Collateral Constraints and Macroeconomic Asymmetries*

Luca Guerrieri†
Federal Reserve Board

Matteo Iacoviello‡
Federal Reserve Board

May 27, 2014

Abstract

We show that occasionally binding collateral constraints on housing wealth drive an asymmetry in the relationship between house prices and economic activity. The sensitivity of macroeconomic aggregates to movements in housing prices can be large when housing wealth is low, and small when housing wealth is high. We develop this argument in a nonlinear general equilibrium model estimated with full information Bayesian methods. As collateral constraints became slack during the housing boom of 2001-2006, expanding housing wealth made little contribution to consumption growth. By contrast, the housing collapse that followed tightened the constraints and sharply exacerbated the recession of 2008-2009. The empirical relevance of this asymmetry is corroborated by the results of panel regressions on state- and MSA-level data.

KEYWORDS: Housing, Collateral Constraints, Occasionally Binding Constraints, Nonlinear Estimation of DSGE Models, Great Recession.

JEL CLASSIFICATION: E32, E44, E47, R21, R31

---

*The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. Replication codes that implement our solution technique for DSGE models with occasionally binding constraints using an add-on to Dynare are available upon request. Stedman Hood and Walker Ray performed superb research assistance.

†Luca Guerrieri, Office of Financial Stability, Federal Reserve Board, 20th and C St. NW, Washington, DC 20551. Email: luca.guerrieri@frb.gov.

‡Matteo Iacoviello, Division of International Finance, Federal Reserve Board, 20th and C St. NW, Washington, DC 20551. Email: matteo.iacoviello@frb.gov.

1 Introduction

Collateral constraints drive an asymmetry in the relationship between house prices and economic activity, and are a central mechanism to explain the collapse of the Great Recession. When housing wealth is high, collateral constraints are slack, and the sensitivity of borrowing and spending to changes in house prices is positive but not large. Conversely, when housing wealth is low, collateral constraints are tight, and borrowing and expenditures move with house prices in a more pronounced fashion. We develop and corroborate this argument in two steps. First, we construct a nonlinear general equilibrium model and estimate it with Bayesian likelihood methods. The estimated model implies that, as collateral constraints became slack during the housing boom of 2001-2006, expanding housing wealth made a small contribution to consumption growth. By contrast, the subsequent housing collapse tightened the constraints and sharply exacerbated the recession of 2008-2009. Second, we present evidence from panel regressions on state- and MSA-level data that corroborates the asymmetry inferred from the estimated model.

The starting point for our analysis is a workhorse macro model along the lines of Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007). The model features nominal price and wage rigidities, a monetary authority that uses an interest rate rule, habit formation in consumption, and investment adjustment costs. To this framework we add three main elements. First, we allow for the dual role of housing, as a durable good, and as collateral for "impatient" households. The total supply of housing is fixed, but housing reallocation takes place across "patient" and "impatient" households in response to an array of shocks which also influence the price of housing. Second, the housing collateral constraint binds only occasionally. The estimation of the model involves inferring when the collateral constraint is binding and when it is slack through observations that do not include the Lagrange multiplier for the constraint. Third, monetary policy is constrained by the zero lower bound. Our assumption that housing is in fixed supply and plays no role in production has the important advantage that the model behaves essentially like a typical model for monetary policy analysis when the borrowing constraint is slack. During these periods, housing prices passively respond to movements in the macroeconomy and only exert a negligible feedback effect on other macro variables. By contrast, when the constraint is found to be binding, the interaction of house prices with borrowing and spending decisions has a first-order effect on the macroeconomy.

We use Bayesian likelihood methods to validate the model against U.S. data. The nonlinear solution of the model allows us to capture the state-dependent effects of shocks based on whether housing wealth is high or low, and whether the zero lower bound on nominal interest

rates binds or not. We quantify the contribution of collateral constraints to business cycles by simulating a version of the model in which parameters are set so that the collateral constraints are slack for all of the agents. The analysis shows that during the 1990-1991 and the 2008-2009 recessions, as collateral constraints became binding, they exacerbated the contraction in consumption substantially. The amplification due to collateral constraints is so large in the 2008-2009 period that, in their absence, the zero lower bound would not have been reached.

The task of isolating the asymmetric effect of changes in house prices using only national data is fraught with difficulty. Figure 1 offers a first look at national house prices. It shows the evolution of U.S. house prices over the period 1976-2011. The top panel superimposes the time series of U.S. house prices and of U.S. aggregation consumption expenditures. The correlation coefficient is 0.55, a substantial but not extreme level. The bottom panel is a scatterplot of changes in consumption and changes in house prices. It highlights that most of the positive correlation seems to be driven by periods when house prices are below average, both during the 1992-1993 period, and during the 2007-2009 recession. When periods with house price decreases are included, there is a strong positive correlation between consumption and house prices. However, excluding periods with declines in house prices results in almost no correlation between consumption and house prices.

Barring the Great Recession, house price declines have been rare at the national level. Regional data exhibit greater variation in housing prices. Accordingly, we corroborate the results of the estimated general equilibrium model using a panel and cross-sectional regressions at the regional level. We verify that the asymmetries uncovered using the estimated model and the national data are just as pronounced when using regional data.[1]

For the regional analysis, we choose measures of activity to match our model counterparts for consumption, employment and credit. Part of our empirical analysis looks for instruments for house price changes as a way to isolate housing preference shocks from other shocks that are more likely to jointly move both housing and other endogenous variables, as done by Mian and Sufi (2011). In all cases, we find statistically significant differences in the reaction of the activity measure of interest to changes in housing prices depending on whether housing prices are high or low.[2]

Our analysis is related to two distinct bodies of work. Numerous recent papers develop general equilibrium models to study the nexus between financial frictions and macroeconomic

[1] We are keenly aware that house prices are endogenous both in theory and in the data. Our modeling strategy attributes most of the variation in house prices to shocks to housing preferences, as in recent work by Liu, Wang, and Zha (2013).

[2] We classify house prices as high in a particular state when house prices are above a state-specific linear trend and have experimented with alternative definitions with little change in the asymmetries uncovered.

outcomes at the national level. Our analysis of regional data builds on an expanding literature that has linked changes in measures of economic activity, such as consumption and employment, to changes in house prices.

A spate of recent papers has quantified the importance of financial shocks in exacerbating the Great Recession using a general equilibrium framework. For instance, see Del Negro, Eggertsson, Ferrero, and Kiyotaki (2011), Jermann and Quadrini (2012), Christiano, Motto, and Rostagno (2013). The common thread among these papers is that financial shocks are key drivers of the Great Recession. The occasionally binding nature of the constraints we consider sets our work apart. In our model, financial constraints endogenously become slack or binding, so that financial shocks are not required to effectively counteract or enhance an otherwise constant set of financial constraints. In this respect, our work extends the basic mechanisms in Mendoza (2010) who also considers occasionally binding financial constraints in a calibrated small open economy setting with an exogenous interest rate. Our extensions make it possible to construct quantitatively meaningful counterfactual exercises and to consider policy alternatives in an empirically validated model for the United States. We use the model to gauge the effects of policies aimed at the housing market in the context of a deep recession.

Regarding the regional analysis, other papers also point towards an prominent role for housing as collateral in influencing both consumption and employment. Recent contributions include Case, Quigley, and Shiller (2005), Campbell and Cocco (2007), Mian and Sufi (2011), Mian, Rao, and Sufi (2012), and Abdallah and Lastrapes (2012). Despite the emphasis on collateral constraints, this literature has failed to recognize that such a channel implies asymmetric relationships for house price increases and declines with other measures of aggregate activity and has not embedded this channel in a model for policy analysis.[3]

Section 2 presents a basic, partial-equilibrium model that illustrates how collateral constraints may imply an asymmetry in the relationship between house prices and consumption. Section 3 considers an empirically-validated general equilibrium model. Sections 4 and 5 describe the estimation method and results, respectively. Section 6 presents additional evidence on asymmetries in the relationship between house prices and other measures of economic activity based on state and MSA-level data. Section 7 considers an experiment which highlights how the same policy – a transfer to indebted borrowers – can have opposite effects depending on whether house prices are high or low. Section 8 concludes.

[3] Our paper is also related to the work of Lustig and van Nieuwerburgh (2010), who find that in times when US housing collateral is scarce nationally, regional consumption is about twice as sensitive to income shocks. However, the channel they emphasize – time variation in risk-sharing among regions – is different from ours.

2 The Basic Model: Collateral Constraints and Asymmetries

To fix ideas regarding the fundamental asymmetry introduced by collateral constraints, it is useful to work through a basic model and analyze its implications for how consumption responds to changes in house prices. Throughout this section, we sidestep obvious general equilibrium considerations and assume that the price of housing is exogenous. We relax all these assumptions in the full DSGE model of the next section.

Consider the problem of a household that has to choose profiles for goods consumption c_t, housing h_t, and borrowing b_t. The household's problem is to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t \left(\log c_t + \mathsf{j} \log h_t \right), \tag{1}$$

where E_0 is the conditional expectation operator. The household is subject to the following constraints:

$$\begin{aligned} c_t + q_t h_t &= y + b_t - R b_{t-1} + q_t \left(1 - \delta_h\right) h_{t-1}; & (2) \\ b_t &\leq \mathsf{m} q_t h_t; & (3) \\ \log q_t &= \rho_q \log q_{t-1} + \varepsilon_{q,t}. & (4) \end{aligned}$$

The first constraint is the budget constraint. Income y is fixed and normalized to one. The term b_t denotes one-period debt. The gross one-period interest rate is R. Housing, which depreciates at rate δ_h, has a price q_t in unit of consumption. The second constraint is a borrowing constraint that limits borrowing to a maximum fraction m of housing wealth. The third equation describes the price of housing, q_t, which follows an AR(1) stochastic process, where $\varepsilon_{q,t}$ is a zero-mean, i.i.d. process with variance σ_q^2.

Denoting with λ_t the Lagrange multiplier on the borrowing constraint, the Euler equation for consumption is given by:

$$\frac{1}{c_t} = \beta R E_t \left(\frac{1}{c_{t+1}} \right) + \lambda_t. \tag{5}$$

To develop the intuition for our result, it is useful to consider a log-linear approximation of equation (5) in a steady state without shocks. Under the assumption that $\beta R < 1$, the borrowing constraint binds and leverage (the ratio of debt to housing wealth) is at its upper bound given by the maximum loan-to-value ratio (LTV) m. In that steady state, $\overline{\lambda} > 0$, and $\overline{c} = \overline{y} - \left(\left(R - 1\right) \mathsf{m} - \delta_h \right) \overline{q} \overline{h}$. Solving this equation forward and linearizing, one obtains the following

expression for consumption in percent deviation from steady state, $\widehat{c}_t$:

$$\widehat{c}_t = -\frac{1-\beta R}{\overline{\lambda}} E_t \left(\lambda_t - \overline{\lambda} + \beta R \left(\lambda_{t+1} - \overline{\lambda}\right) + \beta^2 R^2 \left(\lambda_{t+2} - \overline{\lambda}\right) + ...\right). \tag{6}$$

Expressing the Euler equation as above shows that consumption depends negatively on current and future expected borrowing constraints. As shown by equation (3), increases in q_t will loosen the borrowing constraint. So long as they keep λ_t positive, increases and decreases in q_t will have roughly symmetric effects on c_t. However, large enough increases in q_t lead to a fundamental asymmetry. The multiplier λ_t cannot fall below zero. Consequently, large increases in q_t can bring λ_t to its lower bound and will have proportionally smaller effects on c_t than decreases in q_t. Intuitively, an impatient borrower prefers a consumption profile that is declining over time. A temporary jump in house prices will enable such a profile (high c today, low c tomorrow) without borrowing all the way up to the limit.

More formally, the household's state at time t is its housing h_{t-1}, debt b_{t-1} and the current realization of the house price q_t, and the optimal decision are given by the consumption choice $c_t = \mathcal{C}\left(q_t, h_{t-1}, b_{t-1}\right)$, the housing choice $h_t = \mathcal{H}\left(q_t, h_{t-1}, b_{t-1}\right)$ and the debt choice $b_t = \mathcal{B}\left(q_t, h_{t-1}, b_{t-1}\right)$ that maximize expected utility subject to (2) and (3), given the house price process. Figure 2 shows the optimal leverage and the consumption function obtained from the model outlined above.[4] As the figure illustrates, high house prices are associated with slack borrowing constraints, and with a lower sensitivity of consumption to changes in house prices. Instead, when household borrowing is constrained – an outcome that is more likely when house prices are low and the initial stock of debt is high – the sensitivity of consumption to changes in house prices becomes large. This idea is developed further both in the full model and in the empirical analysis to follow.

3 The Full Model: Demand Effects in General Equilibrium

To quantify the importance of the asymmetric relationship between house prices and consumption, we embed the basic mechanisms described in Section 2 in an estimated general equilibrium model. The starting point for our analysis is a workhorse macro model along the lines of Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007). The model features

[4] The policy functions depicted in Figure 2 are obtained from standard global solution methods (value function iteration). The calibrated parameters are $\beta = 0.988$, $\mathsf{j}= 0.12$, $\mathsf{m}= 0.925$, $R = 1.01$, $\delta = 0.01$. The resulting steady-state housing wealth to quarterly income ratio is 6. For the house price process we set $\rho_q = 0.96$ and $\sigma_q = 0.0175$, in order to match a standard deviation of the quarterly growth rate of house prices equal to 1.77 percent, as in the data.

nominal price and wage rigidities, a monetary authority using a Taylor rule, habit formation in consumption and investment adjustment costs. To this framework we add three main elements. First, housing has a dual role: it is a durable good (which enter the utility function separately from consumption and labor), and it serves as collateral for "impatient" households. The total supply of housing is fixed (its price varies endogenously), but housing reallocation takes place across "patient" and "impatient " households in response to an array of shocks. Second, we allow for the collateral constraint on borrowing to bind only occasionally. The estimation exercise allows us to infer when the constraint binds using observations that do not include the hidden Lagrange multiplier on the constraint.Third, in line with U.S. experience since 2008, monetary policy is potentially constrained by the zero lower bound.

Our assumption that housing is in fixed supply and plays no role in production (unlike in the work of Liu, Wang, and Zha (2013) and Iacoviello and Neri (2010)) has the important advantage that the model behaves essentially like the one in Christiano, Eichenbaum, and Evans (2005) when the borrowing constraint is found to be slack. With a slack borrowing constraint, housing prices only passively respond to movements in the macroeconomy, but play no feedback effect on other macro variables.

Below, we sketch the key features of the model. Appendix A provides additional details as well as the list of all necessary conditions for an equilibrium.

3.1 Households

Within each group of patient and impatient households, a representative household maximizes:

$$E_0 \sum_{t=0}^{\infty} \beta^t \mathsf{z}_t \left(\Gamma \log \left(c_t - \varepsilon c_{t-1} \right) + \mathsf{j}_t \log h_t - \frac{\tau}{1+\eta} n_t^{1+\eta} \right); \tag{7}$$

$$E_0 \sum_{t=0}^{\infty} \left(\beta'\right)^t \mathsf{z}_t \left(\Gamma' \log \left(c'_t - \varepsilon c'_{t-1} \right) + \mathsf{j}_t \log h'_t - \frac{\tau}{1+\eta} n_t'^{1+\eta} \right). \tag{8}$$

Variables accompanied by the prime symbol refer to patient households. The terms c, h, n are consumption, housing, and hours. The discount factors are β and β'. By definition, $\beta' < \beta$. The term j_t captures shocks to housing preferences, and the term z_t captures a shock to intertemporal preferences. These shocks follows:

$$\log \mathsf{j}_t = \left(1 - \rho_j\right) \log \bar{\mathsf{j}} + \rho_j \log \mathsf{j}_{t-1} + u_{j,t}, \tag{9}$$

$$\log \mathsf{z}_t = \rho_z \log \mathsf{z}_{t-1} + u_{z,t}. \tag{10}$$

where $u_{j,t}$ and $u_{z,t}$ are n.i.i.d. processes with variance σ_j^2 and σ_z^2. Above, ε measures habits in consumption. The scaling factors $\Gamma = (1-\varepsilon)/(1-\beta\varepsilon)$ and $\Gamma' = (1-\varepsilon)/(1-\beta'\varepsilon)$ ensure that the marginal utilities of consumption are $1/c$ and $1/c'$ in the non-stochastic steady state.

Patient households maximize their utility subject to:

$$c_t + q_t h_t - b_t + i_t = \frac{w_t n_t}{X_{w,t}} - \frac{R_{t-1} b_{t-1}}{\pi_t} + q_t h_{t-1} + R_{k,t} k_{t-1} + Div_t, \tag{11}$$

where investment and capital are linked by:

$$k_t = \mathsf{a}_t \left(i_t - \phi \frac{(i_t - i_{t-1})^2}{\bar{i}} \right) + (1-\delta) k_{t-1}, \tag{12}$$

$$\log \mathsf{a}_t = \rho_k \log \mathsf{a}_{t-1} + u_{k,t} \tag{13}$$

where $u_{k,t}$ is a n.i.i.d process with variance σ_k^2. Patient agents choose consumption c_t, investment i_t, capital k_t (which depreciates at the rate δ), housing h_t (priced at q_t), hours n_t and borrowing b_t (loans if b_t is negative) to maximize utility subject to (11) and to (12). The term a_t is an investment-specific shock. Loans are set in nominal terms and yield a riskless nominal return of R_t. The real wages is w_t, the real rental rate is $R_{k,t}$. The terms $X_{w,t}$ denotes the markup (due to monopolistic competition in the labor market) between the wage paid by the wholesale firm and the wage paid to the households, which accrues to the labor unions. Finally, $\pi_t = P_t/P_{t-1}$ is the gross inflation rate, Div_t are lump-sum profits from final good firms and from labor unions. The formulation in (12) allows for convex investment adjustment costs, measured by ϕ.

Impatient households do not accumulate capital and do not own finished good firms or land. Their budget constraint is given by:

$$c'_t + q_t h'_t - b'_t = \frac{w'_t n'_t}{X'_{w,t}} + q_t h'_{t-1} - \frac{R_{t-1} b'_{t-1}}{\pi_t} + Div'_t; \tag{14}$$

Impatient households also face a borrowing constraint that limits the amount they can borrow, b'_t, to a fraction m of the house value. We start from the constraint of the basic model of Section 2 and extend it with an eye to empirical realism. Specifically, we allow for – but do not impose – the possibility that borrowing constraints adjust to reflect the market value of the housing

stock only sluggishly.[5] Accordingly, the constraint takes the form:

$$b'_t \leq \gamma \frac{b'_{t-1}}{\pi_t} + (1-\gamma)\, \mathsf{m} q_t h'_t \tag{15}$$

where γ measures the degree of inertia in the borrowing limit, and m is the steady-state loan-to-value ratio. When γ is greater than zero, such specification mimics the common practice that borrowing constraints are fully reset only for those who move or refinance their mortgage. While our model lacks the heterogeneity at the microeconomic level that could capture this phenomenon, this specification captures the empirical finding that measures of aggregate debt lag house prices movements. For instance, a regression of household mortgage debt on its lag and on housing wealth yields coefficients of 0.89 on lagged debt, and of 0.10 on housing wealth. Both coefficients are statistically significant ($t - statistics$ of 45 and 7 respectively), and the R^2 is 0.97. [6]

3.2 Firms

To allow for nominal price rigidities, the model differentiates between competitive flexible price/wholesale firms that produce wholesale goods, and a final good firm that operates in the final good sector under monopolistic competition. Wholesale firms hire labor and capital to produce wholesale goods Y_t. They solve:

$$\max \frac{Y_t}{X_{p,t}} - w_t n_t - w'_t n'_t - R_{k,t} k_{t-1}. \tag{16}$$

Above, $X_{p,t}$ is the price markup of final over wholesale goods. The production technology is:

$$Y_t = n_t^{(1-\sigma)(1-\alpha)} n_t'^{\sigma(1-\alpha)} k_{t-1}^{\alpha}. \tag{17}$$

[5] An interpretation of this form of the borrowing constraint is that, with multi-period debt contracts, the borrowing constraint on housing is reset only for households that acquire new housing goods or choose to refinance. Of course, in the face of home equity line of credits, adjustments of the borrowing constraint may also reflect lenders' perceived changes in the value of the collateral. Justiniano, Primiceri, and Tambalotti (2013), who study the determinants of household leveraging and deleveraging in a calibrated dynamic general equilibrium model, adopt an analogous specification.

[6] Mortgage debt and housing wealth are from Table B.100 of the Financial Accounts. We divide both series by the GDP deflator and detrend them with an HP filter (with $\lambda = 10,000$). The regression – over the 1980Q1–2011Q4 period – might not capture adequately the specification in the borrowing constraint for three reasons. First, the Financial Accounts data are on aggregate housing wealth – housing wealth held both by borrowers and savers. Second, the data on debt include gross mortgage debt – debt held both by borrowers without any other financial assets, and by savers who hold other financial assets alongside mortgage debt. Last, the constraint above may not bind in periods of high housing prices, thus weakening the link between housing wealth and mortgage debt.

In (17), the non-housing sector produces output with labor and capital. The parameter σ measures the labor income share that accrues in the economy to impatient households. When σ approaches zero, the model boils down to a model without collateral effects.[7]

3.3 Nominal Rigidities and Monetary Policy

There are Calvo-style price rigidities and wage rigidities in the final good sector. Final good firms buy wholesale goods Y_t from wholesale firms in a competitive market, differentiate the goods at no cost, and sell them at a markup $X_{p,t}$ over the marginal cost. The CES aggregates of these goods are converted back into homogeneous consumption and investment goods by households. Each period, a fraction $1-\theta_\pi$ of retailers set prices optimally, while a fraction θ_π cannot do so, and index prices to the steady state inflation $\overline{\pi}$. The resulting consumption-sector Phillips curve is:

$$\log\left(\pi_t/\overline{\pi}\right) = \beta E_t \log\left(\pi_{t+1}/\overline{\pi}\right) - \varepsilon_\pi \log\left(X_{p,t}/\overline{X_p}\right) + \mathtt{u}_{p,t}, \tag{18}$$

where $\varepsilon_\pi = (1-\theta_\pi)(1-\beta\theta_\pi)/\theta_\pi$ measures the sensitivity of inflation to changes in the markup, $X_{p,t}$, relative to its steady-state value, $\overline{X}_p$, whereas the term $\mathtt{u}_{p,t}$ denotes an iid price markup shock, $\mathtt{u}_{p,t} \sim N\left(0, \sigma_p^2\right)$.

Wage setting is modeled in an analogous way. Households supply homogeneous labor services to unions. The unions differentiate labor services as in Smets and Wouters (2007), set wages subject to a Calvo scheme and offer labor services to labor packers who reassemble these services into the homogeneous labor composites n_c and n'_c. Wholesale firms hire labor from these packers. The pricing rules set by the union imply the following wage Phillips curves:

$$\log\left(\omega_t/\overline{\pi}\right) = \beta E_t \log\left(\omega_{t+1}\overline{\pi}\right) - \varepsilon_w \log\left(X_{w,t}/\overline{X}_w\right) + \mathtt{u}_{w,t}, \tag{19}$$

$$\log\left(\omega'_t/\overline{\pi}\right) = \beta' E_t \log\left(\omega'_{t+1}/\overline{\pi}\right) - \varepsilon'_w \log\left(X'_{w,t}/\overline{X}'_w\right) + \mathtt{u}_{w,t} \tag{20}$$

where $\omega_t = \frac{w_t \pi_t}{w_{t-1}}$ and $\omega'_t = \frac{w'_t \pi_t}{w_{t-1}}$ denote wage inflation for each household type, and $\mathtt{u}_{w,t} \sim N\left(0, \sigma_w^2\right)$ denotes an iid wage markup shock.

Monetary policy follows a modified Taylor rule that allows for interest rate smoothing and responds to year-on-year inflation and lagged GDP in deviation from steady state, subject to

[7] We have estimated an alternative specification allowing for TFP shocks in equation 17 in a version of the model with variable utilization of capital. The results – further discussed in the Appendix – were similar to those reported here.

the zero lower bound (ZLB):

$$R_t = \max\left[1, R_{t-1}^{r_R}\left(\frac{\pi_t^A}{\overline{\pi}^A}\right)^{(1-r_R)r_\pi}\left(\frac{Y_{t-1}}{\overline{Y}}\right)^{(1-r_R)r_Y}\overline{R}_{r,t}^{1-r_R}\mathtt{u}_{r,t}\right]. \tag{21}$$

The term $\overline{R}$ is the steady-state nominal real interest rate in gross terms, and $\mathtt{u}_{r,t} \sim N(0, \sigma_r^2)$ denotes an iid monetary policy shock.[8]

4 Estimation of the Full Model

We use Bayesian estimation methods to size a subset of the deep structural parameters of the model. Importantly, we estimate the parameter that determines the share of impatient households. The parameters that are not estimated are calibrated in a standard fashion by matching steady–state targets.

4.1 Calibration

Some parameter choices are based on information complementary to the estimation sample. Their values are standard and are reported in Table 1. We set $\beta = 0.995$, implying a steady state 2% annual real interest rate. The capital share $\alpha = 0.3$ and the depreciation rate $\delta = 0.025$ implying a steady-state capital to annual output ratio of 2.1 and an investment to output ratio of 0.2. The weight on housing in the utility function j is set at 0.04, implying a ratio of housing wealth to annual output around 1.5. The maximum loan-to-value ratio m is set at 0.9. The labor disutility parameter η is set at 1, implying a unitary Frisch labor supply elasticity. The steady-state gross price and wage markups $\overline{X}_p$ and $\overline{X}_w$ are both set at 1.2. Finally, we set $\overline{\pi} = 1.005$ implying a 2% annual rate of inflation. The other parameters of the model are estimated using Bayesian estimation methods.

4.2 Priors

We select priors commonly used in the estimation of DSGE models. Our choices hew closely to those of Smets and Wouters (2007), with only minor exceptions. The prior distributions, the posterior modes and confidence intervals for the estimated parameters are reported in Table 1.

A key parameter in determining the extent of the asymmetries is the discount factor of the impatient agents, β'. Values of this parameter that fall below a certain threshold imply

[8] Year-on-year inflation (expressed in quarterly units, like the interest rate) is defined as $\pi_t^A \equiv (P_t/P_{t-4})^{0.25}$.

that impatient agents never escape the borrowing constraint. Then, the model has no asymmetries, regardless of the size of the shocks, and produces a large correlation between housing price growth and consumption growth, since the borrowing constraint holds all the time with equality. Conversely, when β' takes on higher values, closer to discount factor of patient agents, modest increases in house prices suffice to make the borrowing constraint slack (even though the constraint is expected to bind in the long run). We set the prior mean for β' at 0.9875, – corresponding to an annualized discount rate of 5 percent –, with a standard deviation of 0.0025.

4.3 Data

The estimation of the model is based on observations for six series: total real household consumption, price (GDP deflator) inflation, wage inflation, real investment, real housing prices, and the Federal Funds Rate. The observations span the period from 1985Q1 to 2011Q4. Prior to estimation we use a one-sided HP filter (with a value of $\lambda = 100,000$) in order to remove the low-frequency components of consumption, investment and housing prices. The one-sided HP filter has two advantages: first, it yields plausible estimates of the trend and the cycle for the variables we analyze. Second, as argued for instance by Stock and Watson (1999), the one-sided filter is not affected by the correlation of data points with subsequent observations. Appendix Appendix C show that our results are robust to using a linear, deterministic filter to detrend the data.

Our model features six observables and six shocks – investment-specific shocks, wage markup, price markup, monetary policy, intertemporal preferences, and preferences for housing.

4.4 Solution and Likelihood

The proliferation of state variables renders standard global solution algorithms inoperable. Moreover, the occasionally binding constraint on borrowing and the non-negativity constraint on the policy interest rate make first-order perturbation methods inapplicable. We solve the model using the piecewise linear method sketched in Appendix B and described more fully in Guerrieri and Iacoviello (2013). Essentially, depending on whether the zero lower bound binds or not, and depending on whether the collateral constraint on housing binds or not, we identify four regimes. The piecewise linear solution method involves linking the first-order approximation of the model around the same point under each regime. Importantly, the solution that the algorithm produces is not just linear, with different sets of coefficients depending on

each of the four regime, but rather, it can be highly nonlinear. The dynamics in each regime may crucially depend on how long one expects to be in that regime. In turn, how long one expects to be in that regime depends on the state vector.

The solution of the model takes the form:

$$X_t = \mathbf{P}(X_{t-1}, \epsilon_t)X_{t-1} + \mathbf{D}(X_{t-1}, \epsilon_t) + \mathbf{Q}(X_{t-1}, \epsilon_t)\epsilon_t. \tag{22}$$

The vector X_t collects all the variables in the model, except the innovations to the shock processes, which are gathered in the vector ϵ_t. The matrix of reduced-form coefficients $\mathbf{P}$ is state-dependent, as are the vector $\mathbf{D}$ and the matrix $\mathbf{Q}$. These matrices and vector are functions of the lagged state vector and of the current innovations. However, while the the current innovations can trigger a change in the reduced-form coefficients, X_t is still locally linear in ϵ_t.

We represent the solution in Equation (22) below in terms of observed series by premultiplying the state vector X_t by the matrix $\mathbf{H}_t$, which selects the observed variables. Accordingly, the vector of observed series Y_t is simply $Y_t = \mathbf{H}_t X_t$.[9]

Because the solution of our model is only locally linear in ϵ_t, we cannot use the Kalman filter to retrieve the estimates of the innovations in ϵ_t. Instead, following Fair and Taylor (1983) we recursively solve for ϵ_t, given X_{t-1} and the current realization of Y_t, the following system of non-linear equations:[10]

$$Y_t = \mathbf{H}_t\mathbf{P}(X_{t-1}, \epsilon_t)X_{t-1} + \mathbf{H}_t\mathbf{D}(X_{t-1}, \epsilon_t) + \mathbf{H}_t\mathbf{Q}(X_{t-1}, \epsilon_t)\epsilon_t. \tag{23}$$

The vector X_t contains unobserved components, so the filtering scheme requires an initialization. We assume that the initial X_0 coincides with the model's steady state and train the filter using the first 20 observations.

Given that ϵ_t is assumed to be drawn from a multivariate Normal distribution with covariance matrix Σ, a change in variables argument implies that the logarithmic transformation of

[9] The matrix that selects the observed variables is time-varying because we drop the interest rate from the observed vector at the zero lower bound. In that case, we also assume that monetary policy shock is zero, unless the notional rate implied by the model is positive when the observed rate is still zero. In that case, we select the notional rate as observed and reinstate the monetary policy shock.

[10] There is in principle the possibility of multiple solutions for ϵ_t to the extent that Equation (23) is highly nonlinear in ϵ_t. In our specific application we have found little evidence of this multiplicity. In theory, however, our approach to constructing the likelihood does not depend crucially on a one to one mapping between Y_t and ϵ_t. Standard results could be invoked to allow for a general correspondence between Y_t and ϵ_t when constructing the likelihood function.

the likelihood for the observed data Y^T can be written as:

$$log(f(Y^T)) = -\frac{T}{2}\log(\det(\mathbf{\Sigma})) - \frac{1}{2}\sum_t^T \hat{\epsilon}_t' \left(\mathbf{\Sigma}^{-1}\right) \hat{\epsilon}_t - \sum_t^T \log(|\det(\mathbf{H}_t\mathbf{Q}_t)|) \tag{24}$$

In our case, the Jacobian of the inverse transformation for the change in variables, $(\mathbf{H}_t\mathbf{Q}_t)^{-1}$, is known from the model's solution and does not require any additional calculations. By contrast, the general approach in Fair and Taylor (1983) requires the additional numerical evaluation of the Jacobian.[11]

Notice that the inverse transformation needed to form the likelihood is only given implicitly by $(\mathbf{H}_t\mathbf{Q}(X_{t-1},\epsilon_t))\epsilon_t - (Y_t - \mathbf{H}_t\mathbf{P}(X_{t-1},\epsilon_t)X_{t-1} - \mathbf{H}_t\mathbf{D}_t) = 0$. Accordingly, to size the Jacobian of this inverse transformation, we invoke the implicit function theorem. To this purpose, differentiating Equation (23) notice that:

$$\frac{\partial Y_t}{\partial \epsilon_t} = \mathbf{H}_t\mathbf{Q}(X_{t-1},\epsilon_t).$$

To proceed by implicit differentiation, we verify that the determinant of $\mathbf{H}_t\mathbf{Q}(X_{t-1},\epsilon_t)$ is nonzero. Accordingly, the implicit transformation is locally invertible and the Jacobian of the inverse transformation is given by:

$$\frac{\partial \epsilon_t}{\partial Y_t} = (\mathbf{H}_t\mathbf{Q}_t)^{-1}.$$

5 Results

5.1 Estimated Parameters

We combine the evaluation of the likelihood with prior information about the parameters in order to construct and maximize the posterior as a function of the model's parameters, given the data. We then construct the posterior density of the model's parameters using a standard random walk Metropolis-Hastings algorithm.

The posterior modes of the estimated parameters and other statistics are reported in Table 1. Crucially, we find a sizeable fraction of impatient agents, governed by σ. Our choice of prior, a diffuse Beta distribution, simply guarantees that this fraction remain bounded between 0 and 1. The posterior mode is estimated to be 0.42 and the 90% confidence interval ranges from 0.30 to 0.51. Accordingly, $\sigma = 0$, which would imply the exclusion of collateral constraints

[11] To derive Equation (24), recognize that $|\det((H_tQ_t)^{-1})| = 1/|\det(H_tQ_t)|$..

from the model, is highly unlikely. Moving to the parameters that govern nominal rigidities and monetary policy, the posterior modes for the Calvo parameters governing the frequency of price and wage adjustment are both equal to 0.92. This high degree of price and wage rigidity likely compensates the absence of real rigidities, such as variable capacity utilization and partial indexation of prices and wages. The estimated interest rate reaction function gives less weight to output and more weight to inflation than our prior, which was centered around Taylor's canonical values of 0.5 (for output, measured at an annual rate) and 1.5 (for inflation).

Before moving to the key results on the role of housing in recent business cycles, we take comfort from the fact that the key empirical properties of the model variables line up well with their data counterparts. To highlight this point, we focus on some moments in the data, and compare them from those of our estimated model.[12] In the model like in the data, the correlation between the house prices and aggregate consumption – 0.65 in the data, 0.41 in the mode – is higher than the correlation between house prices and investment – 0.40 in the data, 0.10 in the model. The volatility of aggregate consumption is 1.4 percent in the model, compared to 1.1 percent in the data. The model also captures the high volatility of house prices – 3.2 percent in the model, 3.8 percent in the data.

Finally, in variance decomposition exercises, we find that about three quarters of the house price volatility is driven by the housing preference shock. We elaborate on this point with two experiments described below. We will focus first on a comparison of positive and a negative housing shocks, and will move on to present a decomposition that highlights the role of housing shocks in the collapse of the Great Recession.

5.2 Responses to Positive and Negative Shocks

To illustrate the fundamental source of the asymmetry in the model, and to show how the larger model reproduces the key insights of the basic model, Figure 3 considers the effects of a sequence of estimated shocks to housing preferences, the process j_t in (9), which we interpret as a shock to housing demand, while all parameters held at the model's estimated posterior mode.

Between periods 1 and 8, a series of innovations to j_t are set to bring about either an increase or a decrease in house prices of 25 percent. Thereafter, the shock follows its autoregressive process. The dashed line plots a fall in house prices which reduces the collateral capacity of constrained households. Accordingly, those households can borrow less and are forced to curtail their non-housing consumption even further to satisfy their borrowing constraint. On balance,

[12] Here we follow conventional practice and compute the model statistics on simulated series, HP-filtered with a standard two-sided filter with $\lambda = 1,600$. We apply the same two-sided, HP filter also to our observables.

the peak decline in aggregate consumption is about 3 percent. The model's short-run nominal and real rigidities imply that the decline in aggregate consumption translates into a decline in the firms' demand for labor. In equilibrium, the drop in hours worked (not shown) comes close to reaching 3 percent below the balanced growth path. In reaction to a shock of this magnitude, the zero lower bound on interest rates is not attained, therefore the asymmetric responses are only driven by the collateral constraint.

The solid lines plot the responses to a shock of the same magnitude and profile but with opposite sign. In this case, house prices increase 25 percent. Recalling the partial equilibrium model described in Section 2, an increase in house prices can relax borrowing constraints. Accordingly, the borrowing constraint for the impatient household becomes slack, and the Lagrange multiplier in borrowing constraint bottoms out at zero, before taking on positive values again as house prices return to the baseline. When the constraint is slack, the borrowing constraint channel remains operative only in expectation. Thus, impatient households discount that channel more heavily the longer the constraint is expected to remain slack. As a consequence, the response of consumption to the large house price increases considered in the figure is not as dramatic as the reaction to house price declines of an equal magnitude. At peak, the increase both in consumption is 1.5 percent, half as big as the response to the house price decline.

In experiments not reported here, we have found modest asymmetries for other shocks that affect house prices and consumption in our general equilibrium model. These shocks are likely to generate significant asymmetries only insofar as they affect house prices or collateral capacity. However, the asymmetry that we uncover is independent of the particular stochastic structure for the model, and needs not rely on housing demand shocks only. Potentially, in any housing model with occasionally binding constraints, one can find substantial asymmetries as long as the model can match the observed swings in house prices.

5.3 The Asymmetric Contribution of Housing to Business Cycles

Figure 4 shows a counterfactual exercise in which the parameter σ is set to 0, so that all households are patient. By construction, when we feed back the estimated sequence of shocks into the original model, we can get an exact match for the observed data. This is not the case for the counterfactual model. Housing prices are still matched, since housing services are essentially priced by the unconstrained households. However, the evolution of consumption is markedly different from the data in the counterfactual model. When the Lagrange multiplier on the collateral constrained is estimated to be binding, as in the 2008-2009 recession, a large

gap opens up between the observed and counterfactual consumption levels.

Quantitatively, our model estimates imply that, absent collateral constraint, the decline in consumption between 2007 and the end of 2009 would have only been 2 percent. The observed decline was more than three times larger, thus implying that collateral effects can account for more than two thirds of the observed consumption decline. Remarkably, in the absence of collateral constraints the recession would have been curbed to such an extent that the Federal Funds rate would not have reached zero.[13] By contrast, when the Lagrange multiplier on the collateral constraint is estimated to be slack, there is little difference between the counterfactual level of consumption without collateral constraints and the observed level of consumption. We interpret this result as evidence that for most of the housing boom that ended in 2006, the rise in house prices did relatively little to boost consumption.

Figure 5 provides an additional angle to compare our model against a model without collateral constraints. Once more, to exclude the collateral constraints we imposed $\sigma = 0$, but this time we re-estimated the restricted model. Figure 5 highlights the effects of different patterns of shocks needed to match the data by the baseline model with occasionally binding collateral constraints and by the restricted model. The top panels in the figure compare the evolution of consumption and housing prices when only housing shocks are turned on. For both models, the evolution of housing prices are in line with observed housing prices. However, the two models differ drastically in their implications for consumption. Whereas the baseline model comes close to matching the evolution of both housing and consumption with just the housing shocks, housing shocks have little bearing on consumption for the model without the collateral constraints. The bottom panels of Figure 5 compare the evolution of consumption and housing prices from the two models when only consumption preference shocks are turned on. These panels highlight that the restricted model is completely dependent on a sequence of consumption shocks to match the consumption data. Accordingly, the proliferation of shocks that the restricted model calls for is corroborated by a posterior odds ratio of 7 to 1 overwhelmingly in favor of the model with collateral constraints.

In sum, we find it more compelling to argue that a decline in house prices, coupled with a weakening of households balance sheets, were the main culprits for the collapse in consumption during the 2008-2009 recession. By contrast, our results show that a model that excludes

[13] In our sample, the interest rate is at zero from 2009Q1 until the end of sample (2011Q4). According to the estimated model, the interest rate prescribed by the Taylor rule would be at zero only until 2011Q1. The estimated model identifies expansionary monetary shocks in 2011Q2, Q3 and Q4, which bring the model in line with the data. The expansionary shocks occur around the time when the FOMC statements become increasingly explicit in announcing that the committee expected the funds rate to remain at zero for extended periods of time.

collateral constraints has to rely on a contagious attack of patience to explain the depth of the Great Recession.

6 Regional Evidence on Asymmetries

Our model estimated on national- level data motivates additional empirical analysis that we conduct using a panel of data from U.S. states and Metropolitan Statistical Areas (MSA). The advantage of these data is that variation in house prices and economic activity is greater at the regional than at the aggregate level, as documented for instance by Del Negro and Otrok (2007), who find a large degree of heterogeneity across states in regard to the relative importance of the national factors.[14] The use of regional data also allays the concern that little can be learned from national data, given the rarity of declines in house prices at the national level. Note that, in any event, the state-level series aggregated back to the national level track their National Income and Product Accounts (NIPA) counterparts rather well.[15]

To set the stage, Figure 6 shows changes in house prices and several measures of activity, namely changes in employment in the service sector, auto sales, electricity consumption, and mortgage originations. The figure focuses on two points in time, 2005 and 2008 for all the 50 U.S. states and the District of Columbia. For each state there are two dots in each panel: the green dot (concentrated in the north–east region of the graph) shows the lagged percent change in house prices and the percent change in the indicator of economic activity in 2005, at the height of the housing boom.[16] The red dot represents analogous observations for the 2008 period, in the midst of the housing crash. Fitting a piecewise linear regression to these data yields a correlation between house prices and activity that is smaller when house prices are high. This evidence on asymmetry is bolstered by the large cross-sectional variation in house prices across states over the period in question.

[14] In the sample period we analyze, the first principal component for annual house price growth accounts for 64 percent of the variance of house prices across the 50 U.S. states and the District of Columbia. The corresponding numbers for employment in the service sector, auto sales, electricity consumption, and mortgage originations are respectively 73, 90, 44, and 89 percent.

[15] For instance, over the sample period, the correlation between NIPA motor vehicle consumption growth (about 1/3 of total durable expenditure) and retail auto sales growth is 0.89; and the correlation between services consumption growth and electricity usage growth is 0.54.

[16] An analogous relationship is more tenuous for house prices and employment in the manufacturing goods sector. Most goods are traded and are less sensitive to local house prices than services.

6.1 State-Level Evidence

We use annual data from the early 1990s to 2011 from the 50 U.S. states and the District of Columbia on house prices and measures of economic activity. We choose measures of economic activity to match our model counterparts for consumption, employment and credit.

Our main specification takes the following form:

$$\Delta \log y_{i,t} = \alpha_i + \gamma_t + \beta_{POS} \mathcal{I}_{i,t} \Delta \log hp_{i,t-1} + \beta_{NEG} (1 - \mathcal{I}_{i,t}) \Delta \log hp_{i,t-1} + \delta X_{i,t-1} + \varepsilon_{i,t} \quad (25)$$

where $y_{i,t}$ is an index of economic activity and $hp_{i,t}$ is the inflation-adjusted house price index in state i in period t; α_i and γ_t represent state and year fixed effects; and $X_{i,t}$ is a vector of additional controls. We interact changes in house prices with a state-specific indicator variable $\mathcal{I}_{i,t}$ that, in line with the model predictions, takes value 1 when house prices are high, and value 0 when house prices are low. We classify house prices as high in a particular state when house prices are above a state-specific linear trend separately estimated for the 1976-2011 period, a classification that lines up with the findings of the estimated model in Figure 4. Using this approach, the fraction of states with high house prices is about 20 percent in the 1990s, rising gradually to peak at 100 percent in 2005 and 2006, and dropping to 27 percent at the end of the sample. Our results were similar using two alternative definitions of $\mathcal{I}_{i,t}$. Under the first alternative definition, $\mathcal{I}_{i,t}$ equals 1 when real house price inflation is positive. Under the second definition, $\mathcal{I}_{i,t}$ equals 1 when the ratio of house prices to income is high relative to its trend (in log). In our baseline specification, we use one-year lags of house prices and other controls to control for obvious endogeneity concerns. Our results were also little changed when instrumenting current or lagged house prices with one or more lags.

Tables 2 to 4 present our estimates when the indicators of economic activity $y_{i,t}$ are employment in the service sector, auto sales, and electricity usage respectively.

Table 2 presents the results when the measure of regional economic activity is employment in the non-tradeable service sector. We choose this measure (rather than total employment) since U.S. states (and MSAs) trade heavily with each other, so that employment in sectors that mainly cater to the local economy better isolates the local effects of movements in local house prices.[17] The first two columns do not control for time effects. They show that the asymmetry

[17] The BLS collects state-level employment data by sectors broken down according to NAICS (National Industry Classification System) starting from 1990. According to this classification (available at http://www.bls.gov/ces/cessuper.htm), the goods-producing sector includes Natural Resources and mining, construction and manufacturing. The service-producing sector includes wholesale trade, retail trade, transportation, information, finance and insurance, professional and business services, education and health services, leisure and hospitality and other services. A residual category includes unclassified sectors and public adminis-

is strong and economically relevant, and that house prices matter, at statistically conventional levels, both when high and when low. After controlling for time effects in the third column, the coefficient on high house prices is little changed, but the coefficient on low house prices is lower. A large portion of the declines in house prices in our sample took place against the background of the zero lower bound on policy interest rates. As discussed in the model results, the zero lower bound is a distinct source of asymmetry for the effect of change in house prices. Time fixed effects allow us to parse out the effects of the national monetary policy reaching the zero lower bound and, in line with our theory, compress the elasticity of employment to low house prices. In the last two columns, after adding additional variables, the only significant coefficient is the one on low house prices. In column five, the coefficient on high house prices is positive, although it is low and not significantly different from zero. The coefficient on low house prices, instead, is positive and significantly different from zero. These results imply that house prices only matter for economic activity when they are low. The difference in the coefficient on low and high house prices is significantly different from zero.

Table 3 reports our results when the measure of activity is retail automobile sales. Auto sales are an excellent indicator of local demand, since autos are almost always sold to state residents, and since durable goods are notoriously sensitive to business cycles. After adding lagged car sales and personal income as controls, the coefficients on low and high house prices are both positive, but the coefficient on low house prices (estimated at 0.2) is nearly three times as large.

Table 4 reports our results using residential electricity usage as a proxy for consumption. Even though electricity usage only accounts for 3 percent of total consumption, we take electricity usage to be a useful proxy for nondurable consumption.[18] Most activities involve the use of electricity, and electricity cannot be easily stored. Accordingly, the flow usage of electricity may even provide a better measures of the utility flow derived from a good than the actual purchase of the good. Even in cases when annual changes in weather conditions may affect year-on-year consumption growth, their effect can be easily filtered out using state-level observations on heating and cooling degree days, which are conventional measures of weather-driven electricity demand. We use these weather measures as controls in all specifications reported. As the table shows, in all regressions low house prices affect consumption growth more than high house prices. After time effects, lagged income growth and lagged consumption growth

tration. We exclude from the service sector wholesale trade (which on average accounts for about 6 percent of total service sector employment) since wholesale trade does not necessarily cater to the local economy.

[18] Da and Yun (2010) show that using electricity to proxy for consumption produces asset pricing implications that are consistent with consumption-based capital asset pricing models.

are controlled for (last column), the coefficient on high house prices is 0.11, the coefficient on low house prices is nearly twice as large at 0.18, and their difference is statistically larger than 0 at the 10 percent significance level.

Because the effects of low and high house prices on consumption work in our model through tightening or relaxing borrowing constraints, it is important to check whether measures of leverage also depend asymmetrically on house prices. We perform these checks and report the results in Appendix E, which confirms that mortgage originations depend asymmetrically on house prices too.

6.2 MSA-Level Evidence

Tables 5 and 6 present the results of evidence across MSAs. MSAs account for about 80 percent of the population and of employment in the entire United States. In Table 5, the results from the MSA-level regressions reinforce those obtained at the state level. After controlling for income, lagged employment and time effects, the elasticities of employment to house prices are 0.05 and 0.09 when house prices are high and low, respectively. These elasticities are larger than those found at the state level.

A legitimate concern with the panel and time-series regressions discussed so far is that the correlation between house prices and activity could be due to some omitted factor that simultaneously drives both house prices and activity. Even if this were the case, our regressions would still be of independent interest, since – even in absence of a causal relationship – they would indicate that comovement between house prices and activity is larger when house prices are low, as predicted by the model.

To support claims of causality, one needs to isolate exogenous from endogenous movements in house prices. In Table 6, we follow the methodology and insight of Mian and Sufi (2011) and use data from Saiz (2010) in an attempt to distinguish an independent driver of housing demand that better aligns with its model counterpart. The insight is to use the differential elasticity of housing supply at the MSA level as an instrument for house prices, so as to disentangle movements in housing prices due to general changes in economic conditions from movements in the housing market that are directly driven by shifts in housing demand in a particular area. Because such elasticity is constant over time, we cannot exploit the panel dimension of our dataset, and instead use the elasticity in two separate periods by running two distinct regressions of car sales on house prices. The first regression is for the 2002-2006 housing boom period, the second for the 2006-2010 housing bust period. In practice, we rely on the following

differenced instrumental variable specifications:

$$\log hp_t - \log hp_s = b_0 + b_1 \text{ Elasticity} + \varepsilon_b \tag{26}$$

$$\log car_t - \log car_s = c_0 + c_1(\widehat{\log hp_t - \log hp_s}) + \varepsilon_c \tag{27}$$

where $s = 2002$ and $t = 2006$ in the first set of regressions, and $s = 2006$ and $t = 2010$ in the second set.

The first stage, OLS regressions show that elasticity is a powerful instrument in driving house prices, with an R^2 from the first stage regression around 0.20 in both subperiods.[19] The second stage regressions show how car sales respond to house prices dramatically more in the second period, in line with the predictions of the model and with the results of the panel regressions. In the $2002 - 2006$ period, the elasticity of car sales to house prices is 0.24. In the $2006 - 2010$ period, in contrast, this elasticity doubles to 0.46.

Using a higher level of data disaggregation (ZIP-code level data instead of MSAs) and a sample that runs from 2007 to 2009, Mian, Rao, and Sufi (2012) find a large elasticity (equal to 0.74) of auto sales to housing wealth during the housing bust, in line with our findings. Importantly, they also find that this elasticity is smaller in zip codes with a high fraction of non-housing wealth to total wealth. One interpretation of their result – in line with our model – is that households in zip codes with high non-housing wealth might be, all else equal, less likely to face binding borrowing constraints during periods of housing price declines, because they can use other forms of wealth to support their consumption plans.

7 Debt Relief and Borrowing Constraints

So far, our theoretical and empirical results show that movements in house prices can produce asymmetries that are economically and statistically significant. We now consider whether these asymmetries are also important for gauging the effects of policies aimed at the housing market in the context of a deep recession. To illustrate our ideas, we choose a simple example of one such policy, a lump-sum transfer from patient (saver) households to impatient (borrower) households. This policy could mimic voluntary debt relief from the creditors, or a scheme where interest income is taxed and interest payments are subsidized in lump-sum fashion, so that the end result is a transfer of resources from the savers to the borrowers.

We consider this experiment against two different baselines. In one case, house prices are

[19] The F statistics on the first stage regressions are 69.1 and 67.2 for the first and the second period respectively, well above the conventional threshold of 10 for evaluating weak instruments.

assumed to be declining; in the other case, housing prices are assumed to be increasing. The baseline housing price changes are brought about by the same preference shocks considered in Figure 3 and discussed at length above.

Figure 7 shows the cumulative response of house prices to the baseline housing preference shocks and to two transfer shocks from saver households to borrower households. Both transfer shocks are unforeseen. They are sized at the same 1 percent of steady-state total consumption in both cases. Each transfer is governed by an auto-regressive process of order 1, with coefficient equal to 0.5. The first transfer starts in period 1. A series of unforeseen innovations to the shock process phases in the transfer, until it reaches a peak of 1 percent of steady-state consumption. Then, the auto-regressive component of the shock reduces the level of the transfer back to 0. The first transfer happens against a background of housing price declines and tight borrowing constraints. The second transfer, starting in period 50, mimics the first but happens against a baseline with housing price increases and slack borrowing constraints.

The top left panel of Figure 7 shows house prices in deviation from their steady-state level. The path shown is almost identical to the one in Figure 3 because the transfer shocks only have a negligible effect on house prices. The transfer payments are timed to coincide with the series of housing preference shocks that reduce house prices.

The remaining panels in Figure 7 show responses of key variables to the transfer shock in deviation from the baseline path that occurs with the housing preference shock only. Thus, those panels isolate the partial effects of the transfer shocks. The consumption response of borrower households is dramatically different depending on the baseline variation in house prices. When house prices decline, the borrowing constraint is tight and the marginal propensity to consume of borrower households is elevated. When house prices increase, the borrowing constraint becomes slack and the marginal propensity to consume of borrower households drops down closer to that for saver households. In reaction to the transfer, consumption of the savers declines less, and less persistently, against a baseline of housing price declines. In that case, there are expansionary spillover effects from the increased consumption of borrowers to aggregate hours worked and output. Taking together the responses of savers and borrowers, the partial effects of the transfer on aggregate consumption are sizable when house prices are low, and small when house prices are elevated. As a consequence, actions such as mortgage relief can almost pay for themselves through their expansionary effects on economic activity in a scenario of binding borrowing constraints.

8 Conclusions

Numerous recent papers with an empirical focus have emphasized the importance of household debt and the housing market in understanding the Great Recession. Our model provides a framework to analyze these results. The estimated model explains why housing prices matter more during severe recessions and allows the assessment of costs and benefits of alternative policies aimed at restoring the efficient functioning of the housing market.

Our empirical and theoretical results indicate that policy measures aimed at the housing market can produce outsize spillovers to aggregate consumption in periods when collateral constraints are tight. These spillovers are likely to be larger than those that can be found in samples dominated by house price increases, because these periods can severely underpredict the sensitivity of consumption to movements in housing wealth.

Throughout the paper, we have emphasized the role of housing as collateral for households, and on the effects of changes in housing wealth on consumption. However, the mechanism at the heart of our argument has even broader applicability. For instance, to the extent that fixed assets are used for collateral by entrepreneurs, local governments, or exporters, the asymmetries highlighted here for consumption could also be relevant for fixed investment, government spending, or the trade balance.[20]

[20] See for instance Adelino, Schoar, and Severino (2013), Chaney, Sraer, and Thesmar (2012) for investment; Barboza (2011) for government spending; Klapper, Laeven, and Rajan (2012) for trade credit.

Figure 1: House Prices and Consumption in U.S. National Data

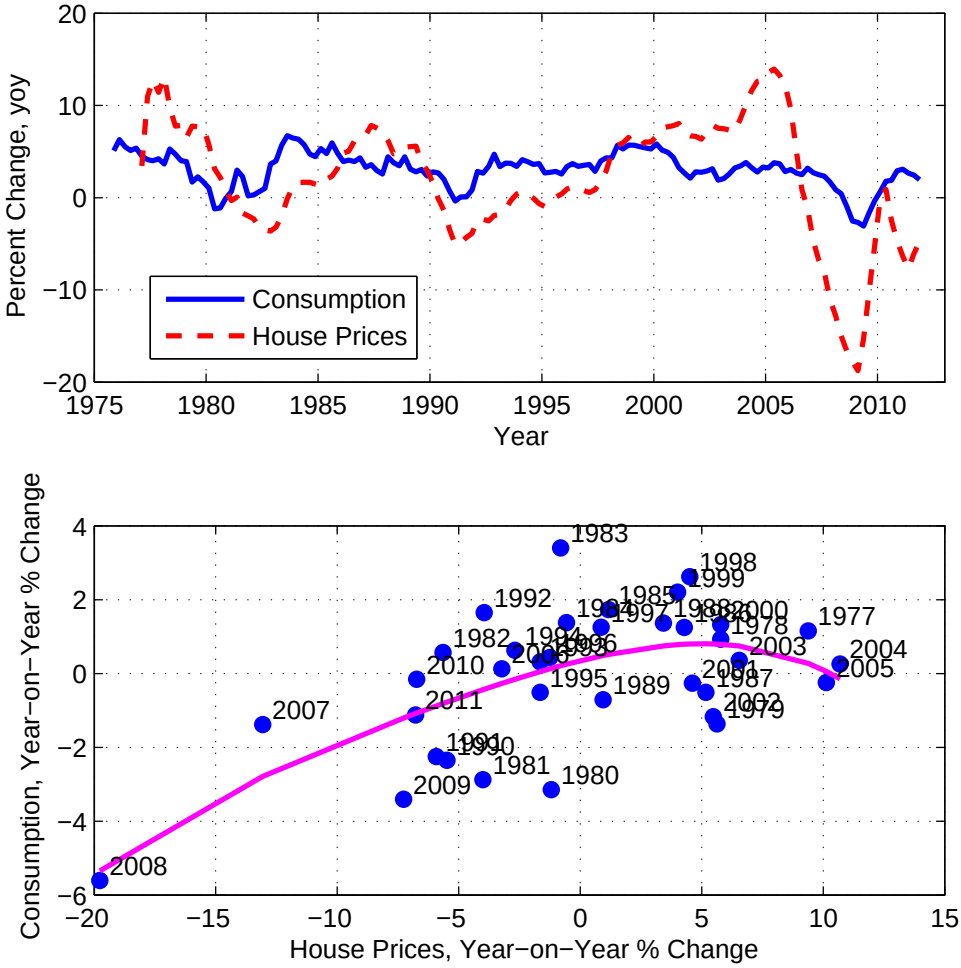

Note: Data sources are as follows. House Prices: Loan Performance National House Price Index (SA), Haver Analytics, USLPHPIS@USECON, divided by the GDP deflator (DGDP@USECON). Consumption: Real Personal Consumption Expenditures, from Department of Commerce, Bureau of Economic Analysis (CH@USECON). In the bottom panel, consumption growth and house price growth are expressed in deviation from their sample mean. The data sample is from 1976Q1 to 2011Q4.

Figure 2: House Prices and Consumption in the Basic Model

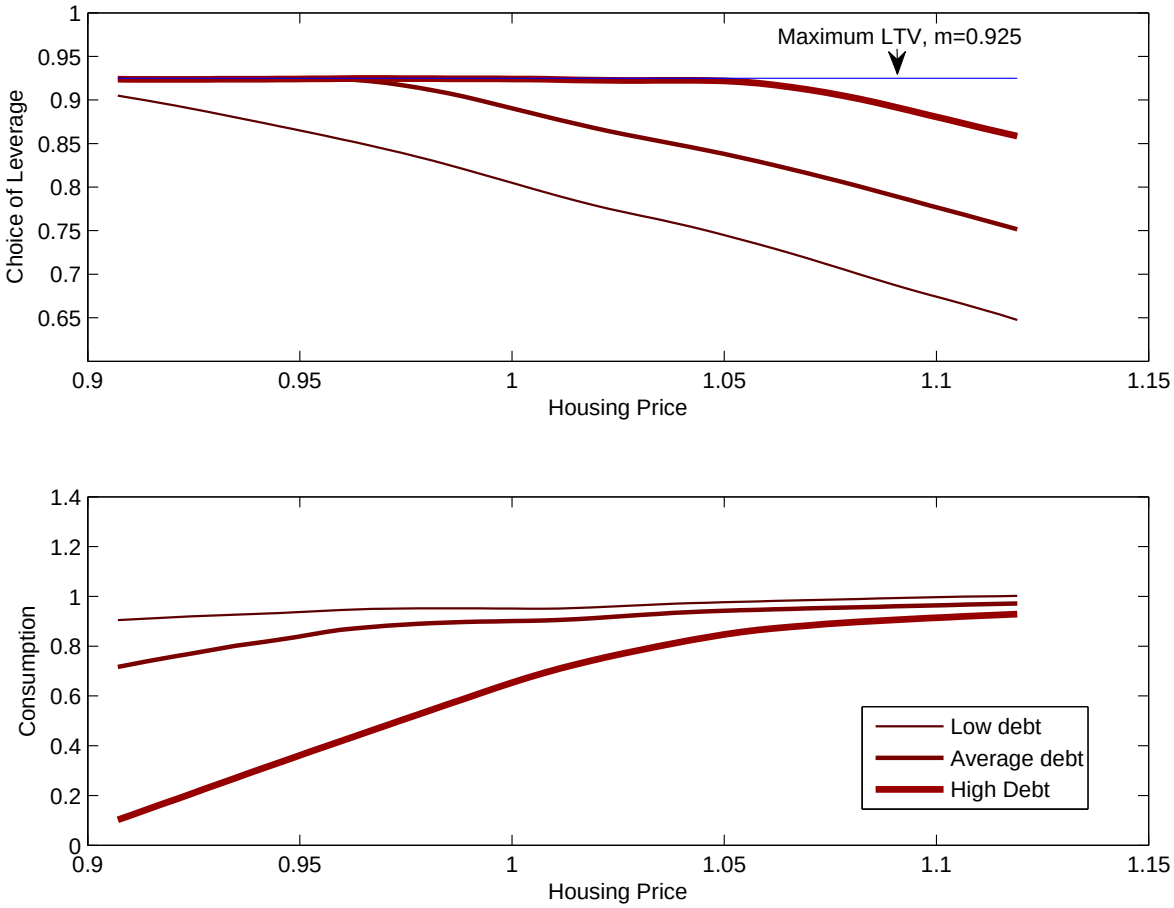

Note: Optimal leverage choice and optimal consumption as a function of the housing price for three different levels of debt, low, normal and high, when housing is at its nonstochastic steady-state value. In the top panel, low levels of house prices move the household closer to the maximum borrowing limit given by $\mathsf{m}= 0.925$. This is more likely to happen at high levels of debt (thick line). In the bottom panel, the higher house prices are, the more likely is the household not to be credit constrained, and the consumption function becomes flatter. At high levels of debt, the household is constrained for a larger range of realizations of house prices, and the consumption function is steeper when house prices are low.

Figure 3: Impulse Responses to Positive and Negative Housing Demand Shocks in the Full DSGE Model

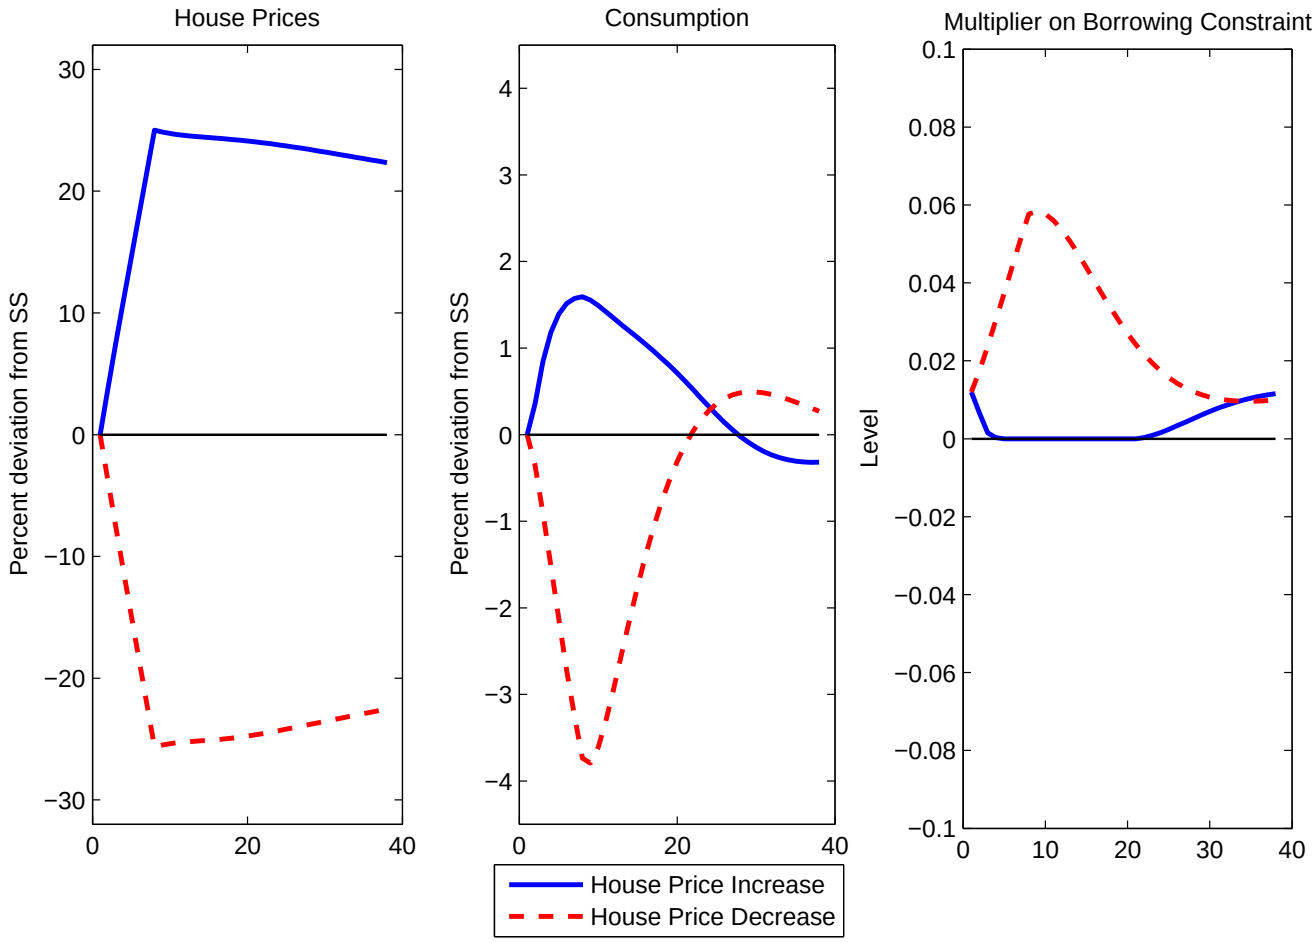

Note: Horizontal axis: horizon in quarters. The simulation shows the dynamic responses to sequence of housing demand shocks of equal size but opposite sign that move house prices up (solid lines) and down (dashed lines) by 25 percent relative to the steady state.

Figure 4: Historical Simulation of the Estimated Model

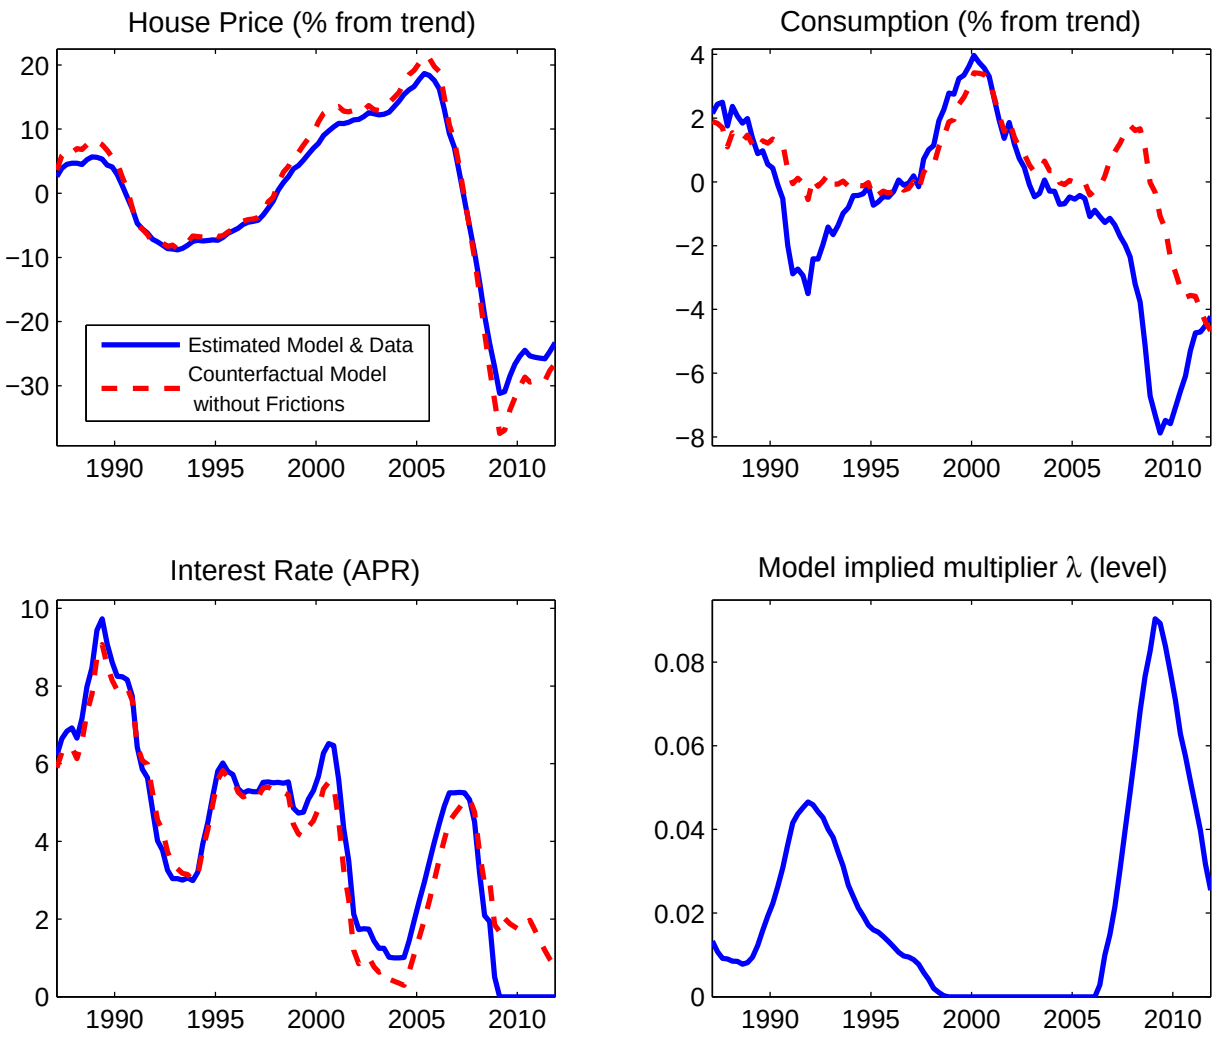

Note: The simulation shows the filtered series for house prices, consumption, interest rates and the Lagrange multiplier in the estimated model. The dashed lines show their counterfactual paths feeding in the same shocks but in absence of constrained households (setting σ=0)

Figure 5: Counterfactual Consumption Paths in the Estimated Model

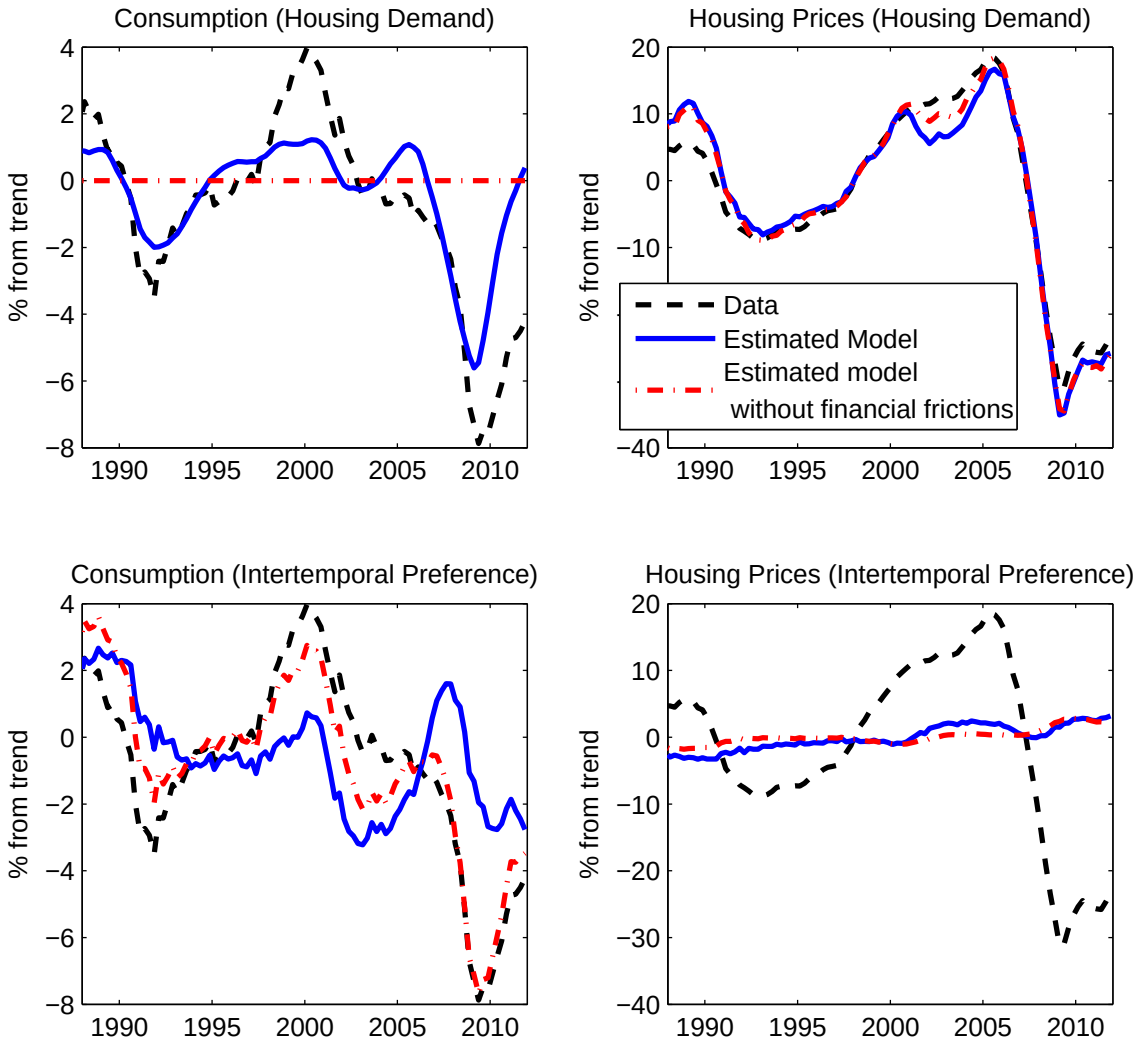

Note: The top panels compare consumption and house prices in the data (dashed line) with their model counterparts when consumption and house prices are driven by the housing preference shock. The dash-dotted line shows the counterfactual model paths for the estimated model with the restriction that $\sigma = 0$. The bottom panels redo the exercise for intertemporal preference shocks.

Figure 6: House Prices and Economic Activity by State

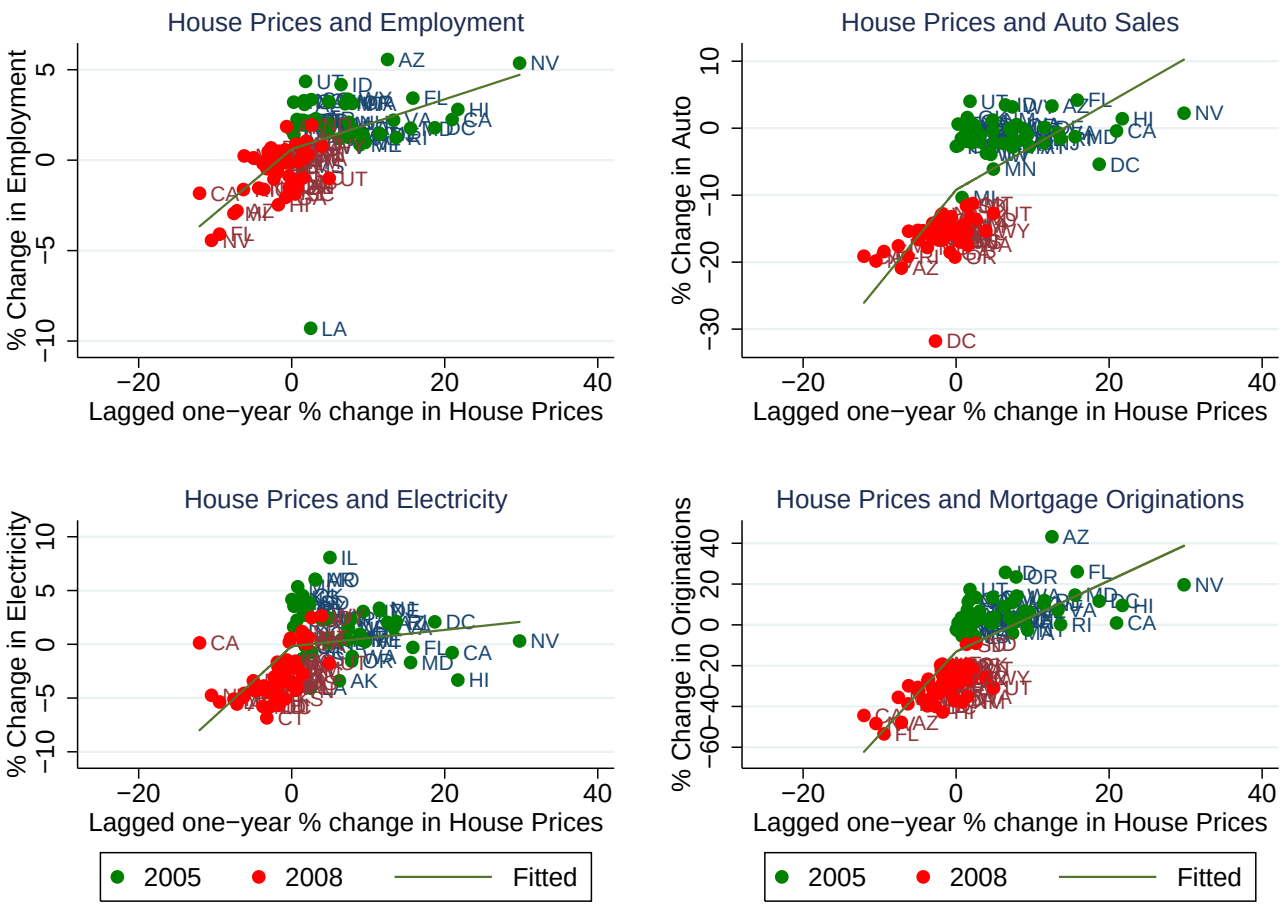

Note: Each panel shows house price growth and activity growth across US states in 2005 and 2008. The "fitted" line shows the fitted values of a regression of activity growth on house prices growth broken down into positive and negative changes.

Figure 7: Transfer from Lenders to Borrowers with Low and High House Prices

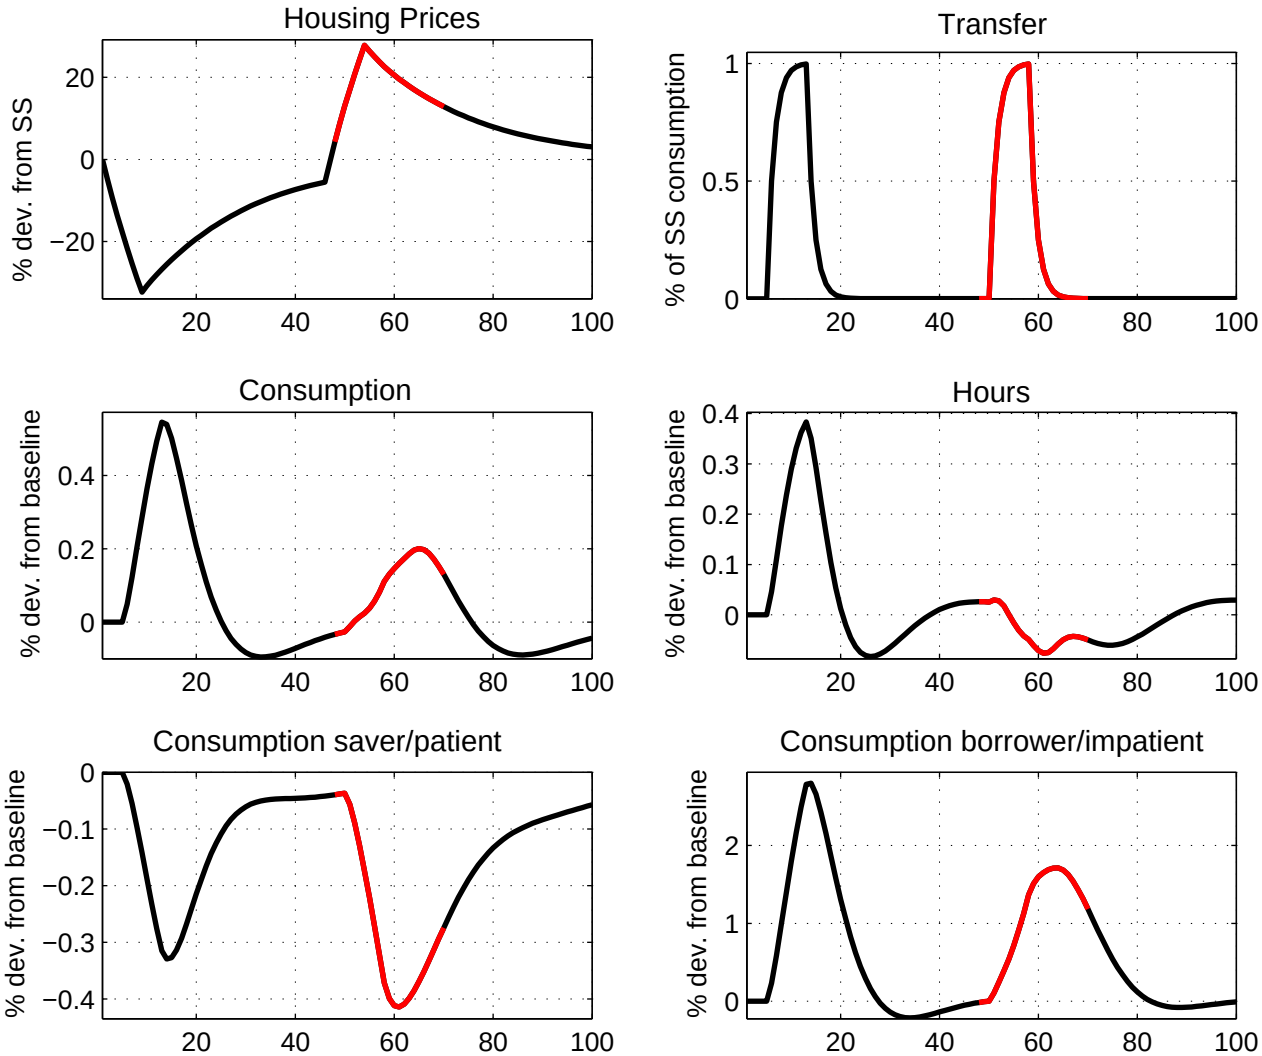

Note: Two unexpected lump-sum transfers from savers to borrowers sized at 1 percent of steady-state total consumption. The first transfer (periods 1-8) happens against a baseline of low house prices and tight collateral constraints. The second transfer (periods 51-58) happens against a baseline of high house prices and slack collateral constraints. Housing price changes in the baseline stem from a housing preference shock. The responses of consumption, hours, savers' and borrowers' consumption are shown in deviation from baseline to isolate the partial effect of the transfer shocks. Variables are plotted in red when the constraint is slack.

Table 1: Parameter Values

	Calibrated Parameters	Value
m	Maximum LTV	0.9
η	labor disutility	1
β	discount factor, patient agents	0.995
$\overline{\pi}$	steady-state gross inflation rate	1.005
α	capital share in production	0.3
δ	capital depreciation rate	0.025
$\overline{j}$	housing weight in utility	0.04
X_p	average price markup	1.2
X_w	average wage markup	1.2

	Estimated Parameters	Prior type, mean, std	Posterior			
			Mode	5%	Median	95%
β'	discount factor, impatient	normal, 0.9875, .0025	**0.9888**	0.9844	0.9886	0.9927
ε	habit in consumption	beta, 0.5, 0.1	**0.6062**	0.5132	0.6205	0.7238
ϕ	investment adjustment cost	gamma, 5, 2	**6.6388**	3.9151	6.2061	9.257
σ	wage share	beta, 0.5, 0.15	**0.4258**	0.332	0.4216	0.5529
r_π	inflation resp. Taylor rule	normal, 1.5, 0.25	**1.6695**	1.4467	1.6817	1.9804
r_R	inertia Taylor rule	beta, 0.75 0.1	**0.7098**	0.6317	0.6999	0.7602
r_Y	output response Taylor rule	beta, 0.125 0.025	**0.0544**	0.0332	0.0543	0.0883
θ_π	Calvo parameter, prices	beta, 0.5, 0.075	**0.9209**	0.8967	0.9187	0.937
θ_w	Calvo parameter, wages	beta, 0.5, 0.075	**0.9242**	0.9021	0.9206	0.938
γ	inertia borrowing constraint	beta, 0.5, 0.2	**0.4871**	0.3317	0.4958	0.673
ρ_j	AR(1) housing shock	beta, 0.5, 0.2	**0.9943**	0.9697	0.9887	0.9974
ρ_K	AR(1) investment shock	beta, 0.5, 0.2	**0.8171**	0.7442	0.7955	0.8533
ρ_Z	AR(1) intertemporal shock	beta, 0.5, 0.2	**0.854**	0.7235	0.816	0.8907
σ_j	std. housing demand shock	invgamma, 0.01, 1	**0.0407**	0.0294	0.0612	0.1377
σ_K	std. investment shock	invgamma, 0.01, 1	**0.0448**	0.0292	0.0444	0.0653
σ_P	std. price markup shock	invgamma, 0.01, 1	**0.0029**	0.0027	0.0031	0.0036
σ_R	std. interest rate shock	invgamma, 0.01, 1	**0.0018**	0.0015	0.0018	0.0021
σ_W	std. wage markup shock	invgamma, 0.01, 1	**0.0098**	0.0087	0.01	0.0115
σ_Z	std. intertemporal shock	invgamma, 0.01, 1	**0.0137**	0.0112	0.0142	0.0186

Note: Calibrated and Estimated Parameters for the Full Model.

Table 2: State-Level Regressions: Employment in Services and House Prices

	% Change in Employment (Δemp_t)				
Δhp_{t-1}	**0.14*** **				
	(0.01)				
Δhp_high_{t-1}		**0.07*** **	**0.08*** **	**0.03***	**0.02**
		(0.01)	(0.01)	(0.02)	(0.01)
Δhp_low_{t-1}		**0.24*** **	**0.12*** **	**0.08*** **	**0.07*** **
		(0.02)	(0.02)	(0.02)	(0.02)
Δemp_{t-1}				**0.26*** **	**0.23*** **
				(0.08)	(0.09)
$\Delta income_{t-1}$					**0.07** **
					(0.03)
pval difference		**0.000**	**0.100**	**0.013**	**0.017**
Time effects	**no**	**no**	**yes**	**yes**	**yes**
Observations	1071	1071	1071	1020	1020
States	51	51	51	51	51
R-squared	0.12	0.16	0.66	0.72	0.73

Note: Regressions using annual observations from 1991 to 2011 on 50 States and the District of Columbia. Robust standard errors in parenthesis. ***,**,*: Coefficients statistically different from zero at 1, 5 and 10% confidence level, respectively. pval is the p-value of the test for difference between low-house price and high-house prices coefficient.

Data Sources and Definitions: Δhp is the inflation–adjusted (using the GDP deflator) percent change in the FHFA House Price Index. Δemp is the percent change in employment in the Non-Tradable Service Sector which includes: Retail Trade, Transportation and Utilities, Information, Financial Activities, Professional and Business Services, Education and Health Services, Leisure and Hospitality, and Other Services (source: BLS Current Employment Statistics: Employment, Hours, and Earnings - State and Metro Area). $\Delta income$ is the percent change in the inflation–adjusted state-level disposable personal income (source: Bureau of Economic Analysis).

Table 3: State-Level Regressions: Auto Sales and House Prices

	% Change in Auto Sales ($\Delta auto_t$)				
Δhp_{t-1}	**0.24*** **				
	(0.03)				
Δhp_high_{t-1}		**-0.05**	**0.16*** **	**0.11*** **	**0.07** **
		(0.04)	(0.04)	(0.03)	(0.03)
Δhp_low_{t-1}		**0.62*** **	**0.33*** **	**0.27** **	**0.20** **
		(0.05)	(0.06)	(0.11)	(0.09)
$\Delta auto_{t-1}$				**0.23**	**0.21**
				(0.17)	(0.17)
$\Delta income_{t-1}$					**0.34*** **
					(0.11)
pval difference		**0.000**	**0.040**	**0.137**	**0.155**
Time effects	**no**	**no**	**yes**	**yes**	**yes**
Observations	969	969	969	918	918
States	51	51	51	51	51
R-squared	0.02	0.06	0.86	0.87	0.88

Note: State–level Regressions using annual observations from 1992 to 2011 on 50 States and the District of Columbia. Robust standard errors in parenthesis. ***,**,*: Coefficients statistically different from zero at 1, 5 and 10% confidence level, respectively. pval is the p-value of the test for difference in the coefficients for low-house prices and high-house prices.

Data Sources and Definitions: $\Delta auto$ is the percent change in inflation–adjusted auto sales, "Retail Sales: Motor vehicle and parts dealers" from Moody's Analytics Database. Auto sales data are constructed with underlying data from the US Census Bureau and employment statistics from the BLS. The two Census Bureau surveys are the quinquennial Census of Retail Trade, a subset of the Economic Census, and the monthly Advance Retail Trade and Food Services Survey. See Table 2 for other variable definitions.

Table 4: State-Level Regressions: Electricity Consumption and House Prices

	% Change in Electricity Consumption ($\Delta elec_t$)				
Δhp_{t-1}	**0.11*** **				
	(0.02)				
Δhp_high_{t-1}		**0.03**	**0.09*** **	**0.14*** **	**0.12*** **
		(0.02)	(0.02)	(0.03)	(0.03)
Δhp_low_{t-1}		**0.24*** **	**0.16*** **	**0.22*** **	**0.19*** **
		(0.03)	(0.03)	(0.04)	(0.04)
$\Delta elec_{t-1}$				**-0.41*** **	**-0.41*** **
				(0.02)	(0.02)
$\Delta income_{t-1}$					**0.15*** **
					(0.05)
pval difference		**0.000**	**0.105**	**0.058**	**0.090**
Time effects	**no**	**no**	**yes**	**yes**	**yes**
Weather Controls*	**yes**	**yes**	**yes**	**yes**	**yes**
Observations	1071	1071	1071	1020	1020
States	51	51	51	51	51
R-squared	0.04	0.04	0.08	0.12	0.12

Note: State–level Regressions using annual observations from 1990 to 2011 on 50 States and the District of Columbia. Robust standard errors in parenthesis. ***,**,*: Coefficients statistically different from zero at 1, 5 and 10% confidence level, respectively. pval is the p-value of the test for difference in the coefficients for low-house prices and high-house prices.

Data Sources and Definitions: $\Delta elec$ is the percent change in Residential Electricity Consumption (source: the U.S. Energy Information Administration's Electric Power Monthly publication. Electricity Power Annual: Retail Sales - Total Electric Industry - Residential Sales, NSA, Megawatt-hours). See Table 2 for other variable definitions. All regressions in the Table control separately for number of heating degree days and number of cooling degree days in each state (source: U.S. National Oceanic and Atmospheric Administration's National Climatic Data Center).

Table 5: MSA Level: Employment in Services and House Prices

	% Change in Employment (Δemp_t)				
Δhp_{t-1}	**0.134*** **				
	(0.006)				
Δhp_high_{t-1}		**0.104*** **	**0.058*** **	**0.049*** **	**0.048*** **
		(0.008)	(0.007)	(0.008)	(0.008)
Δhp_low_{t-1}		**0.183*** **	**0.099*** **	**0.095*** **	**0.094*** **
		(0.009)	(0.008)	(0.010)	(0.010)
Δemp_{t-1}				**0.033**	**0.031**
				(0.041)	(0.041)
$\Delta income_{t-1}$					**0.021***
					(0.011)
pval difference		**0.0000**	**0.0003**	**0.0001**	**0.0000**
Time effects	no	no	yes	yes	yes
Observations	5390	5390	5390	5147	5147
MSA	262	262	262	262	262
R-squared	0.09	0.10	0.37	0.39	0.39

Note: MSA–level Regressions using annual observations from 1992 to 2011 on 262 MSAs (102 MSAs were dropped since they had incomplete or missing data on employment by sector). Robust standard errors in parenthesis. ***,**,*: Coefficients statistically different from zero at 1, 5 and 10% confidence level, respectively. pval is the p-value of the test for difference in the coefficients for low-house prices and high-house prices.

Data Sources and Definitions: $\Delta income$ is the percent change in MSA-level inflation-adjusted personal income (source: BEA, Local and Metro Area Personal Income Release). For employment (Δemp) and house prices (Δhp), see Table 2.

Table 6: MSA Level: Auto Registrations and House Prices

	Cross-sectional Regressions			
	Sample 2002-2006 (Housing Boom)		Sample 2006-2010 (Housing Bust)	
	Δhp	Δcar	Δhp	Δcar
Elasticity	**-7.26*** **		**4.69*** **	
	(0.87)		(0.57)	
Δhp		**0.24*** **		**0.49*** **
		(0.06)		(0.08)
Method	OLS	IV	OLS	IV
Observations	254	254	254	254
R-squared	0.22	0.35	0.21	0.48

Note: Regressions using Housing supply Elasticity at the MSA level as an instrument for house prices in a regression of MSA car registrations on MSA house prices. ***,**,*: Coefficients statistically different from zero at 1, 5 and 10% confidence level, respectively. The housing supply elasticity is taken from Saiz (2010) and measures limits on real-estate development due to geographic factors that affect the amount of developable land, as well as factors like zoning restrictions. The elasticity data are available for 269 cities: we dropped 15 areas because they were covering primary metropolitan statistical areas (PMSA), which are portions of metropolitan areas, rather than complete MSAs.

Data Sources: Car Registrations are retail (total less rental, commercial and government) auto registrations from Polk Automotive Data. Δcar is the percent change in car registrations. See Table 2 for other data sources.

References

Abdallah, C. S. and W. D. Lastrapes (2012). Home equity lending and retail spending: Evidence from a natural experiment in texas. *American Economic Journal: Macroeconomics 4*(4), 94–125. [4]

Adelino, M., A. Schoar, and F. Severino (2013). House prices, collateral and self-employment. Technical report, National Bureau of Economic Research. [24]

Barboza, D. (2011, October 20). China to allow some local governments to issue bonds. *The New York Times*. [24]

Campbell, J. and J. Cocco (2007). How do house prices affect consumption? evidence from micro data. *Journal of Monetary Economics 54*, 591–621. [4]

Case, K., J. Quigley, and R. Shiller (2005). Comparing wealth effects: The stock market versus the housing market. *Advances in Macroeconomics 5*(1), 1235–1235. [4]

Chaney, T., D. Sraer, and D. Thesmar (2012). The collateral channel: How real estate shocks affect corporate investment. *The American Economic Review 102*(6), 2381–2409. [24]

Christiano, L., R. Motto, and M. Rostagno (2013, January). Risk Shocks. NBER Working Papers 18682, National Bureau of Economic Research, Inc. [4]

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy 113*(1), 1–45. [2, 6, 7]

Da, Z. and H. Yun (2010). Electricity consumption and asset prices. *Available at SSRN 1608382*. [20]

DeJong, D. w. C. (2007). *Structural Macroeconometrics*. Princeton: Princeton University Press. [46]

Del Negro, M., G. Eggertsson, A. Ferrero, and N. Kiyotaki (2011). The great escape? A quantitative evaluation of the Fed?s liquidity facilities. Technical report. [4]

Del Negro, M. and C. Otrok (2007). 99 luftballons: Monetary policy and the house price boom across us states. *Journal of Monetary Economics 54*(7), 1962–1985. [18]

Fair, R. C. and J. B. Taylor (1983, July). Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models. *Econometrica 51*(4), 1169–85. [13, 14]

Fernald, J. (2012). A quarterly, utilization-adjusted series on total factor productivity. *Manuscript, Federal Reserve Bank of San Francisco*. [47]

Guerrieri, L. and M. Iacoviello (2013). Occbin: A toolkit for solving dynamic models with occasionally binding constraints. Manuscript, Federal Reserve Board. [12, 40, 42, 43]

Iacoviello, M. and S. Neri (2010). Housing market spillovers: Evidence from an estimated dsge model. *American Economic Journal: Macroeconomics 2*(2), 125–64. [7]

Jermann, U. and V. Quadrini (2012). Macroeconomic effects of financial shocks. *American Economic Review 102*(1), 238–71. [4]

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2013, April). Household leveraging and deleveraging. NBER Working Papers 18941, National Bureau of Economic Research, Inc. [9]

Klapper, L., L. Laeven, and R. Rajan (2012). Trade credit contracts. *Review of Financial Studies 25*(3), 838–867. [24]

Liu, Z., P. Wang, and T. Zha (2013). Land-price dynamics and macroeconomic fluctuations. *Econometrica 81*(3), 1147–1184. [3, 7]

Lustig, H. and S. van Nieuwerburgh (2010). How much does household collateral constrain regional risk sharing? *Review of Economic Dynamics 13*(2), 265 – 294. [4]

Mendoza, E. G. (2010). Sudden Stops, Financial Crises, and Leverage. *American Economic Review 100*(5), 1941–66. [4]

Mian, A. and A. Sufi (2011). House prices, home equity-based borrowing, and the us household leverage crisis. *American Economic Review 101*(5), 2132–56. [3, 4, 21]

Mian, A. R., K. Rao, and A. Sufi (2012). Household balance sheets, consumption, and the economic slump. Mimeo, Princeton University. [4, 22]

Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics 125*(3), 1253–1296. [21, 37]

Smets, F. and R. Wouters (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American Economic Review 97*(3), 586–606. [2, 6, 10, 11]

Stock, J. H. and M. W. Watson (1999). Forecasting inflation. *Journal of Monetary Economics 44*(2), 293–335. [12]

Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters 20*(2), 177–181. [44]

Uhlig, H. (1995). A toolkit for analyzing nonlinear dynamic stochastic models easily. Technical report, Tilburg University, Center for Economic Research. [43]

Appendix

Appendix A Equilibrium Conditions of the Full Model

We summarize here the equations describing the equilibrium of the full model. We use the methods described in Appendix B and more fully developed in Guerrieri and Iacoviello (2013) to solve the model subject to the two occasionally binding constraints.

Let $u_{c,t}$ and $u_{c',t}$ and $u_{h,t}$ and $u_{h',t}$ and $u_{n,t}$ and $u_{n',t}$ denote respectively the time-t marginal utility of consumption, marginal utility of housing and marginal disutility of labor (inclusive of the shock terms: that is, $u_t = \mathsf{z}_t \left(\Gamma \log \left(c_t - \varepsilon c_{t-1} \right) + \mathsf{j}_t \log h_t - \frac{\tau}{1+\eta} n_t^{1+\eta} \right)$, and $u_{c,t}$ is the derivative of u_t with respect to c_t). Let Δ be the first difference operator. The complete system of equations is given by:

- Budget constraint for the patient:

$$c_t + q_t \Delta h_t + ik_t - b_t = \frac{w_t n_t}{X_{w,t}} - \frac{R_{t-1} b_{t-1}}{\pi_t} + R_{k,t} k_{t-1} + Div_t. \tag{A.1}$$

- Capital Accumulation equations for the patient:

$$\begin{aligned} u_{c,t} q_{k,t} \left(1 - \phi \Delta ik_t\right) &= u_{ct} - \beta u_{c,t+1} q_{k,t+1} \phi \Delta ik_{t+1} && \text{(A.2)} \\ u_{c,t} q_{k,t} &= \beta u_{c,t+1} \left(rk_{t+1} + q_{k,t+1} \left(1 - \delta_k\right)\right) && \text{(A.3)} \\ k_t &= ik_t + \left(1 - \delta_k\right) k_{t-1} - \phi_k \left(k_t - k_{t-1}\right)^2 && \text{(A.4)} \end{aligned}$$

- Other optimality conditions for the patient:

$$u_{c,t} = \beta u_{c,t+1} \left(R_t / \pi_{t+1} \right) \tag{A.5}$$

$$\frac{w_t}{X_{w,t}} u_{ct} = u_{nt} \tag{A.6}$$

$$q_t u_{c,t} = u_{h,t} + \beta E_t q_{t+1} u_{c,t+1} \tag{A.7}$$

- Budget and Borrowing Constraint and optimization conditions for the impatient:

$$c'_t + q_t \Delta h'_t + \frac{R_{t-1}}{\pi_t} b_{t-1} = \frac{w'_t}{X'_{w,t}} n'_t + b_t + Div'_t \tag{A.8}$$

$$(1-\lambda_t)\, u_{c,t} = \beta' E_t \left(\frac{R_t - \phi\lambda_{t+1}}{\pi_{t+1}} u_{c,t+1} \right) \tag{A.9}$$

$$\frac{w'_t}{X'_{w,t}} u_{c,'t} = u_{n',t} \tag{A.10}$$

$$b'_t \leq \gamma \frac{b'_{t-1}}{\pi_t} + (1-\gamma)\, \mathsf{m}q_t h'_t \tag{A.11}$$

$$q_t u_{c',t} = u_{h',t} + \beta' q_{t+1} u_{c',t+1} + u_{c,'t} \lambda_t (1-\phi)\, \mathsf{m}q_t \tag{A.12}$$

- Firm Problem and Aggregate production and Phillips curves:

$$Y_t = n_t^{(1-\sigma)(1-\alpha)} n_t'^{\sigma(1-\alpha)} K_{t-1}^{\alpha} \tag{A.13}$$

$$(1-\alpha)(1-\sigma) Y_t = X_{p,t} w_t n_t \tag{A.14}$$

$$\begin{aligned} (1-\alpha)\sigma Y_t &= X_{p,t} w'_t n'_t && \text{(A.15)} \\ \alpha Y_t &= X_{p,t} R_{k,t} K_{t-1} && \text{(A.16)} \end{aligned}$$

$$\begin{aligned} \log\left(\frac{\pi_t}{\overline{\pi}}\right) &= \beta E_t \log\left(\frac{\pi_{t+1}}{\overline{\pi}}\right) - \varepsilon_\pi \log\left(\frac{X_{p,t}}{\overline{X}_p}\right) && \text{(A.17)} \\ \log\left(\frac{\omega_t}{\overline{\pi}}\right) &= \beta E_t \log\left(\frac{\omega_{t+1}}{\overline{\pi}}\right) - \varepsilon_w \log\left(\frac{X_{w,t}}{\overline{X}_w}\right) && \text{(A.18)} \\ \log\left(\frac{\omega'_t}{\overline{\pi}}\right) &= \beta' E_t \log\left(\frac{\omega'_{t+1}}{\overline{\pi}}\right) - \varepsilon'_w \log\left(\frac{X'_{w,t}}{\overline{X}'_w}\right) && \text{(A.19)} \end{aligned}$$

$\omega_t = \frac{w_t \pi_t}{w_{t-1}}$ and $\omega'_t = \frac{w'_t \pi_t}{w_{t-1}}$is wage inflation for each household type, and where $\varepsilon_w = (1-\theta_w)(1-\beta\theta_w)/\theta_w$, $\varepsilon'_w = (1-\theta_w)(1-\beta'\theta_w)/\theta_w$, and

$\varepsilon_p = (1-\theta_\pi)(1-\beta\theta_\pi)/\theta_\pi$.

- Monetary policy:

$$R_t = \max\left(1, R_{t-1}^{r_R} \left(\frac{\pi_t^A}{\overline{\pi}^A}\right)^{(1-r_R)r_\pi} \left(\frac{Y_{t-1}}{\overline{Y}}\right)^{(1-r_R)r_Y} \overline{R}^{1-r_R} \mathtt{u}_{r,t}\right) \tag{A.20}$$

- Definitions/Market clearing:

$$\begin{aligned} Y_t &= c_t + c'_t + k_t - (1-\delta) k_{t-1} \\ H &= 1 = h_t + h'_t && \text{(A.21)} \end{aligned}$$

Equations A.1 to A.21, with the definitions for Div_t and Div'_t below and the functional forms and the laws of motion for the exogenous shocks, define a system of 21 equations in the following variables: c, c', h, h', ik, k, Y, b, n, n', w, w', π, q, R, λ, X_p, X_w, X'_w, R_k, q_k.

We use the methods described in Appendix B and more fully developed in Guerrieri and Iacoviello (2013) to solve the model subject to the two occasionally binding constraints given by equations A.11 and A.20.

Appendix B Solution Method for the Full Model

We use a piecewise linear solution approach to find the equilibrium allocations for a given sequence of unforeseen shocks. This method resolves the problem of computing decision rules that approximate the equilibrium well both when the borrowing constraint binds and when it does not (similar reasoning applies to the nonnegativity constraint on the interest rate, as described at the end of this Section).

The economy features two regimes: a regime when collateral constraints bind; and a regime in which they do not, but are expected to bind in the future.[21] With binding collateral constraints, the linearized system of necessary conditions for an equilibrium can be expressed as

$$\mathcal{A}_1 E_t X_{t+1} + \mathcal{A}_0 X_t + \mathcal{A}_{-1} X_{t-1} + \mathcal{B} u_t = 0, \tag{B.1}$$

where $\mathcal{A}_1$, $\mathcal{A}_0$, and $\mathcal{A}_{-1}$ are matrices of coefficients conformable with the vector X collecting the model variables in deviation from the steady state for the regime with binding constraints; and where u is the vector collecting all shock innovations (and $\mathcal{B}$ is the corresponding conformable matrix). Similarly, when the constraint is not binding, the linearized system can be expressed as

$$\mathcal{A}^*_1 E_t X_{t+1} + \mathcal{A}^*_0 X_t + \mathcal{A}^*_{-1} X_{t-1} + \mathcal{B}^* u_t + \mathcal{C}^* = 0, \tag{B.2}$$

where $\mathcal{C}^*$ is a vector of constants. When the constraint binds, we use standard linear solution methods to express the decision rule for the model as

$$X_t = \mathcal{P} X_{t-1} + \mathcal{Q} u_t. \tag{B.3}$$

When the collateral constraints do not bind, we use a guess-and-verify approach. We shoot back towards the initial conditions, from the first period when the constraints are guessed to

[21] If one assumes that the constraints are not expected to bind in the future, the regime with slack borrowing constraints becomes unstable, since borrowers' consumption falls over time and their debt rises over time until it reaches the debt limit, which contradicts the initial assumption.

bind again. For example, if the constraints do not bind in t but are expected to bind the next period, the decision rule for period t can be expressed, starting from B.2 and using the result that $E_t X_{t+1} = \mathcal{P} X_t$, as:

$$X_t = -\left(\mathcal{A}_1^* \mathcal{P} + \mathcal{A}_0^*\right)^{-1} \left(\mathcal{A}_{-1}^* X_{t-1} + \mathcal{B}^* u_t + \mathcal{C}^*\right). \tag{B.4}$$

We proceed in a similar fashion to compute the allocations for the case when collateral constraints are guessed not to bind for multiple periods or when they are foreseen to be slack starting in periods beyond t. As shown by equation B.4, the model dynamics when constraints are not binding depend both on the current regime (through the matrices $\mathcal{A}_1^*, \mathcal{A}_0^*$ and $\mathcal{A}_{-1}^*$) and on the expectations of future regimes when constraints will bind again (through the matrix $\mathcal{P}$, which is a nonlinear function of the matrices $\mathcal{A}_1$, $\mathcal{A}_0$ and $\mathcal{A}_{-1}$).[22]

It is straightforward to generalize the solution method described above for multiple occasionally binding constraints.[23] The extension is needed to account for the zero lower bound (ZLB) on policy interest rates as well as the possibility of slack collateral constraints. In that case, there are four possible regimes: 1) collateral constraints bind and policy interest rates are above zero, 2) collateral constraints bind and policy interest rates are at zero, 3) collateral constraints do no bind and policy interest rates are above zero, 4) collateral constraints do not bind and policy interest rates are at zero. Apart from the proliferation of cases, the main ideas outlined above still apply.

Appendix C Accuracy of Solution Method for the Full Model, tested using the Basic Model

In the absence of an analytical solution for the models considered in this paper, we assess the solution algorithm used to solve the full general equilibrium model by comparing its performance against standard solution methods. As is well understood, standard global methods are subject to the curse of dimensionality, which renders such methods inoperable for our application. However, the partial equilibrium model of Section 2 of the paper can be solved with both our piecewise-linear algorithm, described in Appendix B, and with standard global solution methods. We use this smaller model to showcase the performance of our solution algorithm.

[22] See for instance Uhlig (1995).

[23] For an array of models, Guerrieri and Iacoviello (2013) compare the performance of the piecewise perturbation solution described above against a dynamic programming solution obtained by discretizing the state space over a fine grid. Their results show that this solution method efficiently and quickly computes a solution that closely mimics the (perfect-foresight) nonlinear solution.

Among standard global methods, we focus on value function iteration since it is reliable, accurate, and well understood.[24] Overall, we find that key aspects of the global solution obtained through value function iteration are matched by the solution from the piecewise-linear algorithm. A key advantage of our algorithm is that it can handle the solution of a model, such as the one described in Section 3 of the paper, for which the curse of dimensionality renders standard global solution methods infeasible.

In Figures A.1 and A.2 we compare the simulated paths for house prices, consumption, leverage and debt using alternative solution methods. In Figure A.1, we consider impulse responses to negative and positive house price shocks. In Figure A.2, we generate a realization of house prices drawing shock innovations for 50 periods from the stochastic AR(1) process described in the text.

The "piecewise-linear" lines are computed using our method. The "nonlinear stochastic" lines refer to the nonlinear model solution obtained using global methods (value function iteration) under the assumption that the agents know and act upon the future distribution of the random shocks. The "nonlinear deterministic" lines refer to the perfect foresight case, solved using global methods under the assumption that agents ignore the future variance of shocks (that is, each period they expect that future shock innovations will equal zero with probability one). Finally, the "linear" lines refer to the model solved used brute force linearization under the – counterfactual – assumption that the borrowing constraint is always binding.

As can be seen from the figures, the nonlinear methods (value function iteration) and the piecewise linear method deliver very similar dynamics for the variables of interest. The similarity of the simulation paths causes the business cycle statistics (reported in Table A.1) to be in broad agreement for those two methods. As expected, leverage and debt are on average lowest in the full stochastic case, since buffer stock motives – ignored by construction or by design in the other cases – cause agents to save more and reduce indebtedness. However, our method – which combines first-order perturbation solutions under two different regimes – comes remarkably close to matching the dynamics of the full nonlinear method under perfect foresight. As first-order perturbation solutions ignore the possibility of future shocks, it is not surprising that our piecewise-linear method would not be able to capture precautionary motives present in the full stochastic non-linear solution. Accordingly, we consider the comparison with a perfect-foresight non-linear solution as more apt. By contrast, the linearized solution that assumes that the constraint is always binding cannot capture the asymmetry of consumption and grossly overestimates its volatility.

[24] Our state variables are the level of debt, the housing stock and the house price process. We discretize the AR(1) house price process with using Tauchen's method (Tauchen (1986)) with 101 grid points. We pick a solution range for housing and debt between -60 and $+60$ percent of their steady state values, discretized over 100 points for debt and 110 points for housing. In between iterations, we use Howard's improvement step. We verified that increasing the number of grid points did not materially change any of the results.

As a further metric to judge to accuracy of our solution method, the last column of Table A.1 reports the welfare cost for a household of using the approximated policy functions instead of the nearly-exact one (which we take to be the solution obtained via value function iteration) in order to solve the problem. The welfare cost of using the piecewise linear policy function is small (about 0.02% of lifetime consumption), and is one order of magnitude smaller than the cost of using the linearized policy function.

Appendix D Additional Details on Estimation

Local linearity of the Policy Functions. The solution of the model can be described by a policy function of the form:

$$X_t = \mathbf{P}(X_{t-1}, \epsilon_t)X_{t-1} + \mathbf{D}(X_{t-1}, \epsilon_t) + \mathbf{Q}(X_{t-1}, \epsilon_t)\epsilon_t. \tag{28}$$

The vector X_t collects all the variables in the model, except the innovations to the shock processes, which are gathered in the vector ϵ_t. The matrix of reduced-form coefficients $\mathbf{P}$ is state-dependent, as are the vector $\mathbf{D}$ and the matrix $\mathbf{Q}$. These matrices and vector are functions of the lagged state vector and of the current innovations. However, while the the current innovations can trigger a change in the reduced-form coefficients, X_t is still locally linear in ϵ_t. To illustrate this point, Figure A.3 shows how the policy function for borrowers' consumption c'_t – one of the elements of X_t – depends on the realization of the housing preference shock $u_{j,t}$ – one of the elements of ϵ_t – when all the other elements of X_{t-1} are held at their steady–state value. The top panel shows the consumption function for the impatient (in deviation from the nonstochastic steady state): this function is piecewise linear, with each of the rays corresponding to a given number of periods in which the borrowing constraint is expected to be slack. The bottom panel shows the derivative of the consumption function with respect to to $u_{j,t}$. As the consumption function is piecewise linear, the derivative is not defined at the threshold values of the shock $u_{j,t}$ that change the expected duration of the regime. However, each of these threshold points for different shocks is a set of measure zero.[25]

Realizations of the shock $u_{j,t}$ above a threshold will imply that the borrowing constraints is temporarily slack. When the constraint is slack, the constraint will be expected to be slack for a number of periods which increases with the size of the shock. Accordingly, consumption will respond proportionally less, and the the $\mathbf{Q}_{c',u_j}$ element of the matrix $\mathbf{Q}$ that defines the impact

[25] It is straightforward to prove that the points in which the derivative of the decision rule is not defined is of measure zero given a choice of process for the stochastic innovations. By construction there are only countably many of these points. If there were uncountably many, a shock could lead to a permanent switch in regimes, which is ruled out by the solution method.

sensitivity of c' to u_j will be smaller.

Robustness to Initial Conditions. We assume that all variables are known and equal to their nonstochastic steady state in the first period, and use the first 20 observations as a training sample for our filtering procedure. By treating the initial distribution of X_0 as known (and equal to its steady–state value), we can eliminate the conditionality of the likelihood function for the observed data Y^T on both X_0 and Y_0. Without this assumption, one needs to integrate the likelihood for Y^T over the distribution for X_0 implied by the specification of the model and the observed data, as discussed for instance by DeJong (2007), and methods such as the particle filter become necessary.

As a robustness exercise, we have estimated our model under different assumptions about the values of the initial state vector X_0. Given a training sample of 20 periods, the initial conditions were essentially irrelevant by period 20, and our estimated parameters were little affected by the value of the initial condition.

Retrieving the Shocks. Our algorithm relies on using a nonlinear equation solver in order to filter in each period t, given X_{t-1}, the sequence of shocks ϵ_t that reproduces the observed behavior of the observables in the vector Y_t. It is possible that small numerical errors in retrieving ϵ_t at each point in time may propagate over time and lead to inaccuracies in computing the filtered shocks. To explore the practical relevance of this possibility, we generate an artificially long sample of observables from our model. At the estimated mode, the borrowing constraint is slack about 30 percent of the time and the ZLB binds about 1 percent of the time. In order to emphasize the asymmetries, we hold all the parameters at their estimated mode, but set the estimated discount factor of impatient agents at 0.9925 (a value slightly higher than its estimated value): with this parameter choice, the borrowing constraint is slack 40% of the time. Drawing from the posterior mode of the shocks, we generate a time series of artificial observables of length $T = 500$. We then use our procedure to filter these shocks back and compare the filtered shocks to the "true" ones used to generate our artificial data set. The correlation between the "true" shocks and the filtered ones is, for all shocks, extremely high, ranging from 0.9996 for the monetary shock to 0.99999998 for the inflation shock.

Detrending Method. We use a one-sided HP filter to detrend the data prior to estimation. As an alternative, we have incorporated deterministic trend growth in the model and estimated the parameters governing the trend jointly with the other model parameters. The results were little affected by this alternative assumption. The model with deterministic trends implies slightly more persistent and more volatile shocks, presumably in order to account for the larger and more persistent deviations of the observables around their constant trends.

Allowing for TFP Shocks. In our baseline specification we include for six observables (inflation, wages, house prices, consumption, investment and the interest rate) and six shocks (investment-specific shocks, wage markup, price markup, monetary policy, intertemporal preferences, and preferences for housing). As a robustness exercise, we have included utilization-adjusted TFP (constructed by Fernald (2012)) in the list of observables and allowed for a seventh, TFP shock, in a model with variable capacity utilization. The estimated model with TFP shocks prefers a slightly lower degree of price rigidity than our benchmark (the Calvo parameters for prices is 0.89 instead of 0.92, and the Calvo parameter for wages is 0.9 instead of 0.92), and prefers a higher fraction of impatient households (0.63 instead of 0.43). These estimates have two effects on the role of the housing collapse in the Great Recession. The first is that smaller nominal rigidities dampen the effect of housing demand shocks. The second is that a higher wage share of impatient households enhances them. All told, the second effect dominates, and the role of the housing collapse in explaining the decline in consumption during the Great Recession is slightly larger than in our benchmark. Most of the other results in the main text are unaffected.

Appendix E State-Level Evidence on House Prices and Mortgage Originations

Because the effects of low and high house prices on consumption work in our model through tightening or relaxing borrowing constraints, it is important to check whether measures of leverage also depend asymmetrically on house prices. Table A.2 shows how mortgage originations at the state level respond to changes in house prices. We choose mortgage originations because they are available for a long time period, and because they are a better measure of the flow of new credit to households than the stock of existing debt. In all of the specifications in Table A.2, mortgage originations depend asymmetrically on house prices, too.

Figure A.1: Accuracy of Solution Method: Impulse Responses

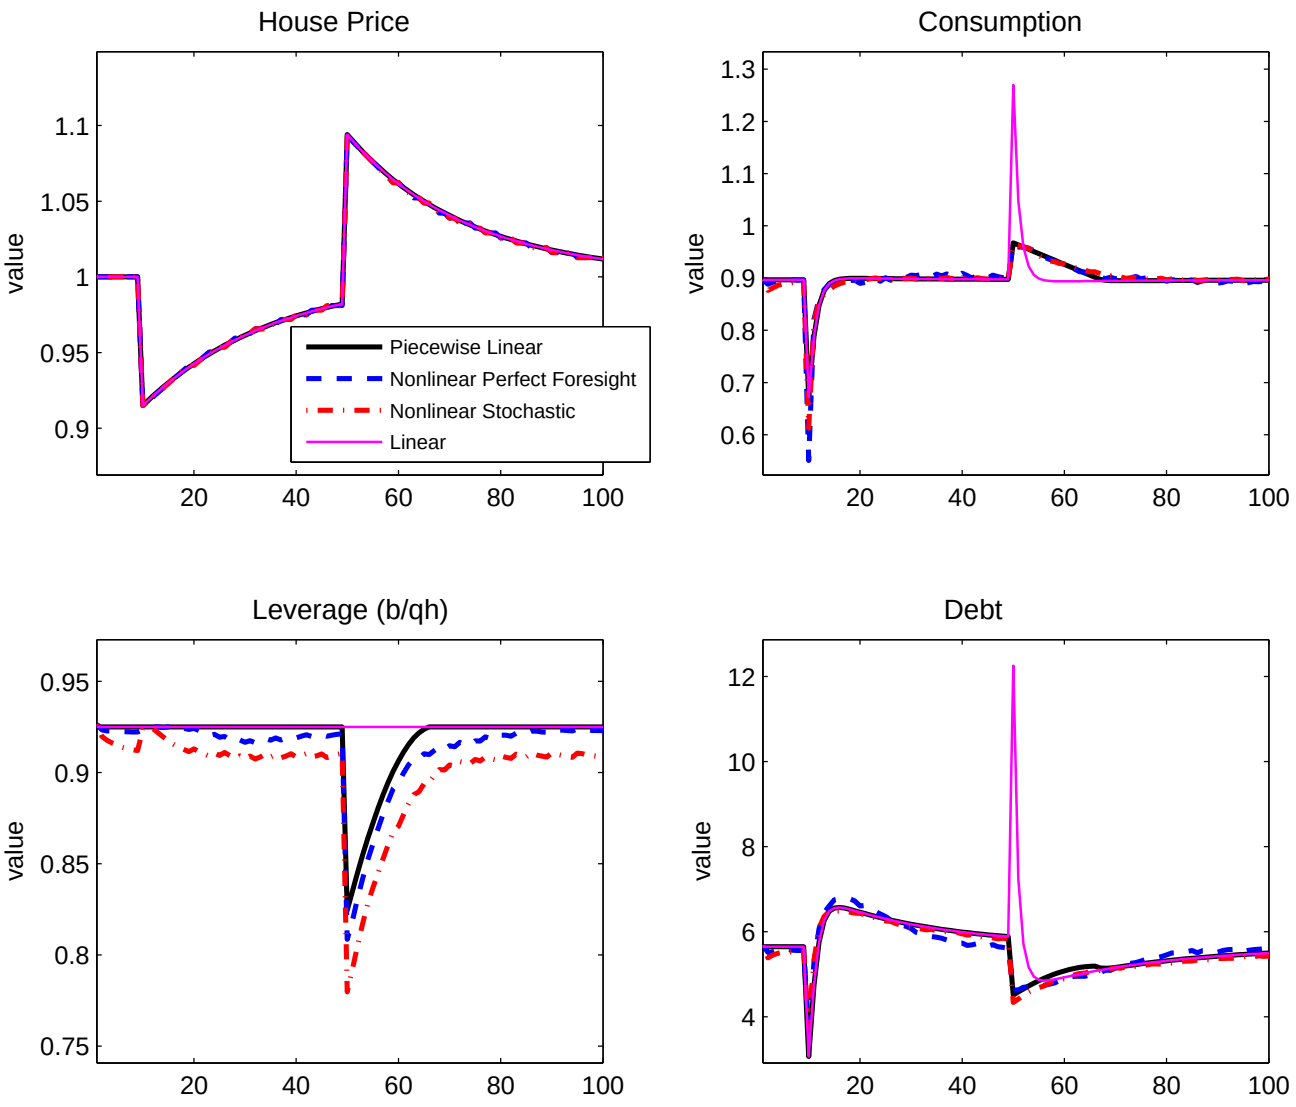

Note: Impulse Responses of the basic model to a negative house price shock in period 10 and a positive house price shock in period 50.

Figure A.2: Accuracy of Solution Method: Simulated Time Series for the Basic Model

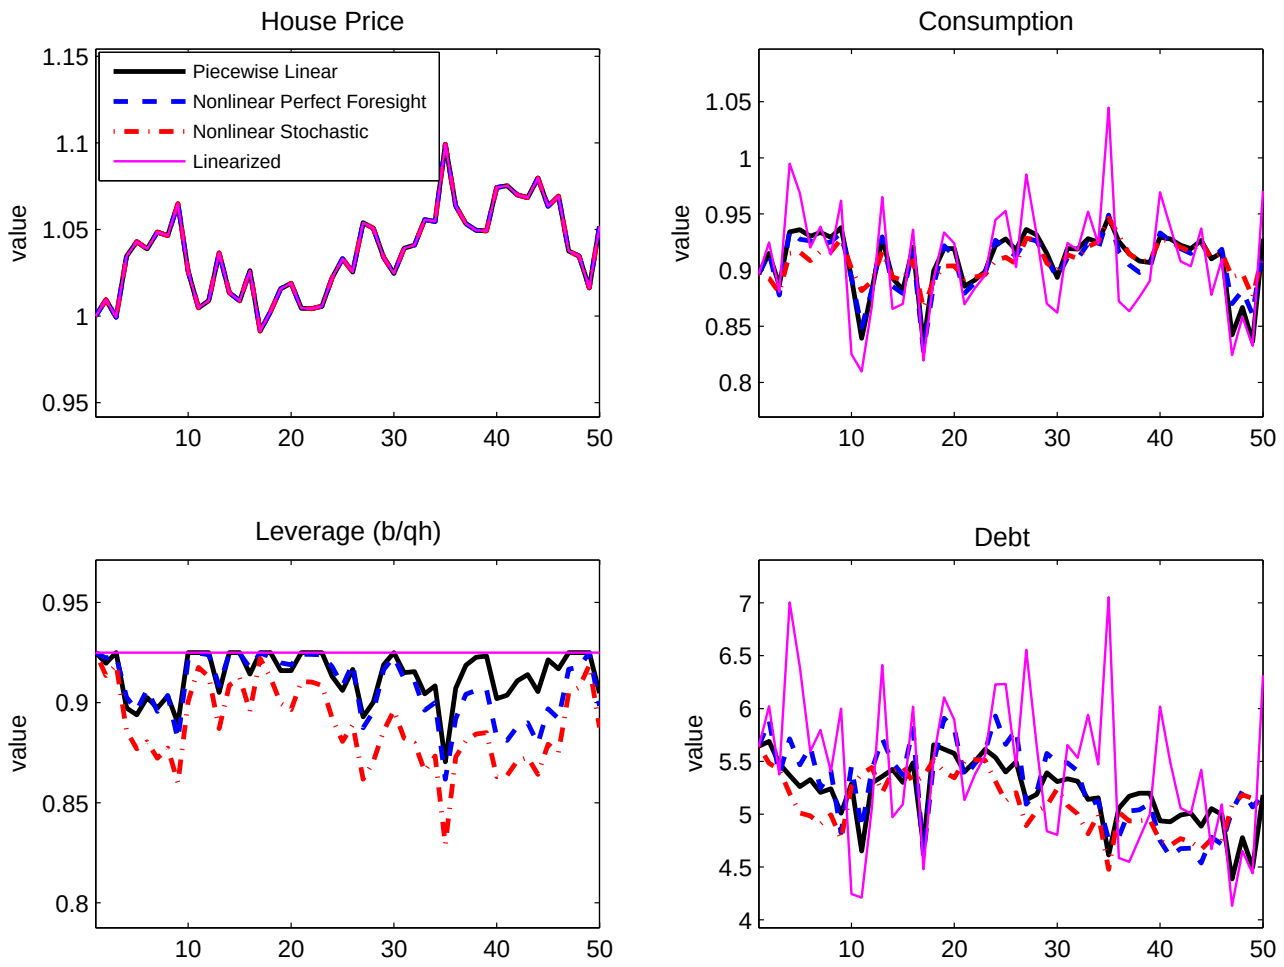

Note: Simulation of the basic model for 50 periods using identical realizations for the exogenous random shock to house prices.

Figure A.3: Local Linearity of the Policy Functions

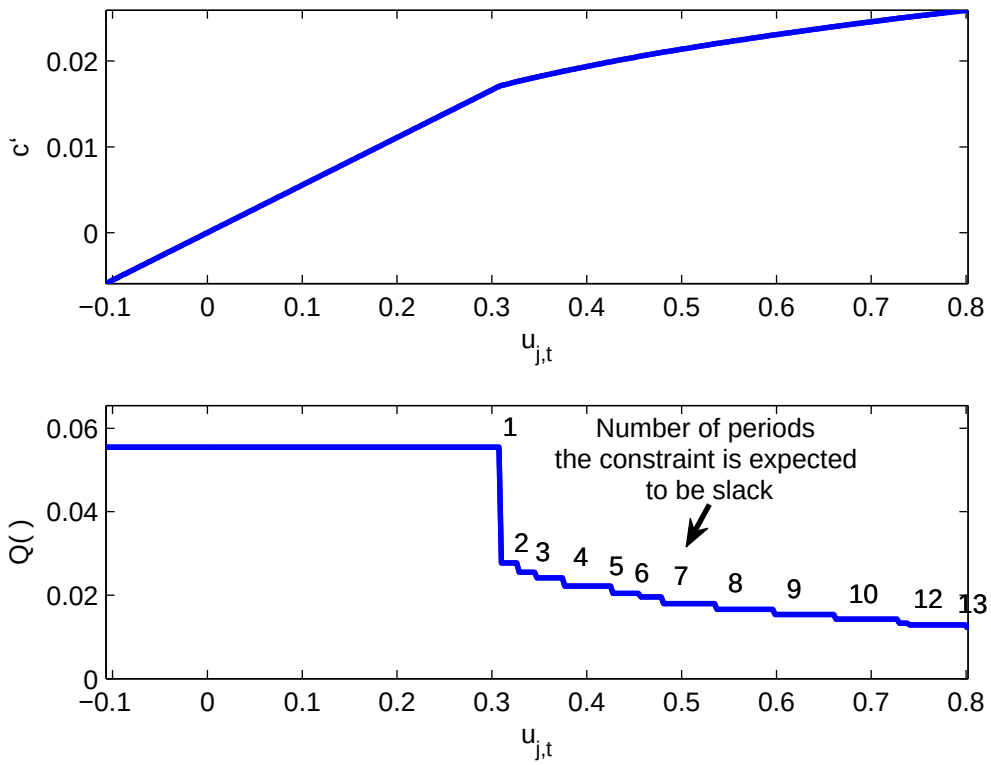

Note: The top panel plots consumption of the impatient agent (in deviation from the steady state) as a function of various realizations of the housing preference shock. The bottom panel plots the slope of the consumption function. The consumption function has a kink when the borrowing constraint becomes binding, and becomes flatter the larger the realization of the housing preference shock.

Table A.1: Accuracy of the Solution Method

Solution Method	Log Consumption		Correlations		$\frac{b}{qh}$	Δ Welfare
	st.dev	skewness	$\log q, \log c$	$\log q, \frac{b}{qh}$	mean	
Linear	6.31%	-0.03	0.40	0.00	0.925	0.199%
Piecewise Linear	4.55%	-1.27	0.55	-0.61	0.908	0.022%
Nonlinear Perfect Foresight	4.54%	-1.18	0.52	-0.64	0.909	0.021%
Nonlinear Stochastic	3.60%	-1.36	0.66	-0.75	0.890	—

Note: Selected properties of the basic model using different solution algorithms. These properties are based on the outcomes of a simulation of 5,000 observations using identical realizations for the exogenous random shocks.

The column labeled "Δ Welfare" indicates the annuity value of the transfer τ (as a percent of current consumption) that would make an agent using the solution method in the first column indifferent between using that method and using the Nonlinear Stochastic solution. Letting (c_t^*, h_t^*) denote the consumption and housing policy in the nonlinear stochastic case, and $\left(\widetilde{c}_t, \widetilde{h}_t\right)$ the consumption policy in the linear case, the two associated value functions are respectively

$$\begin{aligned} W_t^* &= u\left(c_t^*, h_t^*\right) + \beta E_t W_{t+1}^* \\ \widetilde{W}_t &= u\left(\widetilde{c}_t, \widetilde{h}_t\right) + \beta E_t \widetilde{W}_{t+1}. \end{aligned}$$

The transfer τ is the solution to the following equation:

$$\underbrace{u\left(\widetilde{c}_t\left(1+\tau\right), \widetilde{h}\right) + \beta E_t\left(\widetilde{W}_{t+1}\right)}_{\text{value of using suboptimal policy, after transfer}} = \underbrace{W_t^*}_{\text{value of using optimal policy}}$$

By design, the nonlinear stochastic solution attains the highest level of welfare. Note that the linear and piecewise linear solution method could lead to spurious welfare reversals since they linearize the constraints of the original nonlinear problem thus transforming the original problem. To avoid this problem, we use these methods only to compute the borrowing and housing policy, and then obtain the consumption policy c nonlinearly from the budget constraint.

Table A.2: State-Level Regressions: Mortgage Originations and House Prices

	% Change in Mortgage Originations ($\Delta mori_t$)				
Δhp_{t-1}	**1.10*** **				
	(0.18)				
Δhp_high_{t-1}		**-0.41***	**1.08*** **	**1.46*** **	**1.54*** **
		(0.24)	(0.16)	(0.21)	(0.33)
Δhp_low_{t-1}		**3.13*** **	**1.85*** **	**2.53*** **	**2.67** **
		(0.59)	(0.68)	(0.90)	(1.11)
$\Delta mori_{t-1}$				**-0.20*** **	**-0.20*** **
				(0.02)	(0.02)
$\Delta income_{t-1}$					**-0.63**
					(1.04)
pval difference		**0.000**	**0.211**	**0.160**	**0.181**
Time effects	**no**	**no**	**yes**	**yes**	**yes**
Observations	1020	1020	1020	969	969
States	51	51	51	51	51
R-squared	0.01	0.03	0.58	0.53	0.53

Note: State–level Regressions using annual observations from 1992 to 2011 on 50 States and the District of Columbia. Robust standard errors in parenthesis. ***,**,*: Coefficients statistically different from zero at 1, 5 and 10% confidence level, respectively. pval is the p-value of the test for difference in the coefficients for low-house prices and high-house prices.

Data Sources and Definitions: $\Delta mori$ is the percent change in "Mortgage originations and purchases: Value" from the U.S. Federal Financial Institutions Examination Council: Home Mortgage Disclosure Act. See Table 2 for other variable definitions.