, s_s) - c_R & \text{if } \rho_{is} = 1 \\ y_i + w_{is} - pay_{h_i}(z_{is}, s_s) & \text{if } \rho_{is} = 0 \end{cases} \quad (5)$$

$$w_{is+1} \geq 0 \quad (6)$$

$$b'_{is} \leq \bar{b}_{m'_{is}}(x_{is}, s_s) \quad (7)$$

The constraints can be understood as follows. Constraints (2) and (3) are simply accounting relationships describing how the owner's mortgage changes over time given his decisions. Constraint (4) is an assumption about the mortgage market which restricts borrowing to less than 50% payment-to-income ratio. Constraint (5) is the budget constraint and constraint (6) is the no-borrowing constraint. Constraint (7) restricts the choice of new mortgage contracts to only those that are available.

Bellman equations and value functions

Equations (1)-(7) are a complete statement of the homeowner's optimization problem, but it is more easily understood as a Bellman equation. Let $V_{it}^{stay} = V^{stay}(x_{it}, z_{it}, s_t)$ be the maximum of (1). Then, the homeowner's optimization problem can equivalently be written:

$$V_{it}^{stay} = \max_{c, \rho, m', b'} u(\theta^{h_i} c) + \delta E \left[(1 - \lambda) V_{it+1}^{stay} + \lambda V_{it+1}^{move} \mid x_{it}, z_{it}, s_t \right] \quad (8)$$

subject to:

$$z' = (0, r_{m'}(b', x_{it}, s_t), b', m') \quad (9)$$

$$z_{it+1} = \begin{cases} \zeta_{h_i}(z', s_t, s_{t+1}) & \text{if } \rho = 1 \\ \zeta_{h_i}(z', s_t, s_{t+1}) & \text{if } \rho = 0 \end{cases} \quad (10)$$

$$pay_h(z', s_t) \leq 0.5y_i \quad (11)$$

$$c + \frac{1}{1 + rfr_t} w_{it+1} = \begin{cases} y_i + w_{it} - b_{it} + b' - pay_{h_i}(z', s_t) - c_R & \text{if } \rho = 1 \\ y_i + w_{it} - pay_{h_i}(z_{it}, s_t) & \text{if } \rho = 0 \end{cases} \quad (12)$$

$$w_{it+1} \geq 0 \quad (13)$$

$$b' \leq \bar{b}_{m'}(x_{it}, s_t) \quad (14)$$

Equation (8) is a contraction mapping, and therefore a unique solution for V^{stay} exists and can be found by iteratively computing equation (8) (see Stokey et al. (1989)).

2.4 Potential buyers

In each period, a mass 1 of potential buyers enters the market. The buyers are heterogeneous in three dimensions: their income y_i , which is constant over time, their initial wealth w_i , and the utility they would receive from *not* purchasing a home—in other words, their outside option v_i .⁶ Given current

⁶The outside option can be thought of as the net present value of utility from living in a different housing market, or from renting, which will be treated as coming from a separate housing stock.

aggregate conditions and given their individual heterogeneity, each potential buyer must decide whether or not to purchase a house (and which type to purchase) or to take their outside option. If a potential buyer decides not to buy a house, it receives a utility equal to v_i , exits the housing market, and does not re-enter.

It is easy to describe the decision problem of the potential buyer in terms of the decision problem of the homeowner. If the potential buyer decides to buy a house of type h , he chooses mortgage contract m , b and consumption c to maximize:

$$V_h^{buy}(y_i, w_i, s_t) = \max_{c, m, b} u(\theta^h c) + \delta E \left[(1 - \lambda) V_{it+1}^{stay} + \lambda V_{it+1}^{move} \middle| x_{it}, z_{it}, s_t \right] \quad (15)$$

subject to:

$$x_{it} = (y_i, w_i, h) \quad (16)$$

$$z = (0, r_m(b, x_{it}, s_t), b, m) \quad (17)$$

$$z_{it+1} = \zeta_h(z, s_t, s_{t+1}) \quad (18)$$

$$pay_h(z, s_t) \leq 0.5y_i \quad (19)$$

$$c + \frac{1}{1 + rfr_t} w_{it+1} = y_i + w_i - p_h(s_t) + b - pay_h(z, s_t) \quad (20)$$

$$w_{it+1} \geq 0 \quad (21)$$

$$b \leq \bar{b}_m(x_{it}, s_t) \quad (22)$$

As one can see, the decision problem of the buyer is very similar to the decision problem of the refinancing owner. The only additional feature of the buyer's problem is that he also must choose which type of house to buy, if any. Ignoring ties, the potential buyer will purchase a house of type h if and only if:

$$V_h^{buy}(y_i, w_i, s_t) = \max \left\{ V_0^{buy}(y_i, w_i, s_t), V_1^{buy}(y_i, w_i, s_t), v_i \right\} \quad (23)$$

2.5 Housing market demand and supply

Let $d_h(y_i, w_i, v_i, s_t)$ be an indicator function for whether a potential buyer described by (y_i, w_i, v_i) buys a house of type h in state s_t , as determined by equation (23). Now, let the distribution of potential buyer heterogeneity in state s_t be given by the probability density function $\Gamma(y_i, w_i, v_i; s_t)$.

The total demand for house type h in state s_t is therefore:

$$D_h(s_t) = \int_y \int_w \int_v d_h(y, w, v, s_t) \Gamma(y, w, v; s_t) dydw dv \quad (24)$$

The total mass of homes for sale of each type in each period is $\lambda\mu$. In order for the housing market to clear, demand for each house type must equal to the supply on the market. Therefore, in equilibrium, the following must hold:

$$D_h(s_t) = \lambda\mu \text{ for } h = 0, 1 \quad (25)$$

Equation (25) is the housing market clearing condition. It is a simple supply and demand equation from which house prices can be determined in each state. The decision problem of the buyer, as described in equations (15)-(22), clearly shows how mortgage market institutions, through the equilibrium objects r_m and \bar{b}_m can affect demand—and thereby house prices.

2.6 Lenders

Each individual mortgage contract is originated by a risk-neutral and competitive lender. Other than the mortgage contract, the lender is able to invest in single-period risk-free bonds with a return of $rfr_t + a$. $rfr_t + a$ can therefore be thought of as the opportunity cost of funds for a lender in period t . The parameter $a > 0$ represents the better investment opportunities that large institutional lenders may face compared to individual households. After originating a mortgage, the lender re-invests all receipts from that mortgage into single-period bonds at rate $rfr_t + a$. At the time of origination, the lender sets the mortgage interest rate such that the expected return on the mortgage

is equal to its opportunity cost of funds.

Let $\Pi_{it} = \Pi(x_{it}, z_{it}, s_t)$ be the net present value of expected receipts from a mortgage contract z_{it} . Given the decision rules of homeowners, we can write:

$$\Pi_{it} = \lambda \pi_{it}^{move} + (1 - \lambda) \left\{ \pi_{it}^{stay} + \frac{1}{1 + rfr_t + a} E[\Pi_{it+1} | x_{it}, z_{it}, s_t] \right\} \quad (26)$$

where:

$$\begin{aligned} \pi_{it}^{move} &= \tau_{it} \psi_{h_i}(z_{it}, s_t) + (1 - \tau_{it}) b_{it} \\ \pi_{it}^{stay} &= \rho_{it} b_{it} + (1 - \rho_{it}) \text{pay}_{h_i}(z_{it}, s_t) \\ z_{it+1} &= \begin{cases} (0, 0, 0, 0) & \text{if } \rho_{it} = 1 \\ \zeta_{h_i}(z_{it}, s_t, s_{t+1}) & \text{if } \rho_{it} = 0 \end{cases} \end{aligned}$$

For a new origination m, b , the probability of a refinance in the period of origination is zero, so we write:

$$\Pi^{orig}(x_{it}, m, b, s_t) = \text{pay}_{h_i}(z_{it}, s_t) + \frac{1}{1 + rfr_t + a} E[\Pi_{it+1} | x_{it}, z_{it}, s_t]$$

where

$$z_{it} = (0, r_m(b, x_{it}, s_t), b, m)$$

In equilibrium, the expected excess return of the mortgage contract over the opportunity cost of funds is zero, so:

$$\Pi^{orig}(x_{it}, m, b, s_t) - b = 0 \quad (27)$$

2.7 Competitive equilibrium

The economy is in competitive equilibrium when homeowners and potential buyers are behaving optimally, as described by equations (8) and (15), and the housing market and mortgage markets clear. The housing market clearing condition is given by equation (25) and the mortgage market clearing condition is given by equation (27). The equilibrium objects that must be solved for

are all the value functions and policy rules described above. However, all the value functions and policy rules can easily be computed as long as the following four objects are known: $p_h(s_t)$, $r_m(x_{it}, s_t)$, $V^{stay}(x_{it}, z_{it}, s_t, \epsilon_{it}^R)$ and $\Pi(x_{it}, z_{it}, s_t, \epsilon_{it}^R, \epsilon_{it}^D)$.

A competitive equilibrium can be found using an iterative procedure with three nests. In the inner nest, p_h and r_m are taken as given, and V^{stay} and Π are chosen to satisfy their respective Bellman equations. In the middle nest, p_h is given and r_m is chosen to satisfy the mortgage market clearing condition. In the outer nest, p_h is chosen to satisfy the housing market clearing condition. Each nest is described below:

Inner Nest

Note: p_h and r_m are given

1. Guess V_0^{stay} and Π_0 . The subscripts here indicate the step in the algorithm.
2. For $iter \geq 0$:
 - (a) Compute V_{iter+1}^{stay} by solving equations (8)-(14) using V_{iter}^{stay} on the RHS.
 - (b) Compute the optimal policy rules as a result of the solution in (a).
 - (c) Compute Π_{iter+1} using equation (26), with the policy rules from (b) and Π_{iter} on the RHS.
 - (d) Repeat until $V_{iter+1}^{stay} = V_{iter}^{stay}$ and $\Pi_{iter+1} = \Pi_{iter}$

Middle Nest

Note: p_h is given

1. Guess $r_{m,0}$. The subscripts here indicate the step in the algorithm.
2. For $iter \geq 0$

- (a) Compute the **inner nest** until V^{stay} and Π reach convergence, taking $r_{m,iter}$ as given
- (b) For each x_{it}, m, b, s_t , iterate over values of $r_{m,iter+1}(b, x_{it}, s_t)$ until $\Pi^{orig}(x_{it}, m, b, s_t) = b$ (using the values for Π computed in step (a))
- (c) Set $r_{m,iter+1}$ equal to the values computed in step (b)
- (d) Repeat until $r_{m,iter+1} = r_{m,iter}$

Outer Nest

1. Guess $p_{h,0}$. The subscripts here indicate the step in the algorithm.
2. For $iter \geq 0$
 - (a) Compute the **middle nest** until r_m reaches convergence, taking $p_{h,iter}$ as given
 - (b) For each s_t , iterate over values of $p_{h,iter+1}(s_t)$ until $D_h(s_t) = \lambda\mu$ (using the values for V^{stay} , Π , and r_m computed in step (a))
 - (c) Set $p_{h,iter+1}$ equal to the values computed in step (b)
 - (d) Repeat until $p_{h,iter+1} = p_{h,iter}$

3 Implementation

In the previous section, many details about the model were presented abstractly. For example, an aggregate state variable s_t was discussed, but no details were given except that it followed a first order Markov transition process. Additionally, it was mentioned that there are M mortgage product types, but no details were given about the specific product types available.

The advantage of presenting the model abstractly is that it highlights the modular nature of the model and how the model can easily be extended to incorporate new kinds of mortgage contracts or new sources of stochastic shocks. At some point, however, one must stop abstracting and begin implementing the details of the model. In this section, I begin to flesh out the details of the model that will be used for the rest of the paper.

3.1 Mortgage product types

In the baseline model, I allow for three mortgage product types. In addition to $m = 0$, “no mortgage”, the other two types are $m = 1$: agency loans, and $m = 2$: non-agency loans. Agency loans are loans that are securitizable by the government-sponsored agencies, Freddie Mac and Fannie Mae. Freddie Mac and Fannie Mae carry implicit government guarantees on their credit obligations, so it is therefore assumed that lenders treat agency loans as if there were no default risk. In practice, this means that $\psi_h(z_t, s_t) = b_t$ for agency loans. That is, lenders are compensated fully for the remaining balance on the loan in case of a default.

A mortgage must meet certain criteria in order to qualify as an agency loan. These criteria are set by the Federal Housing Finance Agency (FHFA), which regulates Freddie Mac and Fannie Mae. First, a loan must have an original balance less than or equal to what is known as a “conforming loan limit”, which I denote cll_t . In addition, the original balance of an agency loan cannot exceed 80% of the purchase price of the home. In practice, these two requirements translate to the following restriction on available contracts in the model:

$$\bar{b}_1(x_{it}, s_t) = \min \left\{ cll_t, 0.8p_h(s_t), \bar{b}_1^*(x_{it}, s_t) \right\}$$

Note here that I continue to allow for an endogenous maximum borrowing amount, $\bar{b}_1^*(x_{it}, s_t)$, that could potentially be less than the regulatory restrictions.

In contrast to agency loans, non-agency loans carry no guarantees. In the event of a default the lender forecloses on the home and recovers a fraction φ of its value. Therefore, $\psi_h(z_t, s_t) = \varphi p(s_t)$ for non-agency loans. Unlike agency loans, there are no restrictions on the size of a non-agency loan. The only restriction on non-agency loans is that the original balance cannot be greater than the purchase price of the home. Since one of the counterfactual questions I am interested in is the effect of the availability of non-agency loans, I also allow the availability of non-agency loans to depend on an aggregate state variable, $m ps_t$. If $m ps_t = 1$, then non-agency loans are available. If $m ps_t = 0$,

they are not. Therefore:

$$\bar{b}_2(x_{it}, s_t) = \begin{cases} \min \{p_h(s_t), \bar{b}_2^*(x_{it}, s_t)\} & \text{if } mps_t = 1 \\ 0 & \text{otherwise} \end{cases}$$

Both agency and non-agency loans are assumed to be 30-year fixed rate loans. The required payment and transition rules for fixed rate loans are given in the Appendix. Although in the baseline model I only allow fixed rate loans, the model is readily extendable to include other types of products, such as adjustable rate loans, in future work.

3.2 Potential buyer distributions

Recall that potential buyers are heterogeneous in their income y_i , which is constant, their initial wealth w_i , and their outside option v_i . Log-income is normally distributed with mean μ_t^y and variance σ_y^2 . The mean of the income distribution can change over time, but the variance remains constant. Wealth depends on income. For potential buyers with income y_i , wealth is generated from a censored normal distribution with mean:

$$E[w_i^*|y_i, s_t] = \alpha_0^w + \alpha_1^w y_i + \alpha_2^w p_0(s_t)$$

and variance:

$$Var[w_i^*|y_i, s_t] = \sigma_w^2$$

Censoring occurs at 0, so wealth is given by:

$$w_i = \max \{0, w_i^*\}$$

I choose a censored normal distribution for wealth rather than a log-normal distribution in order to generate a mass of potential buyers with zero wealth. It turns out that such a wealth distribution is able to explain loan-to-value ratio patterns in the data fairly well. In addition, wealth is allowed to depend on income and on the baseline level of house prices in the economy. Wealth

is allowed to be correlated with income to reflect the possibility that high-income buyers are also more likely to have higher initial wealth. Wealth is allowed to depend on the baseline level of house prices in order to model, in a reduced-form way, the potential for existing homeowners to re-enter the market as buyers. As discussed in section 2, it is unknown how this reduced-form relationship between buyer wealth and house prices would change under a counterfactual setting. That question is important, but left for future research.

Finally, let v_i^* be such that $v_i = u(v_i^*)$. The transformed variable v_i^* is uniformly distributed between 0 and \bar{v}_t . Because the lower bound of the support of v_i^* is zero, there will always be a positive mass of potential buyers who will always buy a house as long as they can afford it. When \bar{v}_t is high, the average outside option will be high and demand for houses will be lower. When \bar{v}_t is low, the average outside option will be low and the demand for houses will be higher. \bar{v}_t can therefore be thought of as an unobserved demand shock—and I will refer to it as such from now on.

3.3 Aggregate state vector and transitions

So far I have mentioned five aggregate state variables: (1) the risk free rate rfr_t ; (2) the conforming loan limit cll_t ; (3) the mean income of potential buyers μ_t^y ; (4) the unobserved demand shock \bar{v}_t ; (5) the availability of non-agency loans mps_t . In addition to these five, I allow for one more state variable, g_t . The variable g_t determines the transition process of the unobserved demand shock \bar{v}_t . In particular, conditional on \bar{v}_t , $\log(\bar{v}_{t+1})$ is normally distributed with mean:

$$E[\log(\bar{v}_{t+1}) | \bar{v}_t, s_t] = g_t + \log(\bar{v}_t) \quad (28)$$

and variance:

$$Var[\log(\bar{v}_{t+1}) | \bar{v}_t, s_t] = \sigma_v^2$$

From (28) one can see how the state variable g_t controls the evolution and expectation of future demand shocks. I allow g_t to take on two possible values, so the economy can either be in a state of growth: $g_t = g_L$ or a state of decline: $g_t = g_H$. (Note that a high rate of growth for the outside option implies

declining house prices.)

The aggregate state vector is therefore given by $s_t = (rfr_t, cll_t, \mu_t^y, mps_t, g_t, \bar{v}_t)$. Besides \bar{v}_t , I do not allow any of the other variables to vary stochastically. That is, agents in the economy believe that the other state variables remain constant over time, and if they do change, it is completely unexpected. It is not difficult to relax this assumption, but neither is it a bad assumption for the region and period of time that I will be interested in (Los Angeles 2003-2011). Average incomes in Los Angeles are roughly constant over this period. Moreover, risk-free rates and conforming loan limits do not change much over this period, and when they do (in 2008), it is in response to the unexpected financial crisis. Similarly, the disappearance of non-agency loans ($mps_t = 0$) was a result of the unexpected crisis. However, agents do believe that house prices can change over this period, which will be incorporated into the model via the unobserved demand shocks.

3.4 Default and refinance costs

For the default and refinance costs, I choose the conceptually appealing case of *ruthless default* ($c_D = 0$) and *no refinancing* ($c_R \gg 0$). Ruthless default here means that owners who are hit by a moving shock default if and only if the remaining balance of their mortgage exceeds the price of their home. Although this particular choice of default and refinance costs may not accurately reproduce all the moments in the data, it is interesting to study how far this simplified model can take us in explaining the data. As it turns out, ruthless default with no refinancing allow us to explain time-variation in default rates by buyer cohort quite well. Allowing refinancing coupled with ruthless default can help explain default rates even more closely. I discuss these points further in section 5.

3.5 Discretization

The model is solved over discrete grids of its variables. I describe exactly how each variable is discretized in the Appendix. For now, special mention should

be made of two variables for which I only allow one grid point: liquid wealth of homeowners, w_{it} , and mean log-income μ_t^y .

Only one grid point is used for mean income. What this means is that the mean income of potential buyers is constant. This is not an unrealistic assumption for Los Angeles from 2003 to 2011, because the average real income of Los Angeles residents from 2003 to 2011 was roughly constant at \$55,000 (2012 dollars). Note that the assumption does not imply that the mean income of *actual* buyers has been constant, only the mean income of *potential* buyers. In both the model and the data, the average income of actual buyers is correlated with house prices.

I only use one grid point for liquid wealth in order to reduce the computational complexity of the model. With one grid point, homeowners are assumed to save only a minimal amount for precautionary purposes. The results of the paper are robust to this assumption because only households with very high initial wealth relative to income (a small percentage) would have any incentive to save, due to the fact that both income and mortgage payments are constant in the model. Even for the households who do have an incentive to save, they are better off reducing their initial mortgage balance than increasing their level of liquid wealth. I have also computed a version of the model with two grid points for liquid wealth, and found that doing so did not change any results (the large majority of homeowners do not choose to save at a higher rate).

Finally, three other variables with a small number of grid points will be mentioned. First, household income y_i takes one of two values: \$80,000 or \$150,000. These roughly correspond to average buyer income for “low-quality” and “high-quality” homes during the sample period. Second, the risk-free rate rfr_t takes only the values of 0.025 and 0.015. These roughly correspond to the real 10-year treasury rate before 2007 and after 2008. Third, the conforming loan limit cll_t takes only the values \$400,000, \$450,000, and \$750,000. These roughly correspond to the values that the actual conforming loan limit took from 2003 thru 2011. Most other variables, such as initial wealth w_i , have a large number of grid points.

4 Data

Data for the calibration of the model comes from three main sources. The first is an administrative database of housing transactions provided by DataQuick, a real estate consulting company. DataQuick collects data from public records on property transactions. Each time a property is sold, the transaction price, closing date, and any liens against the property are recorded. This allows me to observe the universe of home sale prices and the loan amounts they are purchased with. In addition to sales, DataQuick also contains information about refinances. Each time a new loan is originated against a property, the amount of that loan and the origination date are recorded by DataQuick. DataQuick data for Los Angeles goes back to 1988, and the latest year for which I have data is 2012. Each observation in DataQuick includes a unique property identifier, so it is possible to follow a single property over time and construct an ownership history that includes purchase date, purchase price, purchase loans, sale price, and any refinances along the way. The unique property identifier also allows me to construct price indices using a repeat sales methodology.

Using the DataQuick data, I first decompose the L.A. housing market into two segments, which correspond to $h = 0$ and $h = 1$ in my model. To do this, I first run a fixed-effects regression of log sale price on property fixed effects and year fixed effects.⁷ I then divide the properties into two groups, based on whether the property's estimated fixed effect is above or below the median fixed effect in the data. If the property has an estimated fixed effect higher than the median, it is assigned $h = 1$. It is assigned $h = 0$ otherwise.

After assigning the properties to two groups, I perform a repeat sales regression using the methodology of Case and Shiller (1989) separately on the two groups. Figure 2 reports the resulting price indices with 1999 as the base year, and Figure 3 reports the implied mean house price levels, again using 1999 as the base year. Figure 2 shows that the lower valued homes in Los

⁷All prices are used are real house prices, reported in 2012 dollars. Nominal prices are deflated by the CPI-U-national. All interest rates reported in this section are also deflated by 0.018, which is the average inflation rate from 2003 to 2010.

Angeles appreciated at a considerably higher rate than did the higher value homes. This phenomenon was also observed in San Diego by Landvoigt et al. (2014). Figure 3 reveals that although lower-valued homes appreciated at a faster rate, this is mostly because they started at a lower base level. In fact, it appears that for most of the last decade, yearly changes to the level prices for the two groups followed each other in lock-step. The phenomenon revealed in Figure 3 can be replicated by model in this paper, and is best explained by a model where the price differences between the low and high-valued groups of housing are primarily determined by wealth and borrowing constraint.

In addition to constructing price paths, the DataQuick data is used to construct the average loan-to-value ratios of Los Angeles home buyers from 2003 to 2011. This is done separately for the two groups of housing. Table 2 shows the average LTV of buyers by year and house type. Two facts stand out. First, buyers of low-valued housing generally have to borrow more to purchase their homes. This is likely due to them having lower initial wealth. Second, LTVs for both groups dropped considerably after 2008. This is likely due to the disappearance of the non-agency mortgage market. Both facts can be explained by this paper's model.

Besides DataQuick, the two other sources of data are two loan-level databases; the first provided by Freddie Mac and the second provided by BlackBox, a real estate finance consulting company. The Freddie Mac data is a random sample of 30-year fixed rate loans originated between 1999 and 2012, and purchased by Freddie Mac. The database contains important information about each loan contract, including the original interest rate. The BlackBox data is a loan-level administrative dataset containing information on all kinds of non-agency mortgage products. The data covers over 90% of all non-agency securitization pools. The BlackBox data contains a wealth of information about each mortgage's contractual terms, but for the purposes of this paper I am only interested in the product type and the contract rate. Like the Freddie Mac dataset, the BlackBox data goes back to 1999.

The first thing I would like to do with the mortgage data is to verify that the non-agency market all but disappeared in 2008. Using BlackBox data,

Table 3 shows the number of non-agency fixed-rate loan originations in Los Angeles from 2003 to 2010 (the results do not change when including other product types). It is clear from this table that non-agency mortgages became unavailable after 2008. This fact is also confirmed by external sources. Figure 1 shows the total volume of agency and non-agency mortgage-backed securities issuance in the U.S., as documented by the Securities Industry and Financial Markets Association. The figure shows that non-agency securitization had completely died out by 2008.

The second thing I do with the mortgage data is construct average interest rates for agency and non-agency mortgage originations in Los Angeles, from 2003 to 2010. The averages are computed using only 30-year fixed rate loans. The average real interest rates are shown in Table 4. In the calibration procedure, the average interest rates on the two mortgage product types will be used as moments to which the model will be matched.

5 Calibration

5.1 Data

The model parameters are calibrated by fitting model predictions to aggregate moments observed in the data. The data available are the following:

$$\{rfr_t, cll_t, mps_t, p_{ht}, r_{mt}, \overline{ltv}_{ht}, T_{it}\}$$

for $t = 1, \dots, 8$. That is, I observe the risk-free rate, conforming loan limit, availability of non-agency loans, house prices for each house type, average mortgage interest rate for each product type among buyers, and average LTV of buyers of each house type. In addition, I also see the duration between purchase and sale, T_{it} , for buyers i in each period t . The 8 periods correspond to the housing market in Los Angeles from 2003 to 2010. The values for some of these data are given in Figure 3 and Tables 2 thru 4. The path of the aggregate state variables: rfr_t , cll_t and mps_t are together given in Table 5. Note that in Table 5, g_H and g_L are parameters to be determined, and

the demand shocks \bar{v}_t are not observed. However, the demand shocks can be backed out from realized price paths by inverting $p_h(s_t)$ in each period.

5.2 Calibration method

The parameters to be calibrated are given in Table 1. The first two parameters of the model, the risk aversion parameter γ and the time-discount factor δ are chosen to be 3 and 0.95. The total mass of each house type, μ , is normalized to 1. The variance of the log of potential buyer income is 0.9025, which is equal to the average cross-sectional variation in log-income for Los Angeles residents from 2003 to 2011, as computed using data from the American Community Survey.

The first unknown parameter, λ , can be estimated from T_{it} simply by calculating the per-period hazard rate of a homeowner selling his or her home. Doing this yields an estimate of $\lambda = 0.0952$, which corresponds to an average duration between purchase and sale of about 10 years.

The rest of the parameters are calibrated to fit aggregate moments in the data. Given a guess of the parameters, the model is first solved as described in Section 2. Once the function $p_h(s_t)$ is known, \bar{v}_t can be computed for each period by inverting $p_h(s_t)$ and taking the other state variables as given. Once the state variables in each period are known, the model is simulated for the 8 periods from 2003 to 2010.

The moments to which the model is compared are the following:

1. p_{ht} : price paths in Los Angeles, for $h = 0$ and $h = 1$, as given in Figure 3
2. r_{mt} : path of average interest rate rate by product type among new buyers, as given in Table 4
3. \overline{ltv}_{ht} : path of average LTVs among new buyers, for $h = 0$ and $h = 1$, as given in Table 2
4. Implied transition distribution of the unobserved demand shocks (I elaborate more below)

The simulated house price are computed directly from $p_h(s_t)$. The average interest rate by product type and the average LTV by house type are calculated by averaging over the choices of the buyers in each year.

5.3 Sources of identification

Although all the parameters will affect all the moments in different ways, it is useful to discuss the ways in which particular parameters might affect particular moments, in order to better understand the sources of identification.

First, the price path of a single house type, say $h = 0$, can be used to identify \bar{v}_t in each period, and then θ is identified off the price differences between $h = 0$ and $h = 1$ houses.

Second, the rate premium of agency loans above the risk-free rate can be used to identify the lenders' premium on returns, a . Then, the difference between the interest rates on agency and non-agency mortgages can be used to identify the foreclosure recovery rate ψ .

Third, the path of average LTVs for the two types of housing can be used to identify the parameters governing both the wealth distribution, α^w . For example, larger differences in LTV between $h = 0$ and $h = 1$ buyers imply a higher coefficient of income on wealth. Changes to the average LTV over time can help identify the coefficient of current house prices on wealth. The LTV path can also help identify the weight on the utility over final wealth β . Intuitively, β controls how attractive houses are as investment assets, so when β is higher, there will be more leveraging when house prices are expected to appreciate.

Finally, the parameters governing the evolution of the unobserved demand shock, g_H , g_L and σ_v^2 can also be estimated. For any guess of the parameters g_H , g_L , σ_v^2 , first compute the implied realizations of \bar{v}_t each period by matching the realized price paths. Then, an estimated \hat{g}_H , \hat{g}_L and $\hat{\sigma}_v^2$ can be computed from:

$$\hat{g}_j = \frac{1}{8} \sum_{t=1}^8 1[g_t = g_j] \times \log(\bar{v}_{t+1}/\bar{v}_t) \text{ for } j = H, L$$

$$\hat{\sigma}_v^2 = \frac{1}{7} \sum_{t=1}^8 [\log(\bar{v}_{t+1}/\bar{v}_t) - g_t]^2$$

If g_H , g_L and σ_v^2 are the true parameters that generate the data, then $E[\hat{g}_j] = g_j$ and $E[\hat{\sigma}_v^2] = \sigma_v^2$. Therefore, $\hat{g}_j - g_j$ and $\hat{\sigma}_v^2 - \sigma_v^2$ are moments that can be used to estimate g_H , g_L and σ_v^2 .

5.4 Results and model fit

Table 6 shows the resulting values for each of the unknown model parameters. Figure 4 and Tables 7 and 8 show the model fit for house price paths, buyer LTVs, and mortgage interest rates under these parameters.

Under these parameters, the model does a very good job of fitting the house price paths, as shown in Figure 4. Even though there is a free parameter each period, \bar{v}_t , to match prices, it is not a given that the model would match both price paths well because only one parameter, θ , controls the difference between the prices of $h = 0$ and $h = 1$ houses. The fact that the model matches both price paths well suggests that the price difference between low and high valued housing can be well explained by a single parameter which is constant over time.

Table 7 shows the average buyer LTVs for each house type under both the real and the simulated data. The model fits the LTVs for low-valued housing very well, and also does a good job fitting the LTVs for high-valued housing after 2008. Pre-2008, the model tends to overpredict the LTVs of high-valued house-buyers, but the trend over time is increasing from 2003 to 2006, as it is in the data.

Table 8 shows the average mortgage interest rates by product type under real and simulated data. The model does a good job fitting the mortgage interest rates on agency loans but is unable to explain some of the time-variation seen in the data. This is primarily due to discretization of the risk-free rate into two possible grid points. In contrast, the model is unable to fit well the mortgage interest rates on non-agency loans. Because default rates are generally quite low, the model cannot generate a large difference

between the interest rates on agency and non-agency loans. In fact, to even generate the differences seen in Table 8, the calibrated value for φ had to be pushed down to zero. In the data, the interest rate spread between agency and non-agency loans is being driven by something outside my model—perhaps heterogeneity in the underlying default risk of the borrowers, or differences in liquidity between the agency and non-agency markets—both of which are not modeled. I will leave the resolution of this matter to future work.

Finally, Figures 5 thru 7 show the cumulative default rates over time for the cohort of buyers in 2004, 2005, and 2006. Both the cumulative default rates in the data, and the model’s simulated default rates are plotted. The model appears to underpredict the level of defaults, but predicts changes over time well. In one sense, this is reassuring, because the calibration routine does not use default rates as a moment to be matched. The fact that changes to the cumulative default rate over time are predicted well indicates that default driven by negative equity and moving shocks can explain the time-variation in default rates. However, the model appears to be missing a baseline level of default that is not correlated with house prices. This baseline level of default may be coming from borrowers who experience negative income shocks but are also unable to sell their homes before their lenders initiate foreclosure.

6 Counterfactuals

6.1 The impact of non-agency mortgage credit availability

Figure 3 shows that house prices experienced their largest decline from 2007 to 2008. In between these two periods, three of the observable state variables changed: the risk-free rate went down, the conforming loan limit went up, and the non-agency mortgage market disappeared. The first two changes should have the effect of increasing house prices—so the fact that house prices went down implies that the effect of the disappearing non-agency market was dominant. The non-agency market is very important to house prices because it

is what allows potential buyers with low wealth to borrow loans at high LTV ratios, letting them afford the down payment. Without high LTV ratio loans, these buyers—even if they have a large desire to purchase housing—are priced out of the housing market.

Figures 8-11 illustrate. Figure 8 shows the housing demand profile for potential buyers in 2004, as simulated under the baseline model. Notice a positive mass of buyers at all wealth levels are purchasing housing. Figure 9 shows the mortgage choice profile for potential buyers in 2004. All of the buyers who are purchasing housing with low wealth are using non-agency loans. Contrast these figures to Figure 10, which shows the housing demand profile in 2008. In 2008, only buyers with high wealth are purchasing homes. As Figure 11 shows, this is because low-wealth buyers cannot afford the down-payment required for agency loans. Even if they have a very large desire to purchase housing (i.e. a very low outside option), these buyers are priced out of the market. Prices must be lowered, so that a larger share of high-wealth buyers purchase in 2008, relative to 2004.

A natural question to ask would be: What if the non-agency did not disappear in 2008? The model can be used to simulate outcomes under this counterfactual, simply by changing mps_t to 1 for 2008 thru 2010. The counterfactual outcomes are simulated under the values for \bar{v}_t estimated from the baseline model. Figure 12 shows the resulting price path under the counterfactual (dashed line) compared to the baseline (solid line). The results suggest that house prices would have been much higher in 2008, 2009, and 2010 had non-agency mortgage financing still been available.

Figures 13 and 14 illustrate why. As discussed above, the presence of the non-agency market allows lower-wealth buyers to enter the housing market. Notice, however, that the buyers with the lowest wealth are still being priced out. In fact, low-income buyers are not using non-agency mortgages at all. This is for two reasons. First, prices are dramatically higher in the counterfactual, leading some buyers to be priced out based on their inability meet payment-to-income ratio requirements. Second, non-agency loans continue to be *endogenously* unavailable at certain LTV levels, especially for low-valued

homes. Figure 15 illustrates this by plotting the interest rate on new mortgage originations by product type in 2008. For non-agency loans, interest rate increases rapidly with LTV because when prices are expected to decline, higher LTV dramatically increases the chance of future default. The chance of default and the required interest rate compensation may rise to the point where it becomes no longer profitable for the lender to originate the mortgage at any reasonable interest rate.⁸

6.2 The effectiveness of government responses

The government responded to the disappearance of the non-agency market in two ways. First, it increased the conforming loan limit for various cities in 2008. For Los Angeles, the nominal conforming loan limit was increased from \$417,000 to \$729,750. Second, it reduced the risk-free rate. In this section, I use the model to investigate the effectiveness of these responses.

Figure 16 shows the house price paths that would have resulted if the government had not increased conforming loan limits in 2008. Increasing conforming loan limits had a large effect on high-priced housing in 2008, but the effect becomes smaller in 2009 and 2010. This is expected, because as price levels go down, the importance of a high conforming loan limit is also reduced. The differential effect between house types of increasing the conforming loan limit in 2008 is about 5.3%. This is comparable to the reduced-form difference-in-difference estimate of 6.7% reported in Kung (2014). The relatively small effect of a large increase in conforming loan limits suggests that the larger problem for house prices in 2008 was the difficulty of obtaining high LTV loans rather than large loans in terms of face amount.

Figure 17 shows the house price path that would have resulted if the government had not reduced the risk-free rate in 2008. The results show that reducing the interest rate had only a small effect in 2008, but larger effects in 2009 and 2010. This is reasonable because in 2008, when prices are still fairly high, the determination of house price is driven to a larger extent by

⁸I set the interest rate cap to 10% in the counterfactual. If a lender cannot profitably originate a mortgage at 10% interest rate, that mortgage becomes endogenously unavailable.

the amount of down-payment buyers are able to afford, rather than by considerations of reduced consumption due to higher interest rates. When prices are lower in 2009 and 2010, considerations about consumption become more important relative to liquidity.

Figure 18 shows the house price path that would have resulted if the government had done nothing. The resulting price path appears simply to be the minimum of the two counterfactual price paths when considering either policy intervention separately. Partially, this is due to the coarseness of the discretization (prices are only allowed to adjust in \$25,000 increments). But it is also partly due to the fact that high conforming loan limits simply did not matter much in 2009 and 2010, when prices are already low; and lower interest may not have mattered much in 2008, when prices are high and demand is largely determined by who can afford the required down payments.

Table 9 shows the welfare effect of the government response on buyers in 2008, 2009 and 2010. Average welfare is measured in units of consumption equivalence, which is the amount of promised yearly consumption (to be received forever) that would make a household just as well off as under the model. The table shows that the government response had only a small effect on buyers—but because it raised prices, it benefitted existing owners. This is not surprising, and is a general point, as demand in the model is elastic but supply is inelastic. Future versions of the paper will report the welfare effect on existing owners as well as on buyers.

7 Conclusion

The counterfactual exercises illustrate the utility that this model has in answering a diverse set of questions about how mortgage market institutions affect housing market outcomes (an important, and extremely interesting, set of questions). The model is designed to be very modular, and is therefore easily extended and adapted to a variety of settings suitable for the study of specific questions.

References

- Adelino, Manuel, Antoinette Schoar, and Felipe Severino**, “Credit Supply and House Prices: Evidence from Mortgage Market Segmentation,” Working Paper, Duke University 2014.
- Campbell, John Y. and Joao F. Cocco**, “Household risk management and optimal mortgage choice,” *Quarterly Journal of Economics*, 2003, 118 (4), 1449–1494.
- and – , “A Model of Mortgage Default,” Working Paper, Harvard University 2014.
- Case, Karl E. and Robert J. Shiller**, “The Efficiency of the Market for Single Family Homes,” *The American Economic Review*, Mar 1989, 79 (1), 125–137.
- Corbae, Dean and Erwan Quintin**, “Leverage and the Foreclosure Crisis,” Working Paper, University of Wisconsin - Madison 2014.
- Favara, Giovanni and Jean M. Imbs**, “Credit Supply and the Price of Housing,” CEPR Discussion Paper No. DP8129 2010.
- Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh**, “The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium,” *NBER Working Paper No. 15988*, 2010.
- Glaeser, Edward L., Joshua D. Gottlieb, and Joseph Gyourko**, “Can Cheap Credit Explain the Housing Boom?,” *NBER Working Paper No. 16230*, 2010.
- Himmelberg, Charles, Christopher Mayer, and Todd Sinai**, “Assessing High House Prices: Bubbles, Fundamentals and Misperceptions,” *Journal of Economic Perspectives*, 2005, 19 (4), 67–92.

- Kung, Edward**, “The Effect of Credit Availability on House Prices: An Instrumental Variable and Difference-in-Differences Approach,” Working Paper, UCLA 2014.
- Landvoigt, Tim, Monika Piazzesi, and Martin Schneider**, “The Housing Market(s) of San Diego,” Working Paper, Stanford University 2014.
- Mian, Atif and Amir Sufi**, “The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis,” *Quarterly Journal of Economics*, 2009, *124* (4), 1449–1496.
- **and —**, *House of debt: how they (and you) caused the Great Recession, and how we can prevent it from happening again*, 1st ed., Chicago, IL: University of Chicago Press, 2014.
- Ortalo-Magné, François and Sven Rady**, “Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints,” *Review of Economic Studies*, 2006, *73* (2), 459–485.
- Piskorski, Tomasz and Alexei Tchistyi**, “Optimal Mortgage Design,” *Review of Financial Studies*, 2010, *23* (8), 3098–3140.
- Shiller, Robert J.**, *The Subprime Solution: How Today’s Global Financial Crisis Happened and What to Do about It*, 1st ed., Princeton, NJ: Princeton University Press, 2008.
- Stokey, Nancy L., Robert E. Lucas, and Edward C. Prescott**, *Recursive Methods in Economic Dynamics*, 1st ed., Cambridge, MA: Harvard University Press, 1989.

Table 1: Parameters in the Model

Parameter	Description	Value
γ	coefficient of relative risk aversion	3
δ	time discounting rate	0.95
μ	total mass of each house type	1
σ_y^2	variance of potential buyer income	0.9025
λ	per-period probability of moving	Calibrated
β	weight on utility over final wealth	Calibrated
θ	preference for high quality homes	Calibrated
a	return lender can get over risk-free rate	Calibrated
φ	recovery rate on foreclosures	Calibrated
$(\alpha_0^w, \alpha_1^w, \alpha_2^w, \sigma_w^2)$	parameters governing wealth distribution	Calibrated
g_L, g_H, σ_v^2	evolution of unobserved demand shock	Calibrated

Note: This table lists the free parameters of the model and what their values are set to. If the parameter is to be estimated from data, the estimates are given in Table 6.

Table 2: Average LTVs of Los Angeles Home Buyers

Year	Low-Valued Homes	High-Valued Homes
2003	0.844	0.756
2004	0.849	0.760
2005	0.857	0.760
2006	0.884	0.779
2007	0.842	0.723
2008	0.755	0.617
2009	0.725	0.608
2010	0.723	0.598

Note: This Table shows the average LTV of buyers by purchase year in Los Angeles. The LTVs are computed using DataQuick data.

Table 3: Non-Agency Loan Originations in Los Angeles (BlackBox Data)

Year	# Originations
2003	9,673
2004	11,349
2005	16,039
2006	17,240
2007	7,352
2008	13
2009	0
2010	4

Note: This table shows the number of non-agency mortgage originations in Los Angeles, as reported in the BlackBox data. The table is meant to illustrate that the non-agency market disappears after 2008.

Table 4: Average Interest Rates on 30-Year FRM Originations in Los Angeles (Less Inflation)

Year	Agency Loans	Non-Agency Loans
2003	0.040	0.048
2004	0.041	0.055
2005	0.040	0.062
2006	0.045	0.069
2007	0.045	0.052
2008	0.043	N/A
2009	0.033	N/A
2010	0.031	N/A

Note: This table shows the average interest rate on agency and non-agency 30-year fixed-rate loans in Los Angeles. The interest rates for agency loans are computed from the Freddie Mac loan-level dataset and the interest rate for non-agency loans are computed from the BlackBox data.

Table 5: Aggregate State Variable Paths used in Calibration

Year	rfr_t	cll_t	mps_t	g_t	\bar{v}_t
2003	0.025	0.4	1	g_H	Unobserved
2004	0.025	0.4	1	g_H	\vdots
2005	0.025	0.4	1	g_H	\vdots
2006	0.025	0.45	1	g_H	\vdots
2007	0.025	0.45	1	g_L	\vdots
2008	0.015	0.75	0	g_L	\vdots
2009	0.015	0.75	0	g_L	\vdots
2010	0.015	0.75	0	g_L	\vdots

Note: This table shows the aggregate state variable paths used in model estimation. g_H and g_L , and \bar{v}_t for each period are all parameters to be estimated.

Table 6: Calibration Results

Parameter	Description	Estimate
β	weight on utility over final wealth	0.125
θ	preference for high quality homes	1.3
a	return lender can get over risk-free rate	0.02
φ	recovery rate on foreclosures	0
α_0^w	wealth distribution (constant)	-0.095
α_1^w	wealth dist. (coeff. on income)	1.05
α_2^w	wealth dist. (coeff on house price)	0.075
σ_w^2	wealth dist. (variance)	0.195
g_L	demand shock growth (decline)	0.03
g_H	demand shock growth (growth)	-0.08
σ_v^2	demand shock growth (variance)	0.09

Note: This table shows estimation results from the simulated method of moments. Standard errors will be available in a later version of the paper.

Table 7: Model Fit: LTVs of Home Buyers

Year	Real Data		Simulated Data	
	Low-Valued	High-Valued	Low-Valued	High-Valued
2003	0.844	0.756	0.865	0.811
2004	0.849	0.760	0.859	0.830
2005	0.857	0.760	0.863	0.846
2006	0.884	0.779	0.866	0.842
2007	0.842	0.723	0.854	0.783
2008	0.755	0.617	0.719	0.645
2009	0.725	0.608	0.692	0.617
2010	0.723	0.598	0.692	0.617

Note: This table compares the average LTVs by house type in the actual data and in the simulated data under the estimated parameters.

Table 8: Model Fit: Mortgage Interest Rates

Year	Real Data		Simulated Data	
	Agency	Non-Agency	Agency	Non-Agency
2003	0.040	0.05	0.045	0.045
2004	0.040	0.055	0.045	0.047
2005	0.040	0.06	0.045	0.048
2006	0.045	0.07	0.045	0.047
2007	0.045	0.05	0.045	0.051
2008	0.045		0.035	
2009	0.035		0.035	
2010	0.030		0.035	

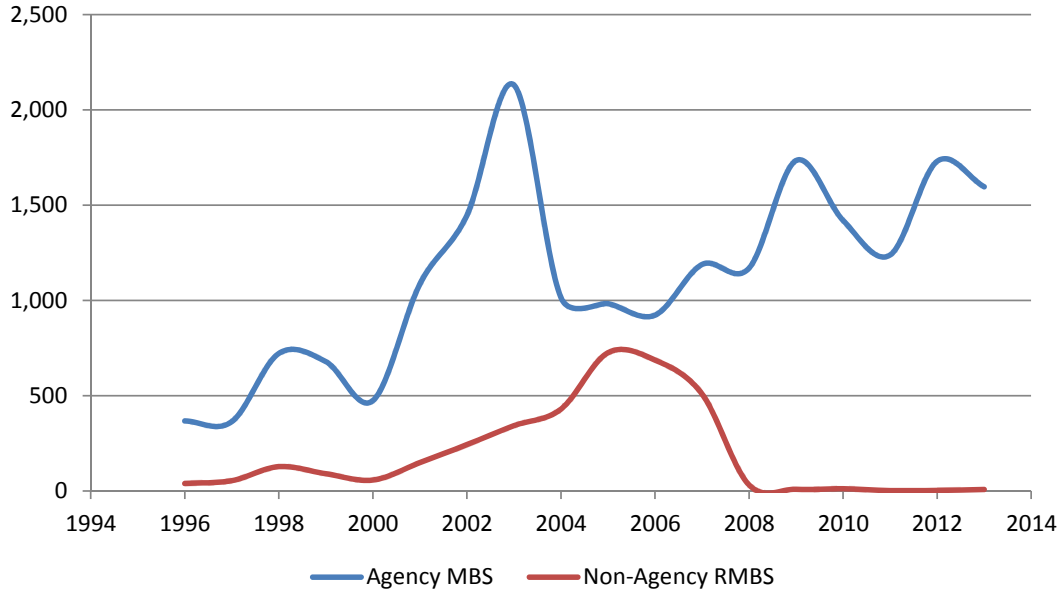
Note: This table compares the average interest rates by mortgage product type in the actual data and in the simulated data under estimated parameters. Reasons for the poor fit of non-agency interest rates are discussed in Section 5.4.

Table 9: Counterfactual 4: No Government Response—Welfare Effects for Buyers

Year	Consumption Equivalent		Difference	
	Baseline	Counterfactual	Levels	Percent
2008	\$162,000	\$159,700	+\$2,300	+1.44%
2009	\$180,500	\$181,400	-\$900	-0.5%
2010	\$180,500	\$181,400	-\$900	-0.5%

Note: This table shows the welfare effects for the counterfactual exercise in which interest rates are not reduced and conforming loan limits are not increased in 2008. Consumption equivalent is defined as the amount of consumption that, if received deterministically every period, would result in the same expected utility as under the model.

Figure 1: Agency and Non-Agency MBS Issuance (USD Billions)




Notes: This graph shows the total volume of agency and non-agency mortgage backed securities issuance from 1994 to 2013. Source: SIFMA. Data can be found at <https://www.sifma.org/uploadedfiles/research/statistics/statisticsfiles/sf-us-mortgage-related-sifma.xls>. Date last accessed: 8/29/2014.

Figure 2: House Price Indices in Los Angeles




Note: This graph shows the repeat-sales price index computed using DataQuick data for the two housing segments as described in Section 4.

Figure 3: House Price Levels in Los Angeles



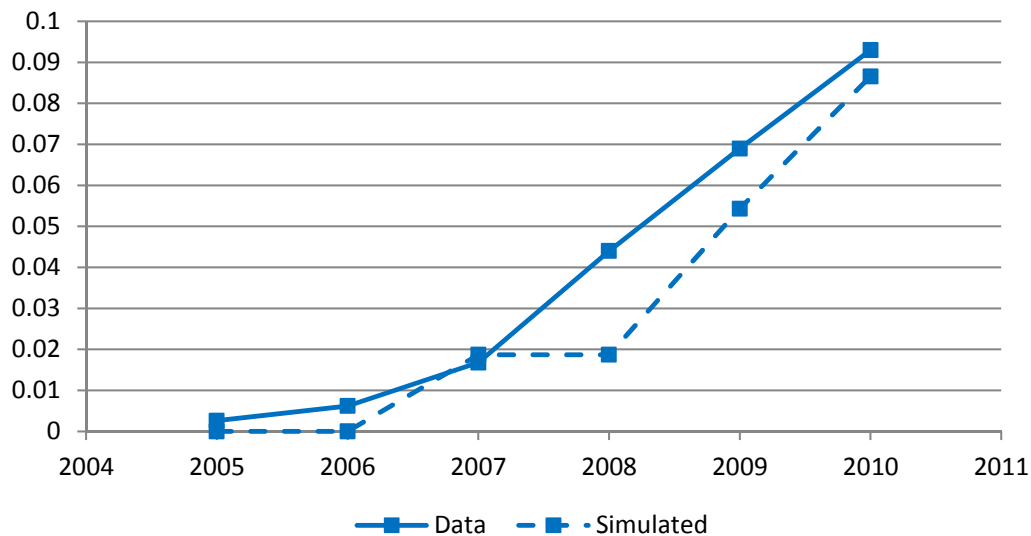
Note: This graph shows the house price levels implied by the repeat-sales house price indices given in Figure 2. 1999 is the base year.

Figure 4: Model Fit: House Price Paths




Note: This graph compared the price paths in the data to the price paths simulated from the baseline model with estimated parameters.

Figure 5: Model Fit: Cumulative Default Rates (2004 Cohort)




Note: This figure shows the cumulative default rate over time for the cohort of buyers who purchased in 2004.

Figure 6: Model Fit: Cumulative Default Rates (2005 Cohort)




Note: This figure shows the cumulative default rate over time for the cohort of buyers who purchased in 2005.

Figure 7: Model Fit: Cumulative Default Rates (2006 Cohort)




Note: This figure shows the cumulative default rate over time for the cohort of buyers who purchased in 2006.

Figure 8: Housing Demand Profile in 2004 (Baseline)




Note: This figure shows purchasing decision by buyer type using the baseline model parameters simulated in 2004.

Figure 9: Mortgage Demand Profile in 2004 (Baseline)




Note: This figure shows mortgage decision by buyer type using the baseline model parameters simulated in 2004.

Figure 10: Housing Demand Profile in 2008 (Baseline)




Note: This figure shows purchasing decision by buyer type using the baseline model parameters simulated in 2008.

Figure 11: Mortgage Demand Profile in 2008 (Baseline)




Note: This figure shows mortgage decision by buyer type using the baseline model parameters simulated in 2008.

Figure 12: Counterfactual 1: Non-Agency Availability—Price Paths




Note: This graph reports counterfactual and baseline price paths for the counterfactual exercise in which non-agency mortgage financing is made available from 2008 onwards. The solid lines show the baseline price paths for $h = 0$ and $h = 1$ housing respectively, and the dashed lines show the counterfactual price paths.

Figure 13: Counterfactual 1: Non-Agency Availability—Housing Demand in 2008




Note: This figure shows purchasing decision by buyer type in 2008 for the counterfactual exercise in which non-agency mortgage financing is made available.

Figure 14: Counterfactual 1: Non-Agency Availability—Mortgage Demand in 2008




Note: This figure shows mortgage decisions by buyer type in 2008 for the counterfactual exercise in which non-agency mortgage financing is made available.

Figure 15: Counterfactual 1: Non-Agency Availability—Mortgage Rates in 2008



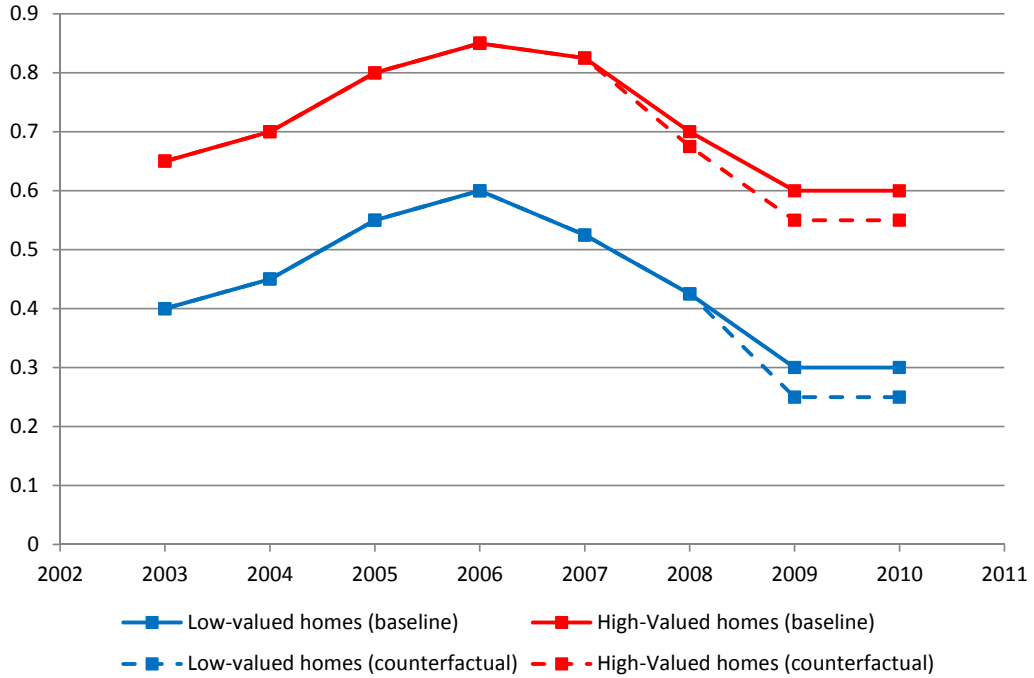
Note: This figure shows mortgages rates by LTV for low-income buyers in 2008, for the counterfactual exercise in which non-agency mortgage financing is made available. Lack of data points at certain LTV levels indicates non-availability of the product type at those LTV levels.

Figure 16: Counterfactual 2: No CLL Response—Price Paths




Note: This graph reports counterfactual and baseline price paths for the counterfactual exercise in which conforming-loan limits are not increased in 2008. The solid lines show the baseline price paths for $h = 0$ and $h = 1$ housing respectively, and the dashed lines show the counterfactual price paths.

Figure 17: Counterfactual 3: No Interest Rate Response—Price Paths



Note: This graph reports counterfactual and baseline price paths for the counterfactual exercise in which interest rates are not reduced in 2008. The solid lines show the baseline price paths for $h = 0$ and $h = 1$ housing respectively, and the dashed lines show the counterfactual price paths.

Figure 18: Counterfactual 4: No Government Response—Price Paths



Note: This graph reports counterfactual and baseline price paths for the counterfactual exercise in which interest rates are not reduced and conforming loan limits are not increased in 2008. The solid lines show the baseline price paths for $h = 0$ and $h = 1$ housing respectively, and the dashed lines show the counterfactual price paths.

A Online Appendix

A.1 Examples of mortgage contracts

30-Year Fixed Rate Mortgages

A 30-year fixed rate mortgage with $a_t < 29$ is characterized by the following transition rules:

$$\begin{aligned}a_{t+1} &= a_t + 1 \\r_{t+1} &= r_t \\b_{t+1} &= (1 + r) b_t - \text{pay}_h(z_t, s_t) \\m_{t+1} &= m_t\end{aligned}$$

And a 30-year fixed mortgage with $a_t = 29$ is characterized by:

$$\begin{aligned}a_{t+1} &= 0 \\r_{t+1} &= 0 \\b_{t+1} &= 0 \\m_{t+1} &= 0\end{aligned}$$

So we can write:

$$\zeta_h(z_t, s_t, s_{t+1}) = \begin{cases} (a_t + 1, r_t, (1 + r) b_t - \text{pay}_h(z_t, s_t), m_t) & \text{if } a_t < 29 \\ (0, 0, 0, 0) & \text{if } a_t = 29 \end{cases}$$

A feature of fixed rate mortgages is that payment is constant in each period, such that the mortgage is fully amortized over 30 years. This requires that the payment is given by:

$$\begin{aligned}\text{pay}_h(z_t, s_t) &= \frac{r_t (1 + r_t)^{30}}{(1 + r_t)^{30} - 1} b_{t-a_t} \\b_{t-a_t} &= \frac{(1 + r_t)^{30} - 1}{(1 + r_t)^{30} - (1 + r_t)^{30-a_t}} b_t\end{aligned}$$

where b_{t-a_t} is the initial balance.

A.2 Evidence on within-market movers

In this section, I use data from the American Community Survey (ACS) to explore the extent to which my results may be affected by the assumption that homeowners who move will move to a different housing market. Within-city movers are an important subset of total movers. According to the ACS in 2005, about 50% of recent movers (households who moved within 1 year of the survey) moved from within the same metropolitan area. Therefore, I cannot justify my assumption simply by claiming that there are few households who move within a housing market. However, if within-market movers tend to move between houses of similar value, then equilibrium pricing will not be significantly affected because there is no net demand or supply being created within a segment of the housing market (broadly defined).

Using ACS data, I first show that, conditional on owning a home, life-cycle changes to housing value are small compared to initial differences in housing value due to income and education. Figure ??? plots average log house values across the U.S. in 2005 as a function of age, for college educated and non-college educated homeowners. The figure shows that although there are significant increases to housing value from age 25 to age 40, these differences are small compared to initial differences in housing value due to education. The average house value for a 25 year old college-educated homeowner is greater than the average house value of a 40 year old non-college educated homeowner. This evidence supports the mechanism in my model where the *initial* housing decision is more determinative of housing value than changes to housing value over time. Figure ??? shows a similar plot for households above and below the median income level at each age group. Figure ??? tells much the same story as Figure ???—that initial differences in education and income play a larger role in determining house value than changes over the life-cycle.

Although differences in education and income play a large role in determining initial house value, there are still significant changes to housing value

over the life cycle. The question, however, is rather these changes over the life cycle are induced by *moving* between owned homes, or whether they are induced by buying a first home at different points in the life cycle (and therefore at different levels of wealth). Unfortunately, ACS data do not allow me to distinguish between movers who are moving from a previously owned home and movers who are buying a home after having rented. The data therefore precludes a direct test of the change in housing value at the time of a move from one owned home to another. However, the ACS data do allow me to investigate whether the housing value of recent movers differs systematically from the housing value of owners who did not move. If homeowners tend to upgrade their homes significantly at the time of a move, then *ceteris paribus*, recent movers should have higher average housing values than homeowners who did not move. To investigate this, I run the following regression:

$$y_i = \beta_0 + \beta_1 InMove_i + \beta_2 OutMove_i + X_i\beta_3 + \epsilon_i \quad (29)$$

where y_i is the reported log housing value of the owner, $InMove_i$ is an indicator for whether the owner moved within the past year (from within the metro area), $OutMove_i$ is an indicator for whether the owner moved within the past year (from outside the metro area), and X_i is a set of controls including a quadratic for household income and dummies for the age of the homeowner, the year of the survey, the metropolitan area, and the race, education level, and employment status of the homeowner. The results of this regression are reported in Table 10, column 1. The results indicate that homeowners who recently moved from outside the metropolitan area tend to purchase homes of slightly higher value (1 percent) than homeowners in the metropolitan area who did not recently move. In contrast, homeowners who recently moved from within the metropolitan area do not appear to purchase homes that are different from the average homeowner in the metropolitan area. If within-market movers have an overall tendency to upgrade or downgrade when they move, there should have been a significant coefficient on $InMove_i$.

One possibility for the lack of a significant coefficient is that some home-

owners upgrade and some downgrade, and that on net these two effects cancel out. To investigate this possibility, I re-run regression (29) separately from homeowners under age 45 and for homeowners over age 45. Homeowners under the age of 45 are more likely to upgrade when they move and homeowners over the age of 45 are more likely to downgrade when they move. This is confirmed in columns 2 and 3 of Table 10, which shows that recent movers under the age of 45 tend to have higher housing value than observationally similar owners who did not recently move, and vice versa for movers over the age of 45. The results imply that there is a tendency to upgrade for young homeowners and a tendency to downgrade for old homeowners. However, the magnitude of the upgrades and downgrades appear to be fairly small. On average, the housing value of a recent young mover is only 4.5 percent higher than the housing value of a similar non-mover. On average, the housing value of a recent old mover is only 5 percent lower than the housing value of a similar non-mover. These results imply that the large differences in housing value between cohorts aged between 25 and 40 are driven mostly by first-time buyers who are buying at different age and wealth levels.

Overall, it appears that the assumption that homeowners who move move to a different housing market is a reasonable first approximation. Although the data suggests that there are within-market movers, these movers do not appear to make large changes to their housing value each time they move. Therefore, in a model with two housing segments that have large value differences, no supply is created in either segment of the housing market by within-market movers. The data suggests that differences in housing value are driven predominantly by differences initial wealth, income and education at the time of first purchase. These are mechanisms are captured by the model.


A.3 Discretization of the model

Table 11 below lists the variables in the model that are discretized and the corresponding grids for each variable. Of special note is the way in which loan amounts are discretized. First, a distinction is made between “original loan

amount” and “initial amount to borrow”. Because all mortgages in the baseline model are fixed-rate mortgages, the exact current balance can be computed from the original loan amount. The original loan amount is therefore the state variable which is kept track of, rather than current loan amount. Keeping track of the original loan amount allows for a more robust calculation of the exact remaining balance in each period and state-of-the-world.


In order to reduce the size of the decision space, borrowers do not freely choose their initial balance from the entire grid of original loan amounts. Instead, their borrowing is limited to 10 percent increments of the current house price. So, a borrower may choose to borrow 10%, 20%, . . . , 80% of the price of the home at the time of origination when using an agency loan, and additionally may choose 90% or 100% when using a non-agency loan.

Figure 19: Differences in housing value across age cohorts and education (2005)



Note: This figure shows average log housing value for different age cohorts and education groups using U.S. ACS data from 2005. Although there are large life-cycle differences in housing value between the ages of 25 and 40, these differences are small compared to differences in initial housing value due to educational differences.

Figure 20: Differences in housing value across age cohorts and income (2005)



Note: This figure shows average log housing value for different age cohorts and income groups using U.S. ACS data from 2005. Although there are large life-cycle differences in housing value between the ages of 25 and 40, these differences are small compared to differences in initial housing value due to income differences.

Table 10: Differences in housing value between movers and stayers

	(1)	(2)	(3)
	All ages	Age<45	Age≥45
<i>InMove_i</i>	0.0047* (0.0026)	0.0458*** (0.0032)	-0.0488*** (0.0041)
<i>OutMove_i</i>	0.0105*** (0.0027)	0.0561*** (0.0034)	-0.0379*** (0.0041)
<i>N</i>	2,439,293	685,580	1,753,713

Note: This table reports the results from regression (29) where log housing value is regressed on indicators for whether the owner recently moved, from either outside or within the metropolitan area. Controls include a quadratic in household income, and dummies for the year of the survey, the metropolitan area, and the age, employment status, education level, and race of the owner.

Table 11: Discretization of the Model

Variable	Description	Grid	# Grid Points
rfr_t	risk-free rate	{0.015, 0.025}	2
μ_t^y	mean buyer income	log 0.055	1
cll_t	conforming loan limit	{0.4, 0.45, 0.75}	3
mps_t	availability of non-agency loans	{0, 1}	2
\bar{v}_t	unobserved demand shock	$\left\{0.2 + \left(\frac{0.7-0.2}{17}\right)n\right\}_{n=0}^{n=17}$	18
y_i	buyer income	{0.08, 0.15}	2
w_i	buyer initial wealth	$\{0.025n\}_{n=0}^{n=40}$	41
w_{it}	homeowner savings	{0.02}	1
r_t	contract rate of mortgage	$\left\{0.03, 0.035, 0.04, 0.045, 0.05, 0.055, 0.06, 0.07, 0.08, 0.1\right\}$	10
b_t	original loan amount	$\{0.025n\}_{n=1}^{n=40}$	10
p	house price	$\{0.025n\}_{n=1}^{n=40}$	40
b'	initial amount to borrow	$\{0.1n \times p_h(s_t)\}_{n=1}^{n=10}$	10